

ROUTE

Implementing Cisco IP Routing

Volume 3

Version 1.0

Student Guide

Text Part Number: 97-2816-01



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

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

CCDE, CCENT, CCSI, Cisco Eee, Cisco HealthPresence, Cisco IronPort, the Cisco logo, Cisco Nurse Connect, Cisco Pulse, Cisco SensorBase, Cisco StackPower, Cisco StadiumVision, Cisco TelePresence, Cisco Unified Computing System, Cisco WebEx, DCE, Flip Channels, Flip for Good, Flip Mino, Flipshare (Design), Flip Ultra, Flip Video, Flip Video (Design), Instant Broadband, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn, Cisco Capital, Cisco Capital (Design), Cisco.Financed (Stylized), Cisco Store, Flip Gift Card, and One Million Acts of Green are service marks; and Access Registrar, Aironet, AllTouch, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIW, CCNA, CCNP, CCSP, CCOVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Lumina, Cisco Nexus, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, Continuum, EtherFast, EtherSwitch, Event Center, Explorer, Follow Me Browsing, GainMaker, iLync, IOS, iPhone, IronPort, the IronPort logo, Laser Link, LightStream, Linksys, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, PCNow, PIX, PowerKEY, PowerPanels, PowerTV, PowerTV (Design), PowerVu, Prisma, ProConnect, ROSA, SenderBase, SMARTnet, Spectrum Expert, StackWise, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0910R)

DISCLAIMER WARRANTY: THIS CONTENT IS BEING PROVIDED "AS IS." CISCO MAKES AND YOU RECEIVE NO WARRANTIES IN CONNECTION WITH THE CONTENT PROVIDED HEREUNDER, EXPRESS, IMPLIED, STATUTORY OR IN ANY OTHER PROVISION OF THIS CONTENT OR COMMUNICATION BETWEEN CISCO AND YOU. CISCO SPECIFICALLY DISCLAIMS ALL IMPLIED WARRANTIES, INCLUDING WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT AND FITNESS FOR A PARTICULAR PURPOSE, OR ARISING FROM A COURSE OF DEALING, USAGE OR TRADE PRACTICE. This learning product may contain early release content, and while Cisco believes it to be accurate, it falls subject to the disclaimer above.

Table of Contents

Volume 3

<i>Implementing Path Control</i>	5-1
Overview	5-1
Module Objectives	5-1
<i>Assessing Path Control Network Performance Issues</i>	5-3
Overview	5-3
Objectives	5-3
Assessing Path Control and Network Performance	5-4
Using Filters to Determine Path Selection	5-6
Using PBR to Determine Path Selection	5-13
Configuring and Verifying PBR	5-15
Configuring and Verifying PBR Operations on a Cisco Router	5-18
Summary	5-32
<i>Lab 5-1 Debrief</i>	5-33
Overview	5-33
Objectives	5-33
Lab Overview and Verification	5-34
Sample Solution and Alternatives	5-38
Summary	5-41
<i>References to Additional Path Control in E-Learning</i>	5-43
Overview	5-43
Objectives	5-43
Preview of E-Learning on Implementing Path Control	5-44
Summary	5-48
Module Summary	5-49
Module Self-Check	5-51
Module Self-Check Answer Key	5-54
<i>Connect an Enterprise Network to an ISP Network</i>	6-1
Overview	6-1
Module Objectives	6-1
<i>Planning the Enterprise-to-ISP Connection</i>	6-3
Overview	6-3
Objectives	6-3
Connecting Enterprise Networks to an ISP	6-4
Exchanging Routing Updates with an ISP	6-6
Defining the Types of Connections to an ISP	6-11
Summary	6-16
<i>Considering the Advantages of Using BGP</i>	6-17
Overview	6-17
Objectives	6-17
Using BGP to Connect to an ISP	6-18
BGP Multihoming Options	6-20
BGP Routing Between Autonomous Systems	6-26
Comparison with IGPs	6-28
Path Vector Functionality	6-29
Features of BGP	6-32
Summary	6-39

Comparing the Functions and Uses of EBGP and IBGP	6-41
Overview	6-41
Objectives	6-41
BGP Neighbor Relationships	6-42
Establishing External BGP Neighbor Relationships	6-43
Establishing Internal BGP Neighbor Relationships	6-45
Summary	6-46
Configuring and Verifying Basic BGP Operations	6-47
Overview	6-47
Objectives	6-47
Specifications for Implementing BGP	6-48
Establishing Internal and External BGP Neighbors	6-50
Shutting Down a BGP Neighbor	6-58
BGP Configuration Considerations	6-59
Identifying BGP Neighbor States	6-65
BGP Authentication	6-72
Example of Activating Basic BGP	6-75
BGP Verification	6-77
Summary	6-88
Lab 6-1 Debrief	6-89
Overview	6-89
Objectives	6-89
Lab Overview and Verification	6-90
Sample Solutions and Alternatives	6-94
Summary	6-97
Using the BGP Attributes and Path Selection Process	6-99
Overview	6-99
Objectives	6-99
BGP Path Selection	6-100
Path Selection with Multihomed Connection	6-104
Characteristics of BGP Attributes for Path Selection and Path Manipulation	6-106
Filtering of BGP Routing Updates	6-120
Summary	6-127
Lab 6-2 Debrief	6-129
Overview	6-129
Objectives	6-129
Lab Overview and Verification	6-130
Sample Solutions and Alternatives	6-134
Summary	6-137
E-Learning Training on IPv6 and Routing for Branch Offices and Remote Workers	6-139
Overview	6-139
Objectives	6-139
Preview of E-Learning Modules “Implementing IPv6” and “Routing Facilities for Branch Offices and Mobile Workers”	6-140
Summary	6-146
Module Summary	6-147
References	6-147
Module Self-Check	6-149
Module Self-Check Answer Key	6-157

Implementing Path Control

Overview

If you implement redundancy in a network, then multiple paths exist and the network administrator can manipulate or control the path across the network. Path control is not the same as quality of service (QoS) or MPLS traffic engineering (MPLS TE). Path control is a way to change the default packet forwarding across the network. This module explains why it is necessary to manipulate routing information in order to obtain better resiliency, performance, and availability.

Path control involves the use of a collection of tools, or set of commands, to manipulate the routing protocol forwarding table and bypass the default packet forwarding. This module will provide an overview of these tools, configuration examples of policy-based routing, and performance issues related to path control.

Module Objectives

Upon completing this module, you will be able to evaluate common network performance issues and identify the tools needed to provide Layer 3 path control that uses Cisco IOS features to control the path. This ability includes being able to meet these objectives:

- Evaluate common network performance issues caused by inefficient path control selection and; also identify the resources necessary to develop a Layer 3 solution that uses Cisco IOS features to control path selection, given a set of network requirements, and implement Path Control between multiple IP routing protocols.
- Discuss the lab results for configuring and verifying path control between multiple IP routing protocols.
- Preview the e-learning products to teach the use of path control through additional methods.
- Preview the available e-learning training on implementing route control by other methods.

Lesson 1

Assessing Path Control Network Performance Issues

Overview

This lesson evaluates common network performance factors such as redundancy, resiliency, performance, availability, and adaptability. Issues related to these factors occur when control of path selection is inefficient. Several solutions exist to avoid these issues. This lesson identifies what resources are needed to provide a Layer 3 solution that uses Cisco IOS features to control path selection under a given set of network requirements.

This lesson discusses the use of policy-based routing (PBR), filters (prefix lists, distribute lists, offset lists, and so on), Cisco IOS IP SLAs, and route tags as tools for implementing path control in large networks.

Policy-based routing is described in detail, because it allows administrators to route traffic along specific paths according to their needs. PBR gives network designers greater flexibility in determining traffic patterns and best routes. Sometimes, simple destination-based routing is not sufficient. In these cases, network designers may route packets based on source address, protocol type, or application type so that they can optimally shape traffic patterns. This lesson discusses how to configure PBR on a Cisco router.

Objectives

Upon completing this lesson, you will be able to explain common network performance issues caused by inefficient path control selection. You will also be able to identify the resources you need to best provide path control, as well as implementation and verification of PBR. This ability includes being able to meet these objectives:

- Assess path control network performance.
- Use filters to determine path selection.
- Use PBR to determine path selection.
- Configure and verify PBR.
- Configure and verify PBR operations on a Cisco router.

Assessing Path Control and Network Performance

This topic describes how to assess path control network performance.

Assessing Path Control Network Performance

- Multiple paths
 - Suboptimal routing
 - Availability of redundant paths
 - Application-specific traffic path optimization
 - Load sharing

```
graph LR; R1 --- R2; R2 --- R3; R2 --- R4; R3 --- R4; R3 --- R5; R3 --- R6; R4 --- R5; R5 --- R6; R6 --- DMZ_100;
```

Networks are designed to provide high availability and redundancy. The use of different routing protocols and different connectivity options may result in inefficient paths for forwarding packets to their destinations. Not only is redundancy implemented, but each routing protocol has a different administrative distance, metric, and convergence time.

When more than one routing protocol is implemented, there is high probability of inefficient routing. One major mistake is to implement incorrectly configured two-way multipoint redistribution. Two-way redistribution requires careful planning, not to mention multipoint redistribution, which is highly problematic.

Convergence is also important. First, protocols converge in different ways from each other and in different ways for different network designs. Second, slow convergence can result in an application sending traffic time-outs before a backup path is found to a destination. Path control is required to avoid performance issues and to optimize paths.

Backup is a high priority in modern networks. However, having redundancy does not guarantee resiliency. For redundancy to be effective, you must configure the network properly. This is true even if you are only a fine-tuning a routing protocol or managing path-by-path control tools.

Network engineers should also distinguish between switching all of the traffic to the backup link if there is a primary link failure, and switching some traffic to back up the link if the primary link is congested. Path control mechanisms can improve performance in such a situation as well. Similarly load balancing can divide traffic among parallel paths.

Suboptimal routing can occur after redistribution or simply because different routing protocols use different metrics. Redistribution between two routing protocols, if done based on administrative distance, is likely to result in a suboptimal path. A routing protocol with a lower administrative distance likes its own routes more than routes coming from another routing protocol with a higher administrative distance value. The situation is similar for metrics. Redistribution resets the metric, and the original metric might not be translated correctly. You can use path control tools to change the default destination forwarding and optimize the path of the packets.

There is no easy way to solve all of the issues with path control. It is important to establish a global strategy for path control that provides predictable and deterministic control over traffic patterns, taking into account the physical connectivity as well as the services running over the network infrastructure.

Using Filters to Determine Path Selection

This topic describes how to adjust the path selection process using filters.

Path Selection Process Using Filters

- Manipulating path control by manipulating routing protocols and the routing table
- Tool availability is protocol-dependent:
 - Route maps..... ✓
 - Prefix lists ✓
 - Distribute lists..... ✓
 - Administrative distance ... ✓
 - Route tagging ✓
 - Offset lists
 - Cisco IOS IP SLA
 - PBR

✓ – Previously covered
 – Not covered yet

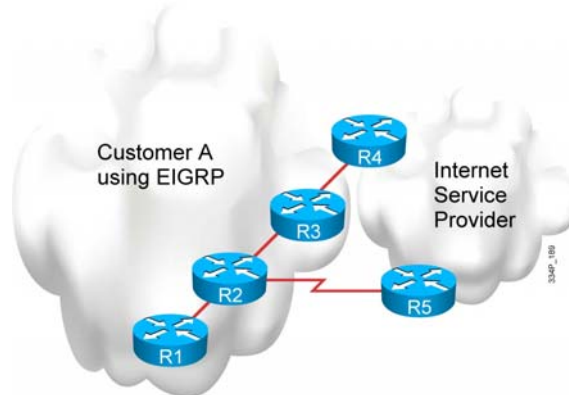
Paths are selected using the IP routing table. The contents of the IP routing table are provided by either a routing protocol or static route entries. The forwarding engine uses these contents to switch the packets from the incoming interface to the outgoing interface. You use filters to manipulate the routing protocols and routing table, in order to influence the desired path.

Several filters or tools can be used to manipulate the routing table:

- Route maps
- Prefix lists
- Distribute lists
- Administrative distance
- Route tagging
- Offset lists
- Cisco IOS IP SLA
- PBR

Path Control Tools: Offset List

- Routers R4 and R5 receive a subset of routes from the ISP.
- The link between R2 and R5 is slow.
- How do you make the path toward R4 the primary way out of the EIGRP network for a set of destinations?



In the figure in the slide, the customer is using RIP and is connected to the Internet Service Provider (ISP) via edge routers R4 and R5. A subset of routes is received from each of the edge routers, and RIP increases the cost at every new hop. The cost between routers R2 and R5 is smaller than that between R4 and R5, because there is only one hop, even though this is very slow link. What configuration should be applied to router R2 in order to prefer the path toward edge router R4, which is must faster? The configuration should take into account that this rule is only valid for a set of destinations.

Path Control Tools: Offset List (Cont.)

- An offset value to incoming and outgoing metrics to routes learned is added
- Supported protocols:
 - EIGRP
 - RIP

R2 (config-route-map) #

```
offset-list 21 in 2 serial 0/0
```

- The router applies an offset of 2 to routes learned from the serial 0/0 interface with the match on access list 21.

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE-01-005-08

In order to configure a route to avoid a slow link on router R2, use an offset list on the router..

The **offset-list** command adds an offset to incoming and outgoing metrics to routes learned via Enhanced Interior Gateway Routing Protocol (EIGRP) or Routing Information Protocol (RIP). The offset value is added to the routing metric. An offset list with an interface type and interface number is considered extended and takes precedence over an offset list that is not extended. Therefore, if an entry passes an extended offset list and a normal offset list, the offset of the extended offset list is added to the metric.

In the figure in the slide, the **offset-list 21 in 2 serial 0/0** command is used to apply an offset of 2 to routes learned from serial interface 0/0 with the match on access list 21. The command in the slide is entered into the RIP process configuration mode on router R2. With this command, the cost of a selection of routes coming from the Serial 0/0 interface is increased, effectively making it worse than the path toward router R4. Therefore, for a selection of routes, router R4 becomes the way out from the customer network.

For more details, about the **offset-list** (EIGRP) and **offset-list** (RIP) commands, please check the Cisco IOS IP Routing Protocols Command Reference via the following link:

http://www.cisco.com/en/US/docs/ios/proute/command/reference/irp_book.html

Path Control Tools: Cisco IOS IP Service Level Agreement

- End-to-end network performance tests based on clear measurement metrics
- Can be used for path control
- Configuration:
 - Define one or more probes
 - Define one or more tracking objects
 - Define the action on the tracking object

Cisco IOS IP Service Level Agreements (SLAs) perform network performance measurement within Cisco Systems devices. They use active traffic monitoring (the generation of traffic in a continuous, reliable, and predictable manner) for measuring network performance. Cisco IOS IP SLAs actively send data across the network to measure performance between multiple network locations or across multiple network paths. They use timestamp information to calculate performance metrics, such as jitter, latency, network and server response times, packet loss, and mean opinion score.


The following steps are required to configure Cisco IOS IP SLA functionality:

- Define one or more probes
- Define one or more tracking objects
- Define the action for each tracking object

Note E-learning training is available for path control topics that provide more information about Cisco IOS IP SLAs.

Cisco IOS IP SLA Example

- Customer A is multihoming to ISP-1 and ISP-2
- The link to ISP-1 is the primary link for all traffic
- Customer A is using the default routes to the ISPs
- A Cisco IOS IP SLA is used to conditionally announce the default route



The following example describes the use of Cisco IOS IP SLA functionality for path control.

Customer A is multihoming to two service providers and is not using BGP as the routing protocol. Default routes are configured instead.

In order to make one link a primary link and the second one a backup link, two static routes with different administrative distances can be configured. The static default route with a lower administrative distance will be preferred and inserted into the IP routing table. However, it is possible that the interface is up and the ISP has problems with the provider edge (PE) router or with the uplink connectivity toward the Internet. In such a case, the active default route sends traffic to the selected ISP, but all of the traffic is lost. The solution to this issue is to use Cisco IOS IP SLA functionality, which can be used to continuously check the reachability of a specific destination (PE interface, ISP, DNS server, or any other trusted destination) and conditionally announce the default route if the connectivity is verified.

The configuration required is shown in the next figure.

Cisco IOS IP SLA Example

R1(config)#

```
ip sla monitor 11
  type echo protocol ipIcmpEcho 10.1.1.1 source-interface FastEthernet0/0
  frequency 10
ip sla monitor schedule 11 life forever start-time now
```

- Sets the probe to send an ICMP packet every 10 seconds to IP address 10.1.1.1
- Starts sending packets now and continues forever

R1(config)#

```
track 1 ip sla 11 reachability
```

- Defines the tracking of object 1 linked to ip sla 11

R1(config)#

```
ip route 0.0.0.0 0.0.0.0 10.1.1.1 2 track 1
```

- Announces the default route to 10.1.1.1 with an administrative distance of 2 if tracking object 1 is true

In the figure in the slide, three steps to configure Cisco IOS IP SLA functionality are shown.

The first step is needed to define the probe. The probe number 11 is defined by issuing the **ip sla monitor 11** command. The test defined with the **type echo protocol ipIcmpEcho 10.1.1.1 source-interface FastEthernet0/0** command specifies to send the Internet Control Message Protocol (ICMP) echoes to destination 10.1.1.1 to check the connectivity. The FastEthernet0/0 interface is used as a source. The **frequency 10** command schedules the connectivity test to repeat every 10 seconds and, finally, the **ip sla monitor schedule 11 life forever start-time now** command defines the start and end times of the connectivity test for probe 11. The start time is now and the end time is forever.

The second step requires the tracking object, which is linked to the probe from step one. The **track 1 ip sla 11 reachability** command defines the tracking of object 1, which is linked to probe 11, and the reachability of the 10.1.1.1 is tracked.

The last step is an action based on the status of the tracking object. The **ip route 0.0.0.0 0.0.0.0 10.1.1.1 2 track 1** command conditionally announces the default route to 10.1.1.1 with an administrative distance of 2 if the result of tracking object 1 is true.

For more details about the **Cisco IOS IP SLA** commands, please check the Cisco IOS IP SLAs Command Reference via the following link:

http://www.cisco.com/en/US/docs/ios/ipsla/command/reference/sla_01.html

Cisco IOS IP SLA Example (Cont.)

R1(config)#

```
ip sla monitor 22
  type echo protocol ipIcmpEcho 11.1.1.1 source-interface FastEthernet0/1
  frequency 10
ip sla monitor schedule 22 life forever start-time now
```

- Sets the probe to send an ICMP packet every 10 seconds to IP address 11.1.1.1
- Starts sending packets now and continues forever

R1(config)#

```
track 2 ip sla 22 reachability
```

- Defines the tracking of object 1 linked to ip sla 11

R1(config)#

```
ip route 0.0.0.0 0.0.0.0 11.1.1.1 3 track 2
```

- Announces the default route to 11.1.1.1 with an administrative distance of 3 if tracking object 2 is true

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE-0100-55

This is an example in which Customer A is multihoming to two ISPs. It requires the configuration of two probes, two tracking objects, and two conditionally announced default routes.

The configuration in the figure in the slide is almost the same as the configuration in the previous figure. The probe number is different, because the second test condition tests the reachability of the backup ISP destination address. The tracking object number is different, as well, because it is related to the second probe. Finally, the announced default route uses a higher administrative distance (an administrative distance number of 3 is configured in the configuration in the slide) as the backup. The ISP is used only if the primary is not available. Of course, the second tracking object is used as a reference for a conditional announcement for this default route.

Using PBR to Determine Path Selection

This topic describes how policy-based routing is used to implement path control.

Policy-Based Routing

- Allows you to implement policies that selectively cause packets to take different paths.
 - IP routing is destination-based
 - PBR avoids destination-based routing
- Is applied to incoming packets
- Makes traffic marking possibility
- Requires a route map to implement the policy
 - Matched routes are modified by **set** commands

PBR offers significant benefits in terms of implementing user-defined policies to control traffic in the internetwork. It provides solutions in cases where legal, contractual, or political constraints dictate that traffic be routed through specific paths.

PBR adds flexibility in a difficult-to-manage environment by giving the network administrator the ability to route traffic based on network needs. For network managers who implement routing policies in their networks, PBR provides an extremely powerful, simple, and flexible tool.

PBR also provides a mechanism for marking packets. As a result, differentiated preferential service can be provided to different types of traffic in combination with queuing techniques that are available in Cisco IOS software.

Policy-based routing is used to bypass the routing table. It allows the network administrator to configure different routing rules beyond the original IP routing table. One of the ways it can be used is to route packets based on the source IP address instead of the destination IP address. PBR is applied to incoming packets and is implemented using route maps, for which **match** commands are used to match the incoming packets and a subset of the **set** commands are used to change the default destination-based routing.

Policy-Based Routing Benefits

- Source-based transit provider selection
 - Different users go different ways
- QoS
 - Sets the precedence or ToS; used with queuing
- Load sharing
 - Forces load sharing without regard to the routing table
- Cost savings
 - Distributes traffic economically

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE-5-1

These benefits are achieved by implementing PBR in a network:

- **Source-based transit-provider selection:** ISPs and other organizations use PBR to route traffic originating from different sets of users through different Internet connections across the policy routers.
- **Quality of service (QoS):** Organizations provide QoS to differentiated traffic using the following two methods:
 - Set the precedence or type of service (ToS) value in each IP packet header at the periphery of the network.
 - Leverage queuing mechanisms to prioritize traffic in the core or backbone of the network.
- **Cost savings:** Organizations achieve cost savings by distributing interactive and batch traffic among low-bandwidth, low-cost, permanent paths and high-bandwidth, high-cost, switched paths.
- **Load sharing:** In addition to the dynamic load-sharing capabilities that are offered by destination-based routing that is supported by Cisco IOS software, network managers can implement policies to distribute traffic among multiple paths based on the traffic characteristics.

Configuring and Verifying PBR

This topic describes how to configure and verify PBR.

Steps to Implement Path Control


- Choose the path control tool
- Match traffic to manipulate the path
- Define the action for matched traffic
- Apply path control to traffic
 - To incoming traffic
 - To traffic local to the router
- Verify path control results

The figure in the slide shows the steps required to implement path control:

- Choose the path control tool. Path control tools can manipulate the IP routing table or bypass it. You must select the correct tool in order to fulfill the requirements.
- When you select the path control tool, implement a traffic-matching configuration to specify which traffic will be manipulated.
- When traffic is classified and matched, define the action for a specific class of traffic, in order to efficiently control the path of packets.
- After a successful configuration of path control, apply statements to manipulate the packets for incoming traffic or traffic local to the router.
- Verify path control. You can use simple connectivity commands and several **show** commands.

Requirements for Policy Based Routing

- Match packets with the destination network 7.7.7.0 and forward them to the next hop 6.6.6.6
- Match packets between 3 and 200 bytes in size and forward them to the interface Ethernet 0
- Apply the route map to the incoming interfaces
- Verify the configuration



The figure in the slide shows the requirements for PBR implementation. PBR is implemented by using the route maps, for which **match** commands are used to match the traffic and **set** commands are used to set desired action to control the path. In the example in the slide, configuration is required to match packets going to the destination network 7.7.7.0 and forward them to the next hop 6.6.6.6 to avoid suboptimal routing. At the same time, packets with sizes between 3 and 200 bytes must be matched and forwarded to interface Ethernet 0. The route map must be applied to the incoming interfaces. Finally, the configuration must be verified.

Steps to Configure and Verify Policy-Based Routing

- Enable PBR by configuring a route map:
 - Match traffic using the **match** command
 - Define the action for matched traffic using the **set** command
- Enable fast-switched PBR or PBR switched by Cisco Express Forwarding (Optional)
- Apply a route map:
 - To an incoming interface
 - To packets that are generated by the router
- Verify the PBR configuration

The requirements from the previous figure are used to create the implementation plan for PBR configuration. The implementation plan includes the following steps to configure and verify PBR on Cisco routers:

- Enable PBR by configuring the route map. Traffic is matched using the **match** command. Actions for manipulating path control are defined using **set** commands.
- If desired, enable fast-switched PBR or PBR switched by Cisco Express Forwarding. Fast-switched PBR must be enabled manually. For PBR switched by Cisco Express Forwarding, in contrast, is automatically enabled when CEF switching is enabled.
- Apply the route map created to an incoming interface or to locally generated packets in the router.
- Verify PBR configuration with basic connectivity and path verification commands, as well as policy routing **show** commands.

Configuring and Verifying PBR Operations on a Cisco Router

This topic describes how to configure and verify PBR on a Cisco Router.

Matching the Traffic

```
R2 (config) #  
route-map PBRmap permit 10
```

- Configure a route map

```
R2 (config-route-map) #  
match ip address 10
```

- Matches IP addresses for policy routing
- Access list 10 is used to match the IP address

```
R2 (config-route-map) #  
match length 3 200
```

- Matches the Layer 3 length of the packet for policy routing
- Packets between 3 and 200 bytes long are matched

Use the **route-map** command to enable policy routing in the router and to enter the **match** and **set** commands required. IP standard or extended access lists are used to establish the PBR match criteria using the **match ip address** command. The **match ip address 10** command matches traffic based on standard IP access list number 10. A standard IP access list is used to specify the match criteria for the source address of a packet. Extended access lists are used to specify the match criteria based on source and destination addresses, application, protocol type, ToS, and precedence.

The **match length** command matches packets based on packet length, and can be configured with a minimum and maximum value. For example, a network administrator can use the **match length** command to distinguish between interactive traffic and file transfer (bulk) traffic, because file transfer traffic typically has larger packet sizes. The **match length 3 200** command matches packets between 3 and 200 bytes long.

For more details about the **route-map** and **match** commands, please check the Cisco IOS IP Routing Protocols Command Reference via the following link:

http://www.cisco.com/en/US/docs/ios/iproute/command/reference/irp_book.html

Policy Routing set Commands

R2 (config-route-map) #

```
set ip next-hop 6.6.6.6
```

- This command defines where to forward packets that pass a match clause of a route map for policy routing.
- Packets that pass the match clause are forwarded to the router at IP address 6.6.6.6.

R2 (config-route-map) #

```
set interface ethernet 0
```

- This command also defines where to forward packets that pass a match clause of a route map for policy routing.
- Packets that pass the match clause are forwarded to the Ethernet 0 interface.

If the match statements are satisfied, one or more of the following set statements are used to specify the criteria for forwarding packets. The router evaluates the four **set** commands for PBR that are shown in this figure and in the following figure, the order listed.

Once you choose a destination address or interface, other **set** commands for changing the destination address or interface are ignored. Some of these commands affect only packets for which there is an explicit route in the routing table, and others affect only packets for which there is *no* explicit route in the routing table.

A packet that is not affected by any of the **set** commands in a route map statement is not policy-routed; it is forwarded normally. In other words, destination-based routing is performed. The router uses the four **set** commands for PBR in the following order:

1. The **set ip next-hop** command configures a list of IP addresses. The list specifies the adjacent next-hop router in the path toward the destination to which the packets are forwarded. If more than one IP address is specified, then the first IP address that is associated with a currently active interface is used to route the packets.

Note With the **set ip next-hop** command, the routing table is checked only to determine if the next hop is reachable. It is not checked to determine if there is an explicit route for the destination address of the packet.

This **set** command affects all packet types and is always used if it is configured.

The **set ip next-hop 6.6.6.6** command from the figure in the slide shows that packets that pass a match clause of a route map for policy routing are forwarded to 6.6.6.6.

2. The **set interface** command configures a list of interfaces through which the packets are routed. If more than one interface is specified, then the first active interface in the list is used for forwarding the packets.

Note If there is no explicit route in the routing table for the destination network address of the packet (for example, if the packet is a broadcast packet or is destined to an unknown address), the **set interface** command has no effect and is ignored.

A default route in the routing table will not be considered an explicit route for an unknown destination address.

The **set interface ethernet 0** command from the figure in the slide defines that packets that pass a match clause of a route map for policy routing are forwarded to interface Ethernet 0.

For more details about the **set** command, please check the Cisco IOS IP Routing Protocols Command Reference via the following link:

http://www.cisco.com/en/US/docs/ios/iproute/command/reference/irp_book.html

Policy Routing set Commands (Cont.)

R2 (config-route-map) #

```
set ip default next-hop 6.6.6.6
```

- This command defines where to output packets that pass a match clause of a route map for policy routing and for which the Cisco IOS Software has no explicit route to a destination.

R2 (config-route-map) #

```
set default interface ethernet 0
```

- This command defines where to output packets that pass a match clause of a route map for policy routing and have no explicit route to the destination.
- This is recommended only for point-to-point links.

3. The **set ip default next-hop** command configures a list of default next-hop IP addresses if there is no explicit route available to the destination address of the packet being considered for PBR. If more than one IP address is specified, then the first next hop specified that appears to be adjacent to the router is used. The optional specified IP addresses are tried in sequential order.

Note A packet is routed to the next hop specified by the **set ip default next-hop** command if there is no explicit route for the destination address of the packet in the routing table.

A default route in the routing table will not be considered an explicit route for an unknown destination address.

The **set ip default next-hop 6.6.6.6** command from the figure in the slide defines that packets that pass a match clause of a route map for policy routing and for which the Cisco IOS Software has no explicit route are forwarded to 6.6.6.6.

4. The **set default interface** command configures a list of default interfaces. If there is no explicit route available to the destination address of the packet that is being considered for PBR, then it is routed to the first active interface in the list of specified default interfaces.

Note A packet is routed out of the interface that is specified by the **set default interface** command only if there is no explicit route for the destination address of the packet in the routing table.

A default route in the routing table will not be considered an explicit route for an unknown destination address.

The **set default interface ethernet 0** command from the figure in the slide defines that packets that pass a match clause of a route map for policy routing and for which the Cisco IOS software has no explicit route are forwarded to interface Ethernet 0.

PBR also provides a mechanism to mark packets using the **set ip tos** and **set ip precedence** commands:

- The **set ip tos** command is used to set the IP ToS value in the IP packets.
- The **set ip precedence** command is used to set the IP precedence in the IP packets. You must specify either the precedence number or name.

Note The **set** commands can be used in conjunction with each other.

For more details about the **set** command, please check the Cisco IOS IP Routing Protocols Command Reference via the following link:

http://www.cisco.com/en/US/docs/ios/iproute/command/reference/irp_book.html

Apply Route Maps for PBR

R2(config-if)#

```
ip policy route-map PBRmap
```

- This command specifies the route map to use for policy routing on an incoming interface that is receiving packets that need to be policy-routed.

R2(config)#

```
ip local policy route-map PBRmap
```

- Specifies the route map to use for policy routing of all packets that originate on the router.

To identify a route map to use for PBR on an interface, use the **ip policy route-map** interface configuration command. In the figure in the slide, the **ip policy route-map PBRmap** command applies the route map named PBRmap to the incoming traffic on the interface.

Note PBR is specified on the incoming interface that receives the packets that need to be policy-routed, not on the interface from which the packets are sent.

To identify a route map to use for local policy routing, use the **ip local policy route-map** global configuration command. In the figure in the slide, the **ip local policy route-map PBRmap** command applies the route map named PBRmap to packets that originate on the router.

Packets that are generated by the router are not normally policy-routed. However, you can use this command to policy-route such packets. You might enable local-policy routing if you want packets that originate at the router to take a route other than the obvious shortest path.

For more details about the **ip policy route-map** and **ip local policy route-map** commands, please check the Cisco IOS IP Routing Protocols Command Reference via the following link: http://www.cisco.com/en/US/docs/ios/iproute/command/reference/irp_book.htm

Enable Fast-Switched PBR or PBR Switched by Cisco Express Forwarding

- Optional command
- Fast-switched PBR must be enabled manually
- Switching of PBR by Cisco Express Forwarding is enabled automatically

```
R2(config-if)#
```

```
ip route-cache policy
```

- Enables fast-switched policy routing

© 2009 Cisco Systems, Inc. All rights reserved.

EN115110-55-10

PBR must be configured before PBR fast switching can be enabled. Fast-switching of PBR is disabled by default. To configure a fast-switched PBR, use the **ip route-cache policy** command in interface configuration mode.

A fast-switched PBR supports all of the **match** commands and most of the **set** commands, except for these restrictions:

- The **set ip default next-hop** command is not supported.
- The **set interface** command is supported over point-to-point links, unless a route cache entry exists that uses the same interface that is specified in the **set interface** command in the route map. Also, at the process level, usually the routing table is checked to determine if the interface is on an appropriate path to the destination.

During fast switching, the software does not make this check. Instead, if the packet matches, the software automatically forwards the packet to the specified interface.

Note Since the release of Cisco IOS Release 12.0, IP PBR can now be fast-switched. Prior to this feature, PBR was process-switched, which meant that on most platforms, the switching rate was approximately 1000 to 10,000 packets per second (p/s). This rate was not fast enough for many applications.

Users who require PBR to occur at faster speeds can now implement it without slowing down the router.

Beginning with Cisco IOS Release 12.0, PBR is supported in the Cisco Express Forwarding switching path. PBR switched by Cisco Express Forwarding has better performance than fast-switched PBR and, therefore, is the optimal way to perform PBR on a router. No special configuration is required to enable PBR to be switched by Cisco Express Forwarding. It is on by default as soon as you enable Cisco Express Forwarding and PBR on the router.

Note The **ip route-cache policy** command is strictly for fast-switched PBR; it is not required for a PBR switched by Cisco Express Forwarding.

For more details about the **ip route-cache policy** command, please check the Cisco IOS IP Routing Protocols Command Reference via the following link:

http://www.cisco.com/en/US/docs/ios/iproute/command/reference/irp_book.htm

Verifying Policy-Based Routing

R1#

```
show ip policy
```

- Displays route maps that are configured on the interfaces.

R1#

```
show route-map [map-name]
```

- Displays a route map.

© 2009 Cisco Systems, Inc. All rights reserved.

5011-10-5-01

To display the route maps used for PBR on the interfaces of the router, use the **show ip policy** EXEC command.

To display configured route maps, use the **show route-map** EXEC command.

For more details about the **show ip policy** and **show route-map** commands, please check the Cisco IOS IP Routing Protocols Command Reference via the following link:

http://www.cisco.com/en/US/docs/ios/iproute/command/reference/irp_book.html

Verifying Policy-Based Routing (Cont.)

R1#

```
debug ip policy
```

- Enables the display of IP policy routing events.

R1#

```
traceroute
```

- The extended **traceroute** command allows for the specification of the source address.

R1#

```
ping
```

- The extended **ping** allows for the specification of the source address.

Use the **debug ip policy** EXEC command to display the IP PBR packet activity. This command shows, in detail, the activities that the PBR is performing. It also displays information that indicates whether a packet matches the criteria. If the criteria match, the resulting routing information for the packet is displayed as well.

Note Because the **debug ip policy** command generates a significant amount of output, use it only when traffic on the IP network is low, so that other activity on the system is not adversely affected.

To discover the routes that packets follow when traveling to their destination from the router, use the **traceroute** privileged EXEC command. To change the default parameters and invoke an extended traceroute test, enter the command without a destination argument. You will be guided through a dialog to select the required parameters.

To check host reachability and network connectivity, use the **ping** privileged EXEC command. You can use the extended command mode of the **ping** command to specify the supported header options by entering the command without any arguments.

For more details about the **debug ip policy** command, please check the Cisco IOS Debug Command Reference via the following link:

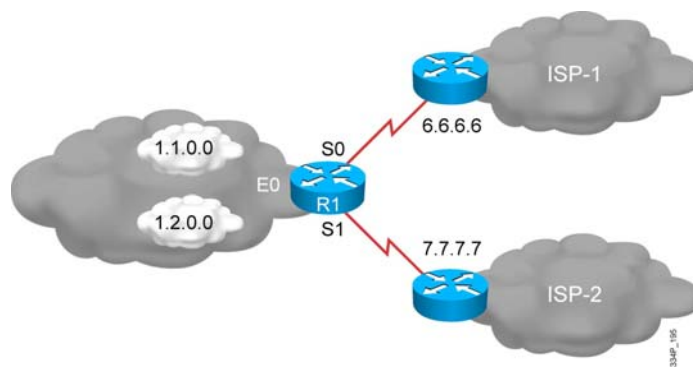
http://www.cisco.com/en/US/docs/ios/debug/command/reference/db_book.html

For more details about the **ping** and **traceroute** commands, please check the Cisco IOS Configuration Fundamentals Command Reference via the following link:

http://www.cisco.com/en/US/docs/ios/fundamentals/command/reference/cf_book.html

Example: PBR Equal Access

- All traffic that uses a default route and is sourced from subnet 1.1.0.0 should go through ISP 1.
- All traffic that uses a default route and is sourced from subnet 1.2.0.0, should go through ISP 2.



The example in the figure in the slide describes a common scenario in which a private company that is attached to more than one ISP must build a traffic policy.

In this figure, router R1 provides Internet access for a private enterprise. Router R1 is connected to two different ISPs. This router advertises a 0.0.0.0 default route into the enterprise network to avoid a large Internet routing table.

The problem is that when traffic from the enterprise networks 1.1.0.0 and 1.2.0.0 reaches router R1, the traffic can go to ISP 1 or ISP 2. The company prefers to have ISP 1 and ISP 2 receive approximately equal amounts of traffic. PBR is used to shape or load-balance traffic from router R1 to each of the ISPs.

PBR is implemented at router R1. All traffic sourced from the 1.1.0.0 subnet is forwarded to ISP 1 if there is no specific route to the destination in the routing table (the default route is not used). All traffic sourced from the 1.2.0.0 subnet is forwarded to ISP 2 if there is no specific route to the destination in the routing table.

Note Remember, this policy provides for an outbound traffic policy from the enterprise to its ISPs only. It does not determine the inbound traffic policy for router R1. It is possible that traffic from 1.1.0.0 going out to ISP 1 will receive responses from ISP 2.

Example: PBR Equal Access (Cont.)

```
R1 (config) # access-list 1 permit 1.1.0.0 0.0.255.255
R1 (config) # access-list 2 permit 1.2.0.0 0.0.255.255

R1 (config) # route-map equal-access permit 10
R1 (config-route-map) # match ip address 1
R1 (config-route-map) # set ip default next-hop 6.6.6.6
R1 (config-route-map) # route-map equal-access permit 20
R1 (config-route-map) # match ip address 2
R1 (config-route-map) # set ip default next-hop 7.7.7.7
R1 (config-route-map) # route-map equal-access permit 30
R1 (config-route-map) # set default interface null0

R1 (config) # interface ethernet 0
R1 (config-if) # ip address 1.1.1.1 255.255.255.0
R1 (config-if) # ip policy route-map equal-access

R1 (config) # interface serial 0
R1 (config-if) # ip address 6.6.6.5 255.255.255.0

R1 (config) # interface serial 1
R1 (config-if) # ip address 7.7.7.6 255.255.255.0
```

The configuration that is shown in the figure in the slide has been set in router R1. The name of the route map is “equal-access.”

The **ip policy route-map equal-access** command has been applied to the Ethernet 0 interface, which is the incoming interface that receives the packets to be policy-routed.

Sequence number 10 in the route map equal-access is used to match all packets that are sourced from any host in subnet 1.1.0.0. If a packet matches, and if the router has no explicit route for its destination, then the packet is sent to the next-hop address 6.6.6.6 (ISP 1 router).

Sequence number 20 in the route map equal-access is used to match all packets that are sourced from any host in subnet 1.2.0.0. If a packet matches, and if the router has no explicit route for its destination, then the packet is sent to the next-hop address 7.7.7.7 (ISP 2 router).

Sequence number 30 in the route map equal-access is used to drop all traffic that is not sourced from subnet 1.1.0.0 or 1.2.0.0. The null0 interface is a route to nowhere (thus, the traffic is dropped).

Verifying Policy-Based Routing: Examples

```
R1#show ip policy
Interface      Route map
Ethernet0      equal-access

R1#show route-map
route-map equal-access, permit, sequence 10
  Match clauses:
    ip address (access-lists): 1
  Set clauses:
    ip default next-hop 6.6.6.6
Policy routing matches: 3 packets, 168 bytes
route-map equal-access, permit, sequence 20
  Match clauses:
    ip address (access-lists): 2
  Set clauses:
    ip default next-hop 7.7.7.7
route-map equal-access, permit, sequence 30
Set clauses:
  default interface null0
```

© 2009 Cisco Systems, Inc. All rights reserved.

5-30 (REV 10-5-09)

In the figure in the slide, the output provides examples of two **show** commands issued on router R1. The **show ip policy** command output indicates that the route map called equal-access is used for PBR on the Ethernet0 interface of the router. The **show route-map** command output indicates that three packets have matched sequence 10 of the route map equal-access.

Verifying Policy-Based Routing: Examples (Cont.)

```
R1# debug ip policy
Policy routing debugging is on

11:51:25: IP: s=1.1.1.1 (Ethernet0), d=190.168.1.1, len 100,
policy match
11:51:25: IP: route map equal-access, item 10, permit
11:51:25: IP: s=1.1.1.1 (Ethernet0), d=190.168.1.1
(Serial0), len 100, policy routed
11:51:25: IP: Ethernet0 to Serial0 6.6.6.6
```

In the figure in the slide, the output provides an example of the **debug ip policy** command output. The **show logging** command shows the logging buffer, including the output of the **debug** command.

The output indicates that a packet from 1.1.1.1 destined for 190.168.1.1 has been received on interface Ethernet0 and that it is policy-routed on Serial0 to next hop 6.6.6.6. The source address of 1.1.1.1 matches line 10 of the route map equal-access.

Summary

This topic summarizes the key points that were discussed in this lesson.

Summary

- Redundant paths (multiple paths), redistribution, and the selected routing protocol all affect network performance. Path control must be enabled to improve performance and avoid suboptimal routing.
- A route map with a group of **match** and **set** commands is one of the tools that can be used for path control.
- The path selection process can be accomplished using filters such as: route tagging, prefix lists, distribute lists, administrative distance, offset lists, and Cisco IOS IP SLAs.
- In order to bypass the routing table destination-based forwarding, policy-based routing is used to determine path selection.

Summary (Cont.)

- Policy-based routing uses route maps for configuration. It can be fast-switched or switched by Cisco Express Forwarding.
- Path control **match** commands match incoming traffic. Path control **set** commands manipulate the path; manipulation can be applied to incoming traffic or to traffic generated by the router.

Lesson 2

Lab 5-1 Debrief

Overview

In Lab 5-1, students manipulate the path for a selection of the packets traveling across the network. First they configure the routing protocol to assure reachability and to establish the base for destination-based forwarding. Next they implement path control to manipulate the path for a set of packets.

Students define several groups of packets and use policy-based routing (PBR) to manipulate the path across the network. Finally, they verify the configuration.

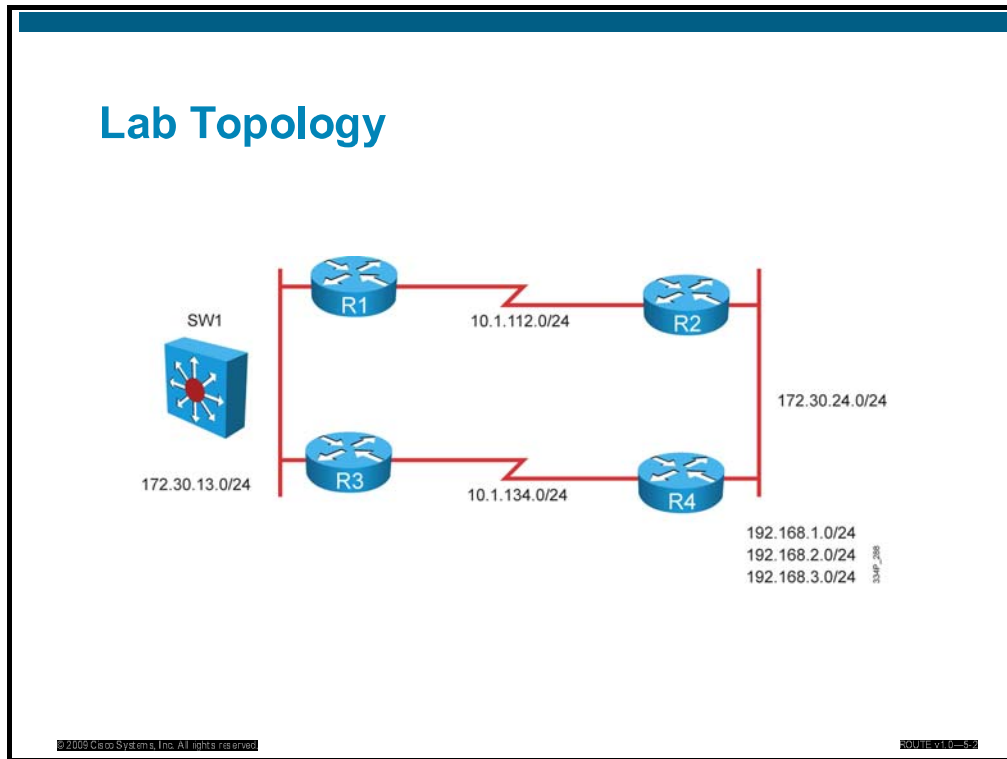
Objectives

Upon completing this lesson, you will be able to explain destination-based forwarding and configure path control using a PBR tool. This ability includes being able to meet these objectives:

- Complete the lab overview and verification
- Describe a sample solution and alternatives

Lab Overview and Verification

This topic describes the lab topology and key checkpoints used to create a solution and to start with the verification.



This figure shows the physical lab topology used for Lab 5-1: Configure and verify Path Control between multiple IP routing protocols. The topology uses four pod routers and one switch. Routers R1 to R4 build the network where path manipulation is required, and switch SW1 is one source of packets. The other source is locally-originated traffic on router R1.

Based on this topology, students will identify the required parameters and configure the EIGRP routing protocol in order to establish Layer-3 reachability in the network. Because some packets use suboptimal paths, they will use PBR as a path, control mechanism to manipulate the path of the packets.

Lab Review: What Did You Accomplish?

- Task 1: Implementing EIGRP Routing
 - What were the steps you took to configure EIGRP routing protocol?
- Task 2: Implementing Policy-Based Routing (PBR)
 - Which tools did you use to configure PBR?
 - Where did you apply the PBR configuration for the traffic passing through the router?
 - Where did you apply the PBR configuration for the traffic locally originated in the router?

In the first task, you implement EIGRP routing protocol on all routers in the lab (from R1 to R4). You establish basic reachability and use verification steps to prove that all of the EIGRP networks are in the IP routing table. Once you finish the configuration, you can identify the path used for forwarding packets.

In the second task, you identify the suboptimal path used for a group of packets and use policy-based routing (PBR) as a path-control tool to manipulate packet forwarding. You use route maps to match the packets and set the output parameters. Then you apply the route maps to traffic on the incoming interface (for packets passing the router) or to locally generated traffic on routers R1 and R3.

Verification

- Did you have enough information to create the implementation plan?
- Did you successfully configure the EIGRP routing protocol?
- Did you check which path the packets took?
- Did you successfully configure PBR?
- Did the packets take the same path they took before PBR was applied?

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE-01-05-36

A common approach to verifying the implementation process for a routing protocol is to follow these steps:

- In order to create the implementation plan, gather enough information.
- Check that routes exist in their local IP tables; they will be in the tables after any successful EIGRP configuration.
- Check the IP routing table for the routes to the destinations. The IP routing table provides all of the information needed to forward incoming packets toward their destinations. Based on the routing protocol used, the path for a packet is defined as the route to its destination, provided that exists in the IP routing table.
- To meet the requirements for path manipulation, configure the PBR path-control tool. You must specify a path for the selected group of packets within the requirements for path manipulation. You can also perform verification.
- If PBR is configured, the path for the selection of packets is different.

Checkpoints

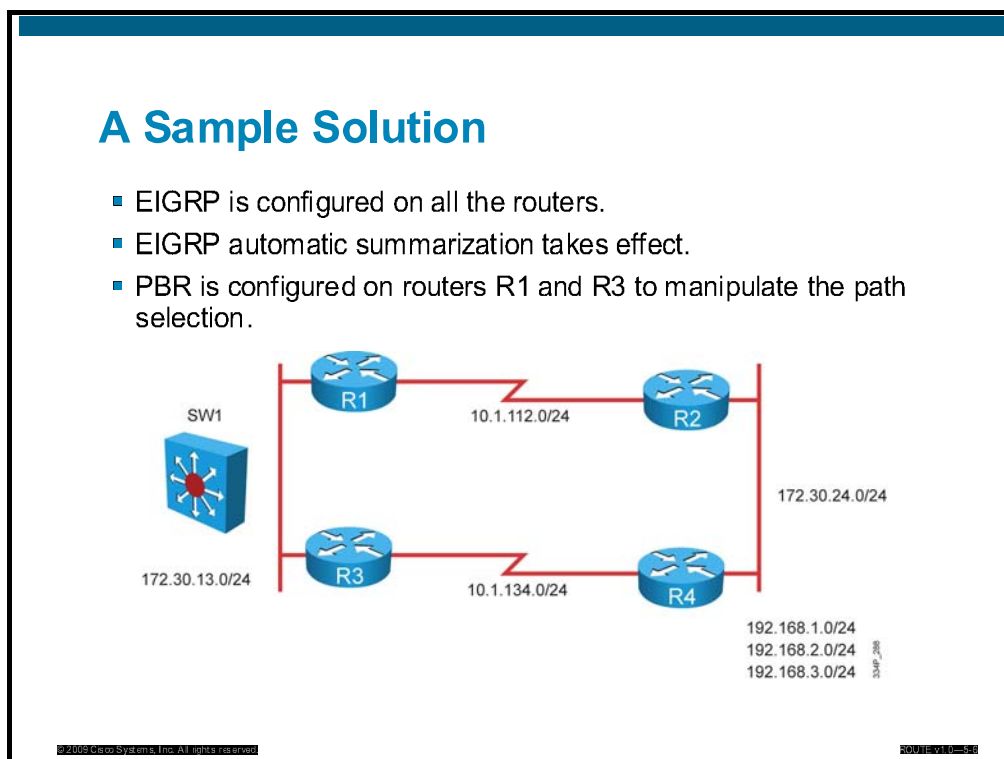
- Configure EIGRP.
- Examine and verify that the EIGRP routing protocol is up and running.
- Verify the IP routing table.
- Examine the path of IP packets.
- Configure PBR.
- Examine the path of the IP packets after the PBR configuration.

With different checkpoints, the network operator can verify for proper configuration. The following checkpoints are used for verification:

- Configure EIGRP.
- Examine and verify that the EIGRP routing protocol is up and running.
- Verify the IP routing table for EIGRP routes.
- Examine the path of IP packets when destination-based forwarding is used (without PBR configuration applied).
- Configure PBR.
- Examine the path of IP packets after PBR configuration.

Sample Solution and Alternatives

This topic describes a sample solution and other alternatives.




The sample solution in the slide includes implementation details and details for each task of the implementation plan. Different solutions are possible; the figure points out a few details of a successful configuration.

The proper implementation of the routing protocol might include the following details:

- EIGRP AS 1 is implemented.
- IP addressing is used with mask /24.
- Automatic summarization is not desired, as the hiding of subnets might prevent a successful PBR implementation.
- After successful EIGRP configuration and verification of the path taken for the packets, PBR needs to be configured. Path manipulation is configured in the router where path manipulation is desired. Many times, this router is close to the source.

Alternative Solutions

- Another routing protocol can be used to provide a different path.
- Static routes can be used.
- The metric and administrative distance can be changed.
- Another path control manipulation tool can be used.



In order to manipulate the path or to influence the destination-based forwarding, another routing protocol can be selected and configured.

Instead of a routing protocol, static routes can be configured. This solution is not scalable, but provides control over packet forwarding. Static routes can also be used as a path manipulation mechanism, because a static route has a lower administrative distance value than any routing protocol.

A network operator can manipulate the path by changing the metric of the link, the administrative distance of the static routes, or the routing protocol.

Instead of using PBR as a path-control tool, a network operator can use one of several other tools, including filtering (prefix lists, distribute lists, administrative distance, route tagging, offset lists, and Cisco IOS IP SLAs).

Q and A

- Why is routing protocol selection important?
- Why is changing the default destination-based forwarding important?
- Which tool is used for PBR implementation?
- How is classification in the route map performed?
- How is the path between the packet source and destination verified?

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE-01-00-50

A routing protocol exchanges routing updates and populates the IP routing table, which is used for destination-based forwarding.

Sometimes, simple destination-based routing is not sufficient. In these cases, network designers may route packets by source address, protocol type, or application type, so that they can optimally shape traffic patterns.

PBR uses route maps to match packets and set where to forward packets that pass a route map match clause.

The **match** commands in route maps are used to match the traffic.

Basic connectivity between the source and destination is verified using the **ping** command. The path is verified hop-by-hop using the **traceroute** command.

Summary

This topic summarizes the key points that were discussed in this lesson.

Summary

- Configure an EIGRP routing protocol to establish reachability in the network and enable destination-based forwarding.
- Configure policy-based routing to manipulate the path for a specific set of packets and verify the configuration.

References to Additional Path Control in E-Learning

Overview

The *Implementing Cisco IP Routing (ROUTE) v1.0* instructor-led training (ILT) course is a comprehensive learning experience. Among the learning tools it makes available to you are e-learning modules, which complement the classroom instructor-led content, as well as self-paced materials and demonstrations. Upon accessing the e-learning content, you will be able to reinforce the knowledge acquired in class and witness real-life scenarios demonstrated in real routers and switches. The content structure is flexible: you can navigate it at your own pace, at the time of your choosing, and at the depth you desire, according to your level of experience.

This lesson provides an overview of the *Implementing Path Control* e-learning module. The content includes demonstrations that cover several topics in path control, such as additional design considerations, route tagging, policy routing, IP SLAs.

Objectives

Upon completing this lesson, you will be able to review the e-learning modules that pertain to path control. This ability includes being able to meet these objectives:

- Describe the contents of the *Implementing Path Control* e-learning module.
- Understand the process of accessing and using e-learning content.

Preview of E-Learning on Implementing Path Control

This topic describes e-learning products that teach the configuration and implementation of path control.

CCNP E-Learning

Cisco CCNP® E-Learning allows you to do the following:

- Complement and enhance your classroom experience
- Reinforce concepts and their application
- Learn at your own pace
- Review advanced topics
- Experience real-life scenarios through directed demonstrations

One of the key ideas behind the design of the Cisco CCNP® curriculum is the understanding that there is no one "best" method of learning for every student. Some students prefer individual labs, while others prefer one-on-one tutoring, hands-on sessions, self-paced computer-assisted instruction, direct-discovery learning, or cooperative learning, among others.

E-learning solutions offer new approaches that complement and enhance classroom-based learning. Designed with flexibility and learning effectiveness in mind, the *Implementing Path Control* e-learning module is based on knowledge you acquired during your Implementing Cisco IP Routing (ROUTE) v1.0 instructor-led training course.

Implementing Path Control

Lesson	Description
Parallel Processes in Path Control	Reinforces design fundamentals through the integration of multiple Cisco IOS tools. Presents advanced topics in path control that you may find useful in certain situations.
Directed Demo: Procedures to Implement Path Control by other Methods	Demonstrates router configuration for scenarios of redistribution with multiple edge routers and dynamic advertisements using IP SLAs.
Self-Check Assessment	Assesses your understanding of the lesson.

© 2009 Cisco Systems, Inc. All rights reserved.

2009-10-23-50

The *Implementing Path Control* e-learning module covers topics that complement the classroom module you are reviewing. After a brief introductory lesson that discusses design issues and the integrated approach to path control, you will be able to follow a directed demonstration of real-life scenarios that combine multiple Cisco IOS features to accomplish specific objectives. These demonstrations include the use of tools such as administrative distance, route tagging, and IP SLA tracking.

Where to Find E-Learning Modules



The “Implementing Path Control” e-learning module is available on a CD as part of your classroom materials. Please contact your instructor with any questions you have about finding and accessing the CD.

The objectives of the *Implementing Path Control* e-learning module are part of your CCNP certification exam. As such, candidates for CCNP certification should review these objectives.

E-Learning Module Structure

The screenshot displays a three-panel interface for an e-learning module. The top-left panel, titled "Addressing and Topology Specifics", shows a network diagram with three routers (R1, R2, R3) and their interconnections. R1 and R2 are connected via a link with IP 10.2.7.0/24. R2 and R3 are connected via a link with IP 10.3.7.0/24. R1 and R3 are connected via a link with IP 192.168.1.0/24. R1 and R2 are also connected via a link with IP 192.168.2.0/24. R2 and R3 are connected via a link with IP 192.168.253.0/24. R1 and R3 are connected via a link with IP 192.168.254.0/24. The diagram is labeled with "OSPF" and "EIGRP". The top-right panel, titled "Debrief : Alternative Configuration", shows a configuration snippet for a route-map named TAGS. The configuration includes:

```
route-map TAGS deny 10
match tag 1000
route-map TAGS permit 20
set tag 1000
....
router ospf 1
redistribute eigrp 1 metric 4 route-map TAGS
....
router eigrp 1
redistribute ospf 1 metric 1000000 0 255
1 800 route-map TAGS
....
```

 A callout bubble next to the configuration says "Keep it simple". The bottom panel shows a console output from a Cisco router. The output includes:

```
R1 | R2 | ISP |
Policy routing matches: 0 packets, 0 bytes
route-map SETTAG, permit, sequence 20
Match clauses:
Set clauses:
Policy routing matches: 0 packets, 0 bytes
R1#trace 192.168.254.1

Type escape sequence to abort.
Tracing the route to 192.168.254.1

 0  1 192.168.2.2 12 msec * 12 msec
R1#
R1#
```

 The interface includes navigation controls at the bottom.

Utilizing a combination of lecture sessions, animated content, lab demonstrations, and assessments, each e-learning module presents a hierarchical structure of lessons and topics. You can navigate through the structure by using an intuitive GUI that includes playback controls, slide selection, and lesson and topic selection. By using these tools, you will be able to navigate the content at your own pace.

The three-panel screen used for directed demonstrations allows you to focus on the device console demonstrations, while still allowing you to look at the command syntax and topology diagram for verification and additional information.

Summary

This topic summarizes the key points that were discussed in this lesson.

Summary

- The *Implementing Path Control* e-learning module complements the current module and is part of the same Cisco certification.
- You can access e-learning content from the Cisco certifications website, and complete it according to the learning style that is most effective for you.

Module Summary

This topic summarizes the key points that were discussed in this module.

Module Summary

- When common network performance issues such as redundancy, resiliency, performance and availability are solved, the proper path control tools must be used for efficient control of path selection. Several solutions exist to avoid these issues.
- Policy based routing (PBR), filters (prefix lists, distribute lists, offset lists, and so on), Cisco IOS IP SLAs, and route tags can be used to implement path control in large networks.

This module covered common network performance issues such as redundancy, resiliency, performance, and availability. A solution to these issues frequently results in an inefficient path being selected, and the default destination-based routing is not the best solution. You must change the routing table by manipulating routing updates that are entered in the routing table or by bypassing the routing table.

Several tools can be used for efficient path control:

- Route maps
- Prefix lists
- Distribute lists
- Administrative distance
- Route tagging
- Offset lists
- Cisco IOS IP SLAs
- Policy based routing

All of these tools must first match traffic, then set actions for the matched traffic.

Module Self-Check

Use the questions here to review what you learned in this module. The correct answers and solutions are found in the Module Self-Check Answer Key.

- Q1) Which four issues are the result of multiple paths? (Choose four.) (Source: Assessing Path Control Network Performance Issues)
- A) suboptimal routing
 - B) unavailable redundant paths
 - C) traffic optimized for only one application
 - D) uneven load sharing
 - E) incomplete route summarization
- Q2) Which three tools can be used for path manipulation? (Choose three.) (Source: Assessing Path Control Network Performance Issues)
- A) filters
 - B) policy-based routing (PBR)
 - C) redistribution
 - D) route maps
 - E) summarization
- Q3) Path control tools can be used to bypass the IP routing table. (Source: Assessing Path Control Network Performance Issues)
- A) true
 - B) false
- Q4) Which two commands are used inside route maps? (Choose two.) (Source: Assessing Path Control Network Performance Issues)
- A) **classify**
 - B) **set**
 - C) **redistribute**
 - D) **match**
 - E) **policy**
- Q5) Which command is used to match traffic inside route maps based on access list 10? (Source: Assessing Path Control Network Performance Issues)
- A) **match ip address 10**
 - B) **match access-list 10**
 - C) **match ip acl 10**
 - D) **match ip access-list 10**
- Q6) Which path control tool uses mask filtering, for which either a defined prefix is matched or part of an address space with a subnet mask longer or shorter than a set number is matched? (Source: Assessing Path Control Network Performance Issues)
- A) distribute lists
 - B) administrative distance
 - C) route tagging
 - D) offset lists
 - E) Cisco IOS IP SLAs
 - F) prefix lists

- Q7) Which tool can be used for end-to-end network performance tests based on clear measurement metrics, as well as for path control? (Source: Assessing Path Control Network Performance Issues)
- A) distribute lists
 - B) administrative distance
 - C) route tagging
 - D) offset lists
 - E) Cisco IOS IP SLAs
 - F) prefix lists
- Q8) Policy-based routing avoids destination-based routing. (Source: Assessing Path Control Network Performance Issues)
- A) true
 - B) false
- Q9) To which two of these are path control tools applied? (Choose two.) (Source: Configuring and Verifying Path Control)
- A) incoming traffic
 - B) outgoing traffic
 - C) traffic locally generated in the router
 - D) MTU-sized packets
- Q10) Which two commands provide a basic connectivity test? (Choose two.) (Source: Configuring and Verifying Path Control)
- A) **show route-map**
 - B) **show ip policy**
 - C) **ping**
 - D) **traceroute**
- Q11) Which command defines the interface on which to output packets that pass a route map match clause for policy routing and for which the Cisco IOS software has no explicit route to a destination? (Source: Configuring and Verifying Path Control)
- A) **set ip next-hop 10.1.1.1**
 - B) **set interface ethernet 0**
 - C) **set ip default next-hop 10.1.1.1**
 - D) **set default interface ethernet 0**
- Q12) Which command displays route maps that are configured on interfaces? (Source: Configuring and Verifying Path Control)
- A) **show ip policy**
 - B) **show route-map map-name**
 - C) **show ip routing table**
 - D) **show ip interface brief**
- Q13) Policy-based routing uses a route map for its configuration and can be fast-switched or switched by Cisco Express Forwarding. (Source: Configuring and Verifying Path Control)
- A) true
 - B) false

- Q14) The **show ip policy** command output shows policy routing matches. (Source: Configuring and Verifying Path Control)
- A) true
 - B) false

Module Self-Check Answer Key

- Q1) A, B, C, D
- Q2) A, B, D
- Q3) A
- Q4) B, D
- Q5) A
- Q6) F
- Q7) E
- Q8) A
- Q9) A, C
- Q10) C, D
- Q11) D
- Q12) A
- Q13) A
- Q14) B

Connect an Enterprise Network to an ISP Network

Overview

The Internet has become a vital resource in many organizations, which requires a single connection and frequently redundant connections to multiple Internet service providers (ISPs). With multiple connections, the Border Gateway Protocol (BGP) is an alternative to using default routes to control path selections. BGP provides several alternatives to using default routes.

Use of BGP as a routing protocol requires that an administrator understand how to properly configure BGP for scalable internetworking. This module discusses BGP configuration and verification when connecting an enterprise network to ISP networks.

Module Objectives

Upon completing this module, you will be able to implement and verify a Layer 3 solution using BGP to connect an enterprise network to an Internet service provider. This ability includes being able to meet these objectives:

- Discuss the components and methods used to connect an enterprise network to an ISP.
- Explain the advantages of using BGP to connect an enterprise network to an ISP, while comparing the functions and uses of EBGP and IBGP.
- Compare and contrast the requirements for establishing EBGP-to-IBGP neighbor relationships.
- Implement BGP operations for ISP connections in an enterprise network.
- Discuss the lab results for configuring BGP operations.
- Configure and verify BGP operations in a multihomed environment using the BGP attributes and route maps to control all BGP routes to and from the router.
- Discuss the lab results for manipulating the EBGP path selections.
- Discuss the e-learning modules that pertain to additional topics in IPv6 and remote access connectivity.

Planning the Enterprise-to-ISP Connection

Overview

Planning needs to be completed before the connection of an enterprise network to an Internet service provider (ISP) can be implemented. To plan the connectivity to an ISP properly, the designer must understand many aspects, including the session origination point—either from the enterprise network toward the Internet or from the Internet toward the enterprise network; the prerequisites for successfully implementing such connectivity; and the available routing options and which routing option should be used in certain cases. The connectivity redundancy considerations must also be taken into account.

Objectives

Upon completing this lesson, you will be able to justify the components and methods used to connect an enterprise network to an ISP. This ability includes being able to meet these objectives:

- Connecting enterprise networks to an ISP.
- Exchanging routing updates with an ISP.
- Define the types of connections to an ISP.

Connecting Enterprise Networks to an ISP

This topic describes the connectivity requirements between an enterprise network and an ISP.

Session Origin Initiation

Enterprise session initiation requirement:

- **One-way:** Connectivity from an enterprise network toward the Internet is the only connectivity required.
- **Two-way:** Connectivity from the Internet to an enterprise network is also required.

Solutions:

- **One-way:** private IP address space with address translation
- **Two-way:** public IP address space (in combination with private) and proper routing

The diagram shows three main components: Company A, an ISP, and the Internet. Company A is represented by a yellow box containing a blue router and several computer icons. The ISP is represented by a yellow box containing a blue router. The Internet is represented by a white cloud containing several server icons. A solid red line connects the two routers. Dashed lines show connectivity from the Company A router to the Internet cloud and from the ISP router to the Internet cloud. A small copyright notice '© 2009 Cisco Systems, Inc. All rights reserved.' is visible at the bottom left of the diagram area.


Modern corporate IP networks are connected to the global Internet and make use of the Internet for some of their data transport needs. Corporations provide many services via the Internet to customers and business partners—systems from web servers to mainframes to workstations can be accessible from anywhere in the world.

In the rare cases in which only connectivity from the clients to the Internet is required, public addresses with network translation are used to allow clients on a private network to communicate with servers on the public Internet.

Typically, not only do clients from an enterprise network need to access external resources located on the public Internet, but clients external to the enterprise network also need to access resources in the enterprise network.

Enterprise Network-to-ISP Connectivity Requirements

- Public IP address space (subpool or whole /24 subnet)
- Link type and bandwidth availability
- Routing options
- Connection redundancy
- Independency in regard to an ISP:
 - Public IP address space
 - AS number



The first parameter that needs to be determined is the number of public IP addresses, which are used to translate client private addresses for those clients that need to access resources on the Internet. These public IP addresses are also used for those enterprise servers that need to be accessible from the Internet. These servers are either equipped with public addresses or addresses that are statically translated from private to public addresses.

In the second parameter, the enterprise network-to-ISP connection link type and speed, which depends on the ISP, needs to be determined. The link types available include leased line, Ethernet over fiber optics or copper, and xDSL. The bandwidth assignment must be done properly to address the enterprise Internet connectivity requirements.

In the third parameter, a proper routing protocol must be selected. The selection is typically made between static and dynamic routing.

The fourth parameter deals with connectivity redundancy. An evaluation is needed to determine what kind of redundancy is required for the enterprise network-to-ISP connectivity. Redundancy includes edge router redundancy, link redundancy, and ISP redundancy, where the ISP redundancy connects to multiple ISPs.

When assessing these parameters, it must be taken into account whether an enterprise network needs to be independent of the selected ISP(s). If independency is required, the public IP address space should not be used from the ISP public address space, but should instead be acquired from the regional Internet authority. Similarly, independency is required for an autonomous system (AS). The AS number must be public and not from the private AS number pool that can be assigned by the ISP.

Exchanging Routing Updates with an ISP

This topic describes the methods for exchanging routing information across an ISP.

Reachability

- Circuit emulation ✓
- Static routes ✓
- MPLS VPNs ✓
- BGP

Static routes and BGP are typically selected for Internet connectivity.

Company A Route exchange ISP

✓ Previously covered
 Not covered yet


Connecting an enterprise network to an ISP requires routing information to be exchanged between the network and the ISP.

The selection of the routing means depends on the answers to these questions:

- Should this routing respond to the changes in a network topology?
- Does routing need to support one link or multiple links to an ISP?
- Is traffic load balancing over multiple links required?
- Will the enterprise network be connected to multiple ISPs?
- Should the ISP offer merely a transport capability that requires connecting different customer locations, perhaps via certain Layer 2 technologies?
- What is the amount of routing information that needs to be exchanged with an ISP?
- What routing options does the ISP offer?

Using Circuit Emulation

- Used to provide different Layer 2 connectivity to customers via the common Layer 3 infrastructure of a service point
 - Ethernet, Frame Relay, PPP, HDLC, ATM, Layer 2 connectivity
 - No routing with the service point from the customer perspective



Sometimes a customer needs Layer 2 connectivity between two or more locations. The following examples illustrate the need for Layer 2 connectivity:

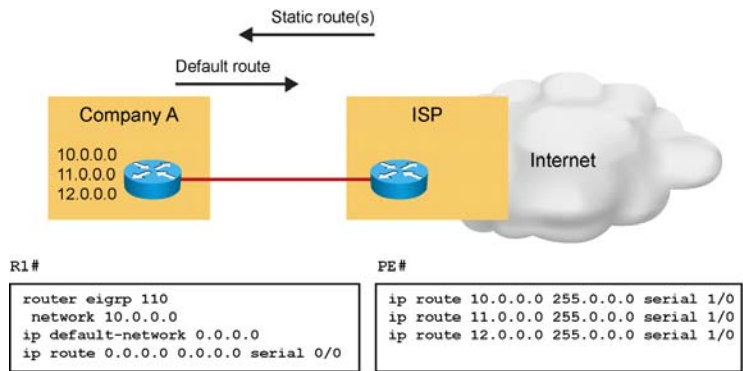
- The locations might include data centers with geographically distributed clusters that require Layer 2 connectivity in order to properly function.
- The customer is in a migration process but still requires Layer 2 connectivity.
- The customer is connecting to another partner and there is a requirement for Layer 2 connectivity.

A Layer 2 connection requirement might range from Ethernet, Frame Relay, PPP, High-Level Data Link Control (HDLC), or even ATM. In such cases, modern service providers employ their IP-based core network enhanced with Multiprotocol Label Switching (MPLS) technologies to provide such connectivity.

With this connectivity, there is no need for a routing exchange between the ISP and the customer. From the customer, it looks like the ISP is providing a Layer 2 port. From the perspective of the ISP, two ports from distant locations must be connected together.

Using Static Routes

- The customer uses the default route toward the ISP.
- The service provider uses static route(s) for customer public networks.
- No automatic adjustment to any changes in the network.




Static routes are the simplest way of exchanging routing information with an ISP. These static routes must agree with the ISP upon routing configuration and they do not change afterward; or if they do change, these changes should not occur very often.

Static routes are typically used for Internet connectivity when a customer is connected through a single connection to an ISP. The customer can use the default route towards the ISP and the ISP must deploy a static route or routes that encompass the public networks of the customer. The ISP would typically also redistribute this information into its Border Gateway Protocol (BGP).

Although static routes are a simple solution, they also have some drawbacks, especially in terms of flexibility and adaptability. For example, if there were a change in a network topology beyond a directly connected link failure, the static routes would not adapt. If adaptation needs to be achieved, either the static routes need to be combined with an IP service level agreement (SLA) functionality that in turn would indicate that a static route is down if a certain condition is met or dynamic routing is used. Note that using an IP SLA cannot completely substitute dynamic routing because it cannot react to all changes in the topology in the Internet.

Using MPLS VPN

- Used to connect multiple customer locations via a service provider's common Layer 3 infrastructure
 - A special VPN can be used to provide Internet connectivity
 - Routing used can be static or dynamic depending on the SP
 - The customer routers are connected to the service provider PE routers



Multiprotocol Label Switching Virtual Private Networks (MPLS VPNs) are used when a customer has multiple locations that need to be interconnected through an ISP and does not want to use expensive Layer 2 technologies such as leased lines.

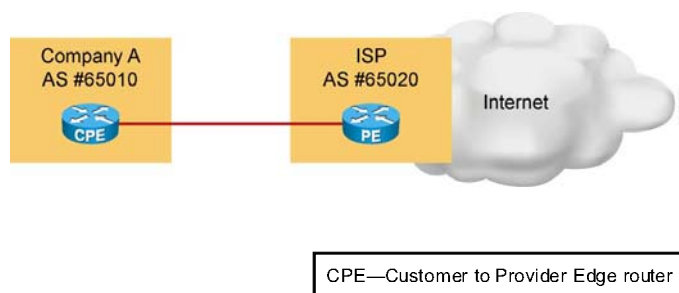
With MPLS VPNs, the ISP uses a common IP-based core network enhanced with MPLS technology to provide secure and manageable connectivity for different customers to their geographically disperse sites. In this way, the traffic from different customers can share the same physical wire, but at the same time, traffic is tagged with the labels in such way that the traffic cannot intermix.

When a customer uses MPLS VPN functionality, the routing between the customer and the Internet service provider is required in order to provide connectivity between the customer locations. Routing options range from static to dynamic and include Routing Information Protocol (RIP), Enhanced Interior Gateway Routing Protocol (EIGRP), Open Shortest Path First (OSPF), Intermediate System-to-Intermediate System (IS-IS), or even BGP, depending on what the ISP offers. Different locations could use different routing protocols, though this situation typically would not occur.

With an MPLS VPN deployment, the service provider can also offer Internet connectivity through the same MPLS core network either through a special Internet VPN or through a global routing table. To exchange the Internet routing information, either Border Gateway Protocol (BGP) or a default routing is used.

Using BGP

- The customer deploys BGP to announce its public networks
- The ISP announces a default route, a subset of Internet routes, or a complete Internet routing table
- Typically used for inter-AS routing



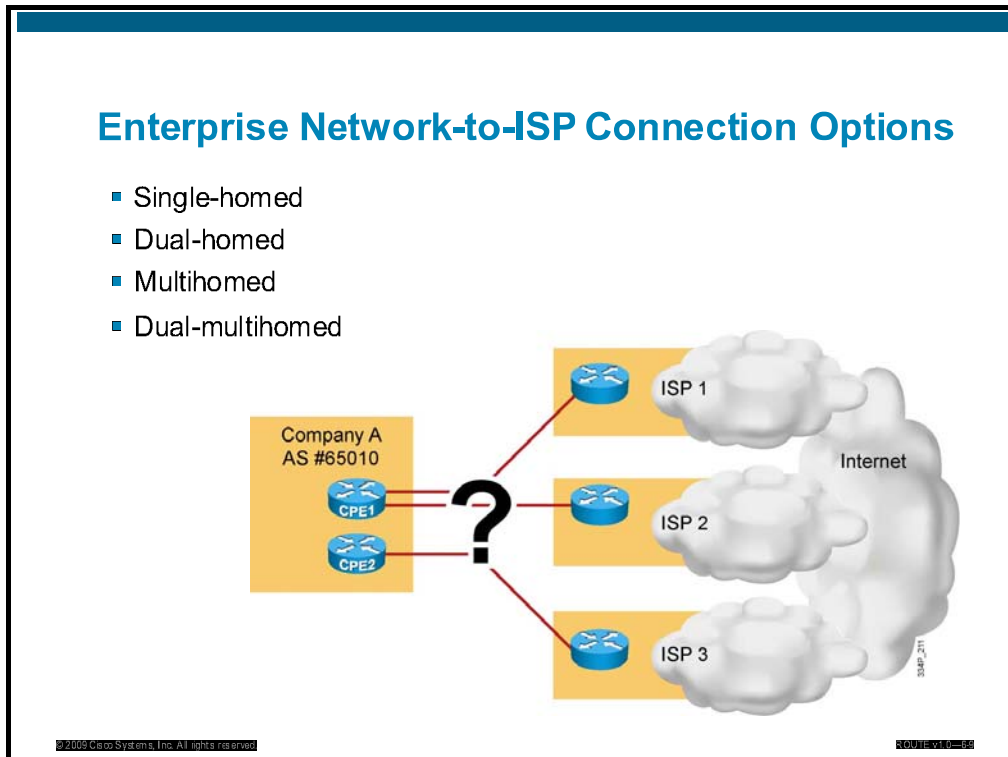
A typical option used to provide dynamic routing exchange when Internet connectivity is deployed is BGP. BGP dynamically exchanges routing information and thus reacts to topology changes including those changes beyond a customer-to-ISP link failure.

From the routing perspective, three options can be used:

- An ISP (or multiple ISPs) announces a default route only. If two ISPs are used, one ISP will be primary and the second ISP will serve as the backup.
- An ISP (or multiple ISPs) announces a default route and a subset of Internet routes, typically its own public address space. This results in the shortest path being used to the directly connected public networks. For public networks that are not directly connected, one of the ISPs is the primary one and the second ISP is the backup.
- An ISP (or multiple ISPs) announces a complete Internet routing table. The shortest path to any destination will be chosen according to the routing table.

Defining the Types of Connections to an ISP

This topic describes the types of enterprise-to-ISP connections and their effect on the selection of an exchange method.



When connecting an enterprise network to an ISP, redundancy is a big concern. The following different aspects can be addressed in order to achieve connectivity redundancy:

- Deployment of redundant links
- Deployment of redundant equipment
- Use of redundancy within a single router

From the connection perspective, a customer can be connected to a single ISP or to multiple ISPs.

With a single ISP connection, redundancy can still be achieved if two links toward the same ISP are used effectively, making a customer dual-homed. When no link redundancy is used, the customer is simply single-homed and in case of a failure within an ISP network, the connectivity to the Internet is interrupted.

With multiple ISP connections, the redundancy is built into the design because the customer is multihomed and thus resistant to a single ISP failure. To enhance the resiliency further, a customer can have two links toward a single ISP making the solution dual-multihomed.

Single-homed ISP Connectivity


Link failure results in broken connectivity.

Option 1: Routing with static routes

- A default route from an enterprise network
- A static route(s) from an ISP for customer networks

Option 2: Routing with BGP

- The customer announces its public network(s)
- The ISP announces the default route to the customer



Single-homed ISP connectivity is used in cases when a loss in Internet connectivity is not as vital or problematic to a customer (though these days the Internet typically is a vital resource).

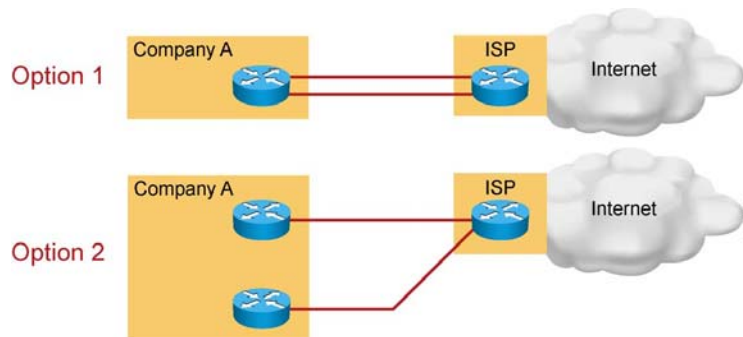
In this case, a customer uses a single connection to a single ISP. The type of connection depends on the ISP offering and can include a leased line, xDSL, or an Ethernet connection. Failure of a link results in a no Internet connectivity.

Such Internet access to a single ISP does not require BGP. Static routes are typically used to handle the routing. If BGP is used, the customer uses it to dynamically announce its public networks to an ISP and the ISP announces only a default route to the customer, because that is sufficient to provide the connectivity through a single link to the Internet.

Dual-homed ISP Connectivity

Characteristics:

- Connected with two links to the same ISP
- Can use a single router or two edge routers
- Can use static routes or BGP



When a customer is connected to a single ISP only, resiliency can still be achieved by deploying a second link to the same ISP. With a second link being used, the routing must be properly configured to allow usage of such a link.

Depending on the SLA signed with the ISP, the routing deployed could achieve these benefits:

- Primary and backup link functionality where a single primary link is used to forward and receive traffic to and from the ISP, and the secondary link is used only when the first link fails.
- Load sharing between the links is achieved with Cisco Express Forwarding switching.


In both cases, the routing used can be based on either static or dynamic routes, which would typically include BGP.

To enhance the resiliency further, the two links can terminate at separate customer routers.

Multihomed ISP Connectivity

Characteristics:

- Connected to two or more different ISPs
- Can use a single router or multiple edge routers
- Dynamic routing with BGP



A proper resiliency is achieved by connecting to two or more different ISPs. The benefits of connecting to two or more different ISPs are the following:


- Resistance to a failure beyond a directly connected link to a single ISP.
- Sharing the load for different destination networks between ISPs based on the network proximity.
- Scaling the solution beyond two ISPs.
- Achieving a solution independent of an ISP. An ISP change requires an update of routing configuration and change of a link. The public IP address space used remains the same.

Connections from different ISPs can terminate on the same router or different routers to further enhance the resiliency. The routing must be capable of reacting to dynamic changes, and BGP is typically used to achieve this flexibility.

Dual-multihomed ISP Connectivity

Characteristics:

- Connected to two or more different ISPs with two links per ISP
- Typically uses multiple edge routers (one per ISP)
- Dynamic routing with BGP



Multihoming exists when an autonomous system has more than one connection to the Internet. Two typical reasons for multihoming are as follows:

- To increase the reliability of the connection to the Internet
 - If one connection fails, the other connection remains available.
- To increase the performance of the connection
 - Better paths can be used to certain destinations.

The benefits of BGP are apparent when an autonomous system has multiple External Border Gateway Protocol (EBGP) connections to either a single autonomous system or multiple autonomous systems. Having multiple connections allows an organization to have redundant connections to the Internet so that if a single path becomes unavailable, connectivity can still be maintained.

Summary

This topic summarizes the key points that were discussed in this lesson.

Summary

- Connecting an enterprise network to an ISP requires, at a minimum, a public IP address pool, a proper link to the ISP, consideration of redundancy requirements, and the proper routing protocol.
- To exchange routing updates with an ISP, the customer can use different options. Static routes and BGP are the options that are most commonly used.
- The way in which the customer connects to an ISP depends on the redundancy requirements, where a single-homed connectivity has no redundancy, and the dual-multihomed connectivity has the most redundancy built in.

Lesson 2

Considering the Advantages of Using BGP

Overview

Border Gateway Protocol (BGP) is an exterior gateway protocol (EGP) and this lesson describes how BGP may be used by enterprises to connect to Internet service providers (ISPs). BGP routes between autonomous systems using path vector attributes upon which routing policy decisions at the autonomous system level are enforced.

Objectives

Upon completing this lesson, you will be able to explain the advantages of using BGP to connect an enterprise network to an ISP, while comparing the functions and uses of EBGP and IBGP. This ability includes being able to meet these objectives:

- Use BGP to connect to an ISP.
- Describe BGP multihoming options.
- Configure BGP routing between autonomous systems.
- Explain path vector functionality.
- Identify features of BGP.

Using BGP to Connect to an ISP

This topic describes connectivity between an enterprise network and an ISP that requires the use of BGP, including a description of the issues that arise when an enterprise decides to connect to the Internet through multiple ISPs.

BGP Terminology

- Autonomous system (AS): a collection of networks under a single administrative domain
- Interdomain routing: routing between the customer and the ISP
- Internal routing: uses IGP protocol (RIP, OSPF, EIGRP, and so on) to exchange routing information inside the AS
- External routing: uses EGP protocol (BGP) to exchange routes between autonomous systems
- Two BGP implementations:
 - Internal BGP (IBGP): when BGP is used inside an AS
 - External BGP (EBGP): when BGP is used between autonomous systems

The Internet is a collection of autonomous systems that are interconnected to allow communication among these autonomous systems. BGP provides the routing between these autonomous systems.

Enterprises that want to connect to the Internet do so through one or more ISPs. If your organization has only one connection to one ISP, then you probably do not need to use BGP; instead, you would use a default route. However, if you have multiple connections to one or multiple ISPs, then BGP might be appropriate because it allows manipulation of path attributes, so that the optimal path can be selected.

To understand BGP, it is necessary to first understand how it is different from the other protocols discussed so far in this course. One way to categorize routing protocols is by whether the protocols are interior or exterior, as follows:

- **IGP:** Interior gateway protocol (IGP) is a routing protocol that exchanges routing information within an autonomous system. Routing Information Protocol (RIP), Open Shortest Path First (OSPF), and Enhanced Interior Gateway Routing Protocol (EIGRP) are examples of IGPs.
- **EGP:** Exterior gateway protocol (EGP) is a routing protocol that exchanges routing information between different autonomous systems. BGP is an example of an EGP.

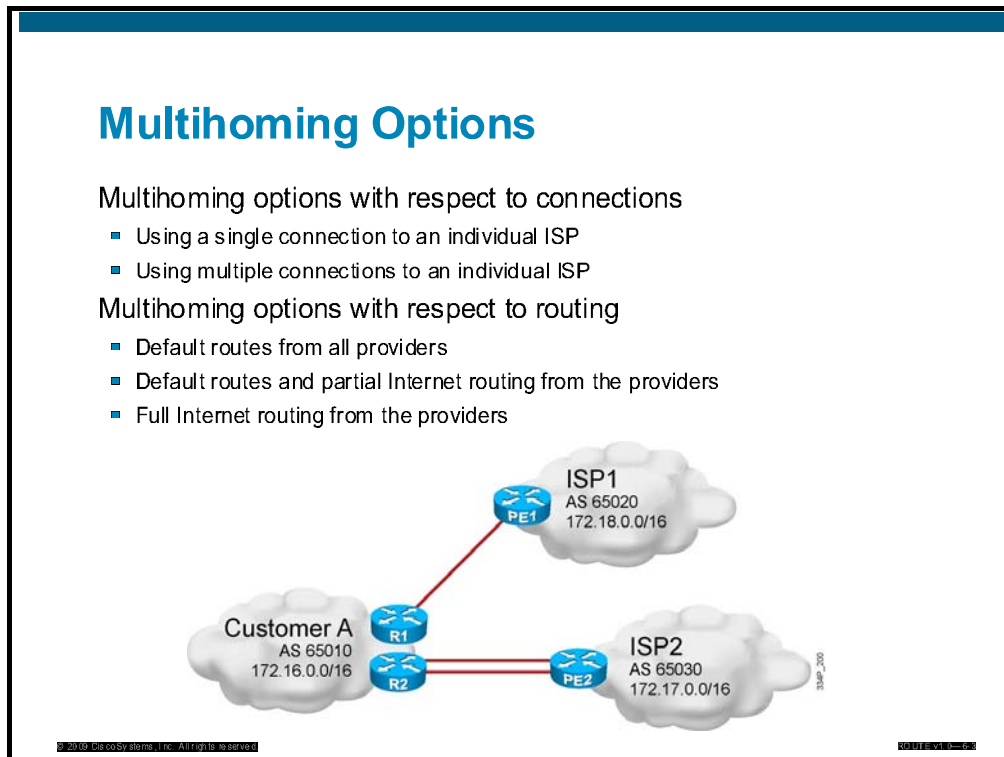
BGP is an Interdomain Routing Protocol (IDRP), also known as an EGP. BGP version 4 (BGP4) is the latest version of BGP and is defined in RFC 4271. As noted in this RFC, the classic definition of an autonomous system is “a set of routers under a single technical administration, using an IGP and common metrics to route packets within the autonomous system, and using an inter-autonomous system routing protocol (also called an EGP) to determine how to route packets to other autonomous systems.”

Autonomous systems can use more than one IGP, potentially with several sets of metrics. From the point of view of BGP, the most important characteristic of an autonomous system is that it appears to other autonomous systems to have a single coherent interior routing plan and presents a consistent picture of reachable destinations. All parts of an autonomous system must connect to each other.

When BGP is running between routers in different autonomous systems, it is called External BGP (EBGP). When BGP is running between routers in the same autonomous system, it is called Internal BGP (IBGP). BGP allows the path that packets take to be manipulated by the autonomous system, as described in this module. It is important to understand how BGP works to avoid creating problems for the autonomous system as a result of running BGP.

BGP Multihoming Options

This topic describes BGP multihoming options.



An organization can be multihomed to either a single ISP or to multiple ISPs. A drawback to having all connections homed to a single ISP is that connectivity issues in that single ISP can cause your autonomous system to lose connectivity to the Internet. By having connections to multiple ISPs, an organization gains the following benefits:

- A redundancy with the multiple connections
- Not tied into the routing policy of a single ISP
- More paths to the same networks for better policy manipulation

A multihomed autonomous system will run EBGP with its external neighbors and might also run IBGP internally.


If an organization has determined that it will perform multihoming with BGP, there are three common ways to do this:

- Each ISP passes only a default route to the autonomous system.
 - The default route is passed to the internal routers.
- Each ISP passes only a default route and provider-owned specific routes to the autonomous system.
 - These routes may be passed to internal routers or all internal routers in the transit path can run BGP and pass these routes between them.
- Each ISP passes all routes to the autonomous system.
 - All internal routers in the transit path run BGP and pass these routes between them.

These options are described in the following pages.

Default Routes from Providers

- Customer A receives the default route from each ISP.



The first multihoming option is to receive only a default route from each ISP. This configuration requires the fewest resources within the autonomous system because a default route is used to reach any external destinations. The autonomous system sends all its routes to the ISPs, which process and pass the routes on to other autonomous systems.

If a router in the autonomous system learns about multiple default routes, the local interior routing protocol installs the best default route in the routing table. From the perspective of this router, it chooses the default route with the least-cost IGP metric. This IGP default route will route packets destined to the external networks to an edge router of this autonomous system, which is running EBGp with the ISPs. The edge router will use the BGP default route to reach all external networks.

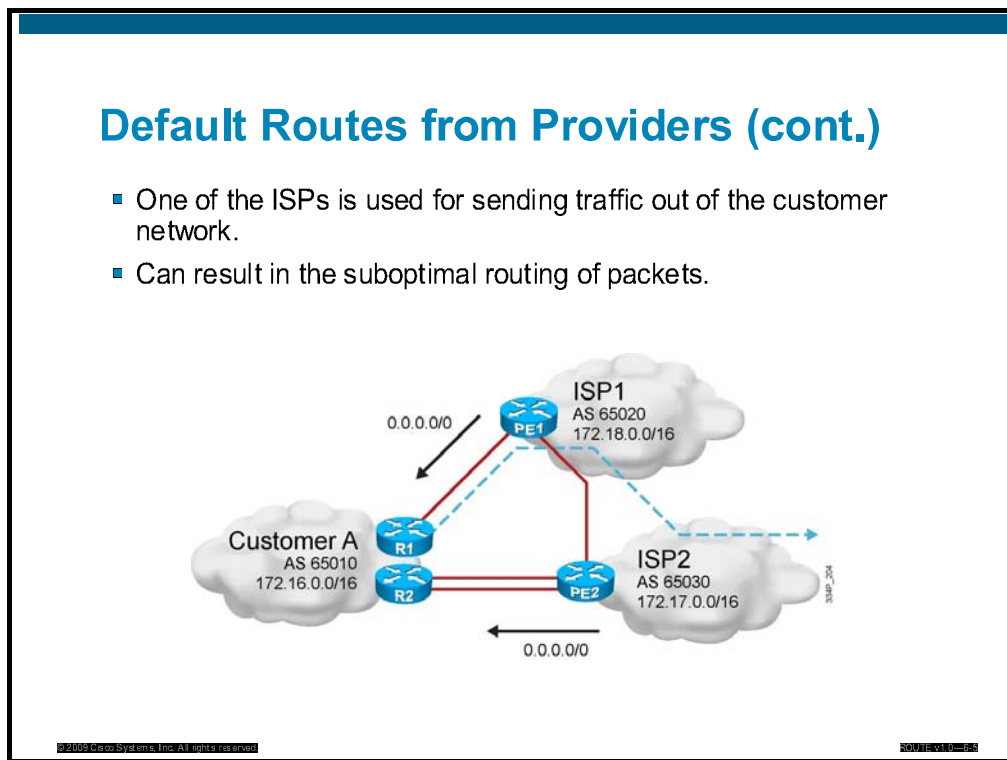
The route that inbound packets take to reach the autonomous system is chosen outside the autonomous system (within the ISPs and other autonomous systems).

Regional ISPs that have multiple connections to national or international ISPs commonly implement this option. The regional ISPs do not use BGP for path manipulation; however, ISPs require the capability of adding new customers as well as the networks of the customers. If the regional ISP does not use BGP, then each time that the regional ISP adds a new set of networks, the customers must wait until the national ISPs add these networks to their BGP process and place static routes pointing at the regional ISP. By running EBGp with the national or international ISPs, the regional ISP needs to add only the new networks of the customers to its BGP process. These new networks automatically propagate across the Internet with minimal delay.

A customer that chooses to receive default routes from all providers must understand the following limitations of this option:

- Path manipulation cannot be performed because only a single route is being received from each ISP.

- Bandwidth manipulation is extremely difficult and can be accomplished only by manipulating the IGP metric of the default route.
- Diverting some of the traffic from one exit point to another is challenging because all destinations are using the same default route for path selection.

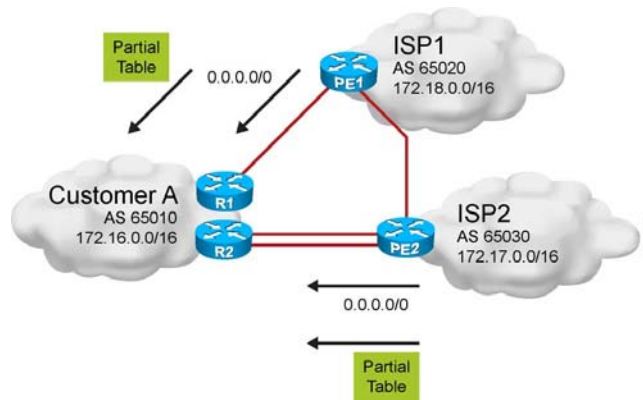


In the above figure, AS 65020 and AS 65030 send default routes into AS 65010—the Customer A network. Because of the IGP metric in the Customer A network and because of the configuration of routers R1 and R2, the provider edge (PE)1 router from ISP 1 is selected as the default gateway in order to reach any external network outside of the Customer A autonomous system.

This procedure can lead to suboptimal routing. For example, if the goal is to reach network 172.17.0.0, the packets will be sent toward ISP 1 to router PE1 because the PE1 router is the preferred way out of the Customer A autonomous system. ISP 1 then forwards the traffic to the final destination to ISP 2.

Default Routes and Partial Table from Providers

- Customer A receives the default route from each ISP.
- Customer A receives a partial routing table from each ISP.



In the second design option for multihoming, all ISPs pass default routes plus the selected specific routes to the autonomous system.


An enterprise running EBGP with an ISP that wants a partial routing table generally receives the networks that the ISP and its other customers own. The enterprise can also receive the routes from any other autonomous system.

Major ISPs are assigned between 2000 and 10,000 classless interdomain routing (CIDR) blocks of IP addresses from the Internet Assigned Numbers Authority (IANA), which the ISPs reassign to their customers. If the ISP passes this information to a customer that wants only a partial BGP routing table, the customer can redistribute these routes into its IGP. The internal routers of the customer (these routers are not running BGP) can then receive these routes via redistribution. The routers can take the nearest exit point based upon the best metric of specific networks instead of taking the nearest exit point based on the default route.

Acquiring a partial BGP table from each provider is beneficial because path selection will be more predictable than when using a default route.

Default Routes and Partial Table from Providers (cont.)

- The partial table is used to forward traffic to the correct ISP.
- If the destination is unknown, then a default route to one of the ISPs is used.

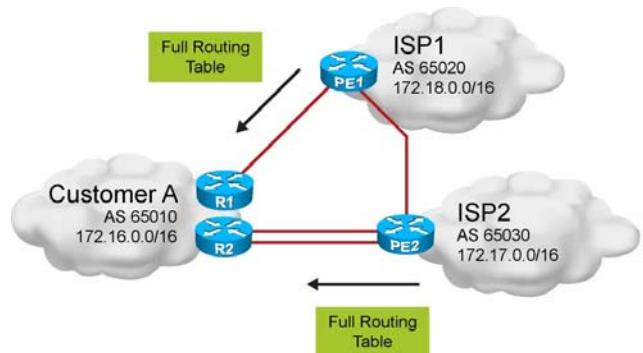


In the figure above, ISPs in AS 65020 and AS 65030 send default routes and the routes that each ISP owns (partial table) to Customer A (AS 65010).

By running IBGP between the internal routers R1 and R2 within AS 65010, AS 65010 can choose the optimal path to reach the destination networks within ISP 1 and ISP 2. If Customer A is sending the traffic to an unknown destination, one of the default routes (from ISP 1 or ISP 2) is used. Again, this process can lead to suboptimal routing as shown in the figure above. The unknown routes to another autonomous system are not shown in the figure because these routes are not specifically advertised to AS 65010 by ISP 1 and ISP 2. The IGP metric that is used to reach the default route within the autonomous system selects the default gateway out of the Customer A network.

Full Internet Routing from Providers

- Customer A receives a full routing table from each ISP.
- Requires that enough memory and CPU resources are available.



In the third multihoming option, all ISPs pass all routes to the autonomous system, and IBGP is run on at least the routers in the transit path in this autonomous system. This option allows the internal routers of the autonomous system to take the path through the best ISP for each route.

This configuration requires many resources within the autonomous system because it must process all the external routes.

The autonomous system sends all its routes to the ISPs, which process the routes and pass the routes to other autonomous systems.

In the above figure, AS 65020 and AS 65030 send all routes into AS 65010. The ISP that a specific router within AS 65010 uses to reach the external networks is determined by the BGP protocol.

The routers in AS 65010 can be configured to influence the path to certain networks. For example, router R1 and router R2 can influence the outbound traffic from AS 65010.

BGP Routing Between Autonomous Systems

This topic describes how BGP routes between autonomous systems.

Autonomous System

- Autonomous system (AS) – a collection of networks under a single technical administration.
 - 16-bit numbers (as of January 2009 32-bit numbers are available)
 - Ranging from 1 to 65535
 - Private AS: 64512–65535
- Internet Assigned Numbers Authority (IANA) allocates AS numbers.
- IGP operates within an AS.
- BGP is used between autonomous systems.



Recall that an autonomous system is a collection of networks under a single technical administration. IGPs operate within an autonomous system, and BGP (specifically BGP4) is used between autonomous systems on the Internet.


IANA is the organization responsible for allocating autonomous system numbers. Specifically, the American Registry for Internet Numbers (ARIN) has the jurisdiction to assign numbers for the Americas, the Caribbean, and Africa. Réseaux IP Européennes Network Information Center (RIPENIC) administers autonomous system numbers for Europe, and the Asia Pacific Network Information Center (APNIC) administers the numbers for the Asia Pacific region.

Autonomous system numbers are 16-bit numbers ranging from 1 to 65535. RFC 1930 provides guidelines for the use of autonomous system numbers. A range of autonomous system numbers, 64512 through 65535, is reserved for private use, much like private IP addresses. The autonomous system numbers used in this course are all in the private range to avoid publishing autonomous system numbers belonging to organizations.

Note Using an IANA-assigned autonomous system number rather than a private autonomous system number is necessary only if your organization plans to use an EGP such as BGP and connect to a public network such as the Internet.

BGP Routing Between Autonomous Systems

- BGP is used to provide an interdomain routing system.
- BGP guarantees the exchange of loop-free routing information.
- BGP works differently than IGP.
 - BGP is a policy-based routing protocol.
 - Control traffic flow using multiple BGP path attributes.



The main goal of BGP is to provide an interdomain routing system that guarantees loop-free exchange of routing information between autonomous systems. Routers exchange information about paths to destination networks.

BGP is a successor of EGP, which was developed to isolate networks from each other as the Internet grew.

There are many RFCs relating to BGP4, which is the current version of BGP. These RFCs include 1772, 1773, 1774, 1930, 1966, 1997, 1998, 2042, 2385, 2439, 2545, 2547, 2796, 2858, 2918, 3065, 3107, 3392, 4223, and 4271.

BGP4 has many enhancements over earlier protocols. The Internet uses BGP4 extensively to connect ISPs and to connect enterprises to ISPs.

BGP4 and its extensions are the only acceptable versions of BGP available for use on the public-based Internet. BGP4 carries a network mask for each advertised network and supports both variable-length subnet masking (VLSM) and classless interdomain routing (CIDR). BGP4 predecessors did not support these capabilities, which are currently mandatory on the Internet.

When using CIDR on a core router for a major ISP, the IP routing table, which is composed mostly of BGP routes, has more than 175,000 CIDR blocks. Not using CIDR at the Internet level would cause the IP routing table to have more than 2,000,000 entries. Using BGP4 and CIDR prevents the Internet routing table from becoming too large to interconnect millions of users.

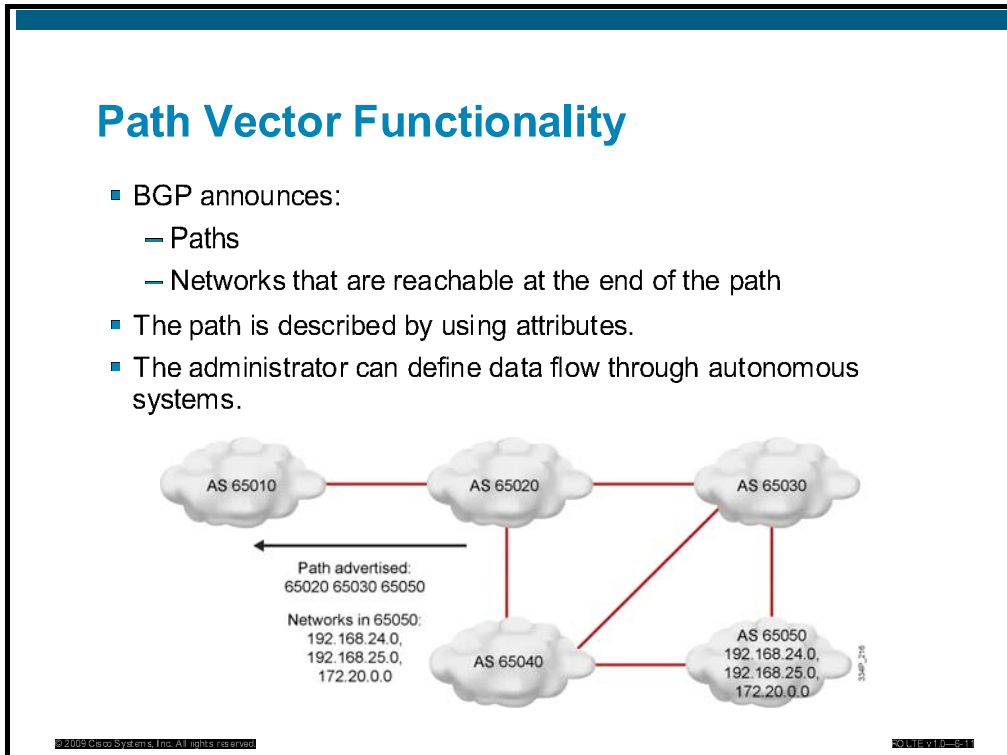
Comparison with IGP

BGP works differently than IGPs. An internal routing protocol looks for the quickest path from one point in a corporate network to another based upon certain metrics. RIP uses hop counts that attempt to cross the least number of Layer 3 devices to reach the destination network. OSPF and EIGRP look for the best speed available according to the bandwidth statement on the interface. All internal routing protocols consider the path cost to get somewhere.

In contrast, BGP, an external routing protocol, does not consider speed to determine the best path. Instead, BGP is a policy-based routing protocol that allows an autonomous system to control traffic flow using multiple BGP path attributes. BGP allows a provider to use all its bandwidth by manipulating these path attributes.

Path Vector Functionality

This topic describes how BGP uses path vector functionality.



Internal routing protocols announce a list of networks and the metrics to get to each network. In contrast, BGP routers exchange network reachability information, called “path vectors,” made up of path attributes (similar to metrics). The path vector information includes a list of the full path of BGP autonomous system numbers (hop by hop) necessary to reach a destination network and the networks that are reachable at the end of the path. Other attributes include the IP address to get to the next autonomous system (the next-hop attribute) and an indication of how the networks at the end of the path were introduced into BGP (the origin code attribute).


This autonomous system path information is useful to construct a graph of loop-free autonomous systems and is used to identify routing policies so that restrictions on routing behavior can be enforced based on the autonomous system path.

The autonomous system path is always loop-free. A router running BGP does not accept a routing update that already includes the router autonomous system number in the path list because the update has already passed through its autonomous system, and accepting it again would result in a routing loop.

An administrator can define policies or rules about how data will flow through the autonomous systems.

BGP Routing Policies

- BGP can support any policy conforming to the hop-by-hop (AS-by-AS) routing paradigm.



BGP allows routing policy decisions at the autonomous system level to be enforced. These policies can be implemented for all networks owned by an autonomous system, for a certain CIDR block of network numbers (prefixes), or for individual networks or subnetworks.

BGP specifies that a BGP router can advertise to neighboring autonomous systems only those routes that it uses itself. This rule reflects the hop-by-hop routing paradigm that the Internet generally uses.

The hop-by-hop routing paradigm does not support all possible policies. For example, BGP does not enable one autonomous system to send traffic to a neighboring autonomous system, intending that the traffic take a different route from that taken by traffic that originates in that neighboring autonomous system. In other words, how a neighboring autonomous system routes traffic cannot be influenced, but how traffic gets to a neighboring autonomous system can be influenced. However, BGP supports any policy that conforms to the hop-by-hop routing paradigm.

Because the Internet currently uses the hop-by-hop routing paradigm only, and because BGP can support any policy that conforms to that paradigm, BGP is highly applicable as an inter-autonomous system routing protocol.

Example: BGP Routing Policies

For example, in the figure above, the following paths are possible for AS 65010 to reach networks in AS 65060 through AS 65020:

- 65020 65030 65060
- 65020 65050 65060
- 65020 65030 65050 65070 65060
- 65020 65050 65030 65060
- 65020 65050 65070 65060

AS 65010 does not see all these possibilities.

AS 65020 advertises to AS 65010 only its best path of 65020 65030 65060, the same way that IGP's announce only their best least-cost routes. This path is the only path through AS 65020 that AS 65010 sees. All packets that are destined for 65060 through 65020 will take this path.

Even though other paths exist, AS 65010 can only use what AS 65020 advertises for the networks in AS 65060. The autonomous system path that is advertised, 65020 65030 65060, is the autonomous system-by-autonomous system (hop-by-hop) path that AS 65020 will use to reach the networks in AS 65060. AS 65020 will not announce another path, such as 65020 65050 65030 65060, because it did not choose that as the best path based on the BGP routing policy in AS 65020.

AS 65010 will not learn about the second-best path or any other paths from AS 65020 unless the best path of AS 65020 becomes unavailable.

Even if AS 65010 were aware of another path through AS 65020 and wanted to use it, AS 65020 would not route packets along that other path because AS 65020 selected 65020 65030 65060 as its best path and all AS 65020 routers will use that path as a matter of BGP policy. BGP does not let one autonomous system send traffic to a neighboring autonomous system, intending that the traffic take a different route from that taken by traffic originating in the neighboring autonomous system.

To reach the networks in AS 65060, AS 65010 can choose to use AS 65020 or it can choose to go through the path that AS 65040 is advertising. AS 65010 selects the best path to take based on its own BGP routing policies.

Features of BGP

This topic describes the features of BGP in terms of deployment, enhancements over other distance vector routing protocols, and database types.

Features of BGP

BGP is a path vector protocol with the following properties:

- Reliable updates: BGP runs on top of TCP (port 179)
- Incremental, triggered updates only
- Periodic keepalive messages to verify TCP connectivity
- Rich metrics (called path vectors or attributes)
- Designed to scale to huge internetworks (for example, the Internet)

It has enhancements over distance vector protocols.

© 2009 Cisco Systems, Inc. All rights reserved.

6-32

BGP is categorized as an advanced distance vector protocol, but it is actually a path vector protocol. BGP is very different from standard distance vector protocols like RIP.

BGP uses TCP as its transport protocol, which provides reliable connection-oriented delivery. BGP assumes that its communication is reliable; therefore, it does not have to implement retransmission or error recovery mechanisms. BGP uses TCP port 179. Two routers using BGP form a TCP connection with one another and exchange messages to open and confirm the connection parameters. These two BGP routers are called “peer routers,” or “neighbors.”

After the connection is made, BGP peers exchange full routing tables. However, because the connection is reliable, BGP peers subsequently send only changes (incremental, or triggered, updates) after that. Reliable links do not require periodic routing updates; therefore, routers use triggered updates instead. BGP sends keepalive messages, similar to the hello messages sent by OSPF, IS-IS, and EIGRP.

BGP is the only IP routing protocol to use TCP as its transport layer. OSPF, IGRP, and EIGRP reside directly above the IP layer, and RIP version 1 (RIPv1) and RIP version 2 (RIPv2) use User Datagram Protocol (UDP) for their transport layer.

OSPF and EIGRP have their own internal function to ensure that update packets are explicitly acknowledged. These protocols use a one-for-one window so that if either OSPF or EIGRP has multiple packets to send, the next packet cannot be sent until OSPF or EIGRP receive an acknowledgment from the first update packet. This process can be very inefficient and cause latency issues if thousands of update packets must be exchanged over relatively slow serial links. OSPF and EIGRP rarely have thousands of update packets to send. EIGRP can hold more than 100 networks in one EIGRP update packet, so 100 EIGRP update packets can hold up to 10,000 networks, and most organizations do not have 10,000 subnets in their corporations.

BGP, on the other hand, has more than 175,000 networks (and growing) on the Internet to advertise, and it uses TCP to handle the acknowledgment function. TCP uses a dynamic window, which allows 65,536 bytes to be outstanding before it stops and waits for an acknowledgment. For example, if 1000-byte packets are being sent, there would need to be 65 packets that have not been acknowledged for BGP to stop and wait for an acknowledgment when using the maximum window size.

TCP is designed to use a sliding window, where the receiver will acknowledge at the halfway point of the sending window. This method allows any TCP application, such as BGP, to continue to stream packets without having to stop and wait, as would be required with OSPF or EIGRP.

When to Use BGP

- BGP should be used if one of the following is true:
 - An AS is a transit AS
 - An AS is multihomed
 - Inter-AS routing policy must be manipulated
- BGP should not be used if one of the following is true:
 - Single-homed AS
 - Insufficient memory and processor resources to handle BGP routing
 - Insufficient understanding of route filtering and BGP path selection process

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-10

BGP allows ISPs to communicate and exchange packets. These ISPs have multiple connections to each other and agreements to exchange updates. BGP is used to implement the agreements between two or more autonomous systems.

Improper controlling and filtering of BGP updates can potentially allow an outside autonomous system to affect the traffic flow to your autonomous system. It is important to know how BGP operates and how to configure it properly to prevent this situation.

For example, if a customer is connected to ISP A and ISP B (for redundancy), the goal is to implement a routing policy to ensure that ISP A does not send traffic to ISP B via the autonomous system of the customer. The customer does not want to waste valuable resources and bandwidth within its autonomous system to route traffic for its ISPs, but will want to be able to receive traffic destined to its autonomous system through each ISP.

BGP should be used in the following cases:

- If an AS is a transit AS, where packets transit the AS to reach other autonomous systems
- If an AS is multihomed with multiple connections to other autonomous systems
- If inter-AS routing policy must be manipulated, where path selection for traffic entering and leaving the AS must be influenced

BGP is not always an appropriate solution to interconnect autonomous systems. There are several cases where BGP should not be used:

- When the network is single-homed AS with a single connection to the Internet or other AS. An AS with one exit path uses the default route as the most appropriate solution. BGP would unnecessarily use router CPU resources and memory.
- When there are insufficient memory and processor resources in the network edge router to handle BGP routing.
- When there is insufficient understanding of route filtering and the BGP path selection process.
- If the routing policy implemented in an autonomous system is consistent with the policy in the ISP autonomous system, it is not necessary or desirable to configure BGP in that autonomous system.

BGP Databases

- BGP neighbor table
 - List of BGP neighbors
- BGP table
 - List of all networks learned from each BGP neighbor
 - Multiple paths to same destination network can be present
 - Each path is associated with BGP attributes
- IP routing table (forwarding database)
 - List of best paths to destination networks used to forward traffic

A router running BGP keeps its own tables to store BGP information that it receives from and sends to other routers, including a neighbor table, a BGP table (also called a forwarding database or topology database), and an IP routing table.

For BGP to establish an adjacency, it must be explicitly configured for each neighbor. BGP forms a TCP relationship with each of the configured neighbors and keeps track of the state of these relationships by periodically sending a BGP/TCP keepalive message.

Note The BGP sends BGP/TCP keepalives by default every 60 seconds.

After establishing an adjacency, the neighbors exchange the BGP routes that are in their IP routing table. Each router collects these routes from each neighbor that successfully establishes an adjacency and then places the routes in its BGP forwarding database. All routes that have been learned from each neighbor are placed into the BGP forwarding database. The best routes for each network are selected from the BGP forwarding database using the BGP route selection process and are then offered to the IP routing table.

Each router compares the offered BGP routes to any other possible paths to those networks, and the best route, based on administrative distance, is installed in the IP routing table.

EBGP routes (BGP routes learned from an external autonomous system) have an administrative distance of 20. IBGP routes (BGP routes learned from within the autonomous system) have an administrative distance of 200.

BGP Message Types

BGP defines the following message types:

- Open, which includes hold time and BGP router ID
- Keepalive
- Update
 - Information for one path only (could be to multiple networks)
 - Includes path attributes and networks
- Notification
 - When an error is detected
 - BGP connection closed after message is sent

The four BGP message types are open, keepalive, update, and notification.

After a TCP connection is established, the first message sent by each side is an open message. If the open message is acceptable, the side that receives the message sends a keepalive message confirming the open message. After the receiving side confirms the open message and establishes the BGP connection, the BGP peers can exchange any update, keepalive, and notification messages.

BGP peers initially exchange their full BGP routing tables. Incremental updates are sent only after topology changes in the network occur. BGP peers send keepalive messages to ensure that the connection between the BGP peers still exists, and send notification packets in response to errors or special conditions.

Here are more details about the different types of BGP messages:

- **Open message:** An open message includes the following information:
 - **Version number:** The suggested version number. The highest common version that both routers support is used. Most BGP implementations today use BGP4.
 - **AS number:** The AS number of the local router. The peer router verifies this information. If it is not the AS number that is expected, the BGP session is ended.
 - **Hold time:** Maximum number of seconds that can elapse between the successive keepalive and update messages from the sender. On receipt of an open message, the router calculates the value of the hold timer by using whichever is smaller: its configured hold time or the hold time that was received in the open message.
 - **BGP router ID:** This 32-bit field indicates the BGP ID of the sender. The BGP ID is an IP address that is assigned to that router, and it is determined at startup. The BGP router ID is chosen in the same way that the OSPF router ID is chosen—it is the highest active IP address on the router unless a loopback interface with an IP address exists. In this case, the router ID is the highest loopback IP address. The router ID can also be statically configured.

- **Optional parameters:** These parameters are type, length, and value (TLV) encoded. An example of an optional parameter is session authentication.
- **Keepalive message:** BGP keepalive messages are exchanged between BGP peers often enough to keep the hold timer from expiring. If the negotiated hold-time interval is 0, then periodic keepalive messages are not sent. A keepalive message consists of only a message header.
- **Update message:** A BGP update message has information on one path only; multiple paths require multiple update messages. All the attributes in the update message refer to that path, and the networks are those that can be reached through that path. An update message can include the following fields:
 - **Withdrawn routes:** This list displays IP address prefixes for routes that are withdrawn from service, if any.
 - **Path attributes:** These attributes include the AS path, origin, local preference, and so on (as described later in this module). Each path attribute includes the attribute TLV. The attribute type consists of the attribute flags, followed by the attribute type code.
 - **Network-layer reachability information:** This field contains a list of IP address prefixes that are reachable by this path.
- **Notification message:** A BGP notification message is sent when an error condition is detected; the BGP connection is closed immediately after this is sent. Notification messages include an error code, an error subcode, and data that are related to the error.

Summary

This topic summarizes the key points that were discussed in this lesson.

Summary

- BGP is typically used for interdomain routing.
- Three common ways to perform multihoming with BGP are as follows:
 - Each ISP passes only a default route.
 - Each ISP passes only a default route and specific provider-owned routes.
 - Each ISP passes all routes.
- BGP is the external routing protocol used between autonomous systems. Forwarding is based on policies and not on best path.
- BGP routers exchange network reachability information called path vectors, made up of path attributes.
- A router running BGP keeps its own tables to store BGP information that it receives from and sends to other routers, including a neighbor table, a BGP table, and an IP routing table.

Lesson 3

Comparing the Functions and Uses of EBGP and IBGP

Overview

This lesson explains important terminology that is used in establishing Border Gateway Protocol (BGP) peering relationships.

This lesson describes and defines External BGP (EBGP) and Internal BGP (IBGP) neighbors as well as the requirements for establishing peering relationships. In addition, this lesson examines the difference between an interior gateway protocol (IGP) and BGP and explains the reason for having all routers in the transit path within an autonomous system (AS) running IBGP.

Understanding the relationship between various types of BGP routers and the common terminology that is used when you are discussing these routers is necessary for troubleshooting connectivity issues between BGP neighbors. The following terms are explained: BGP speaker, BGP router, BGP neighbor, and BGP peer.

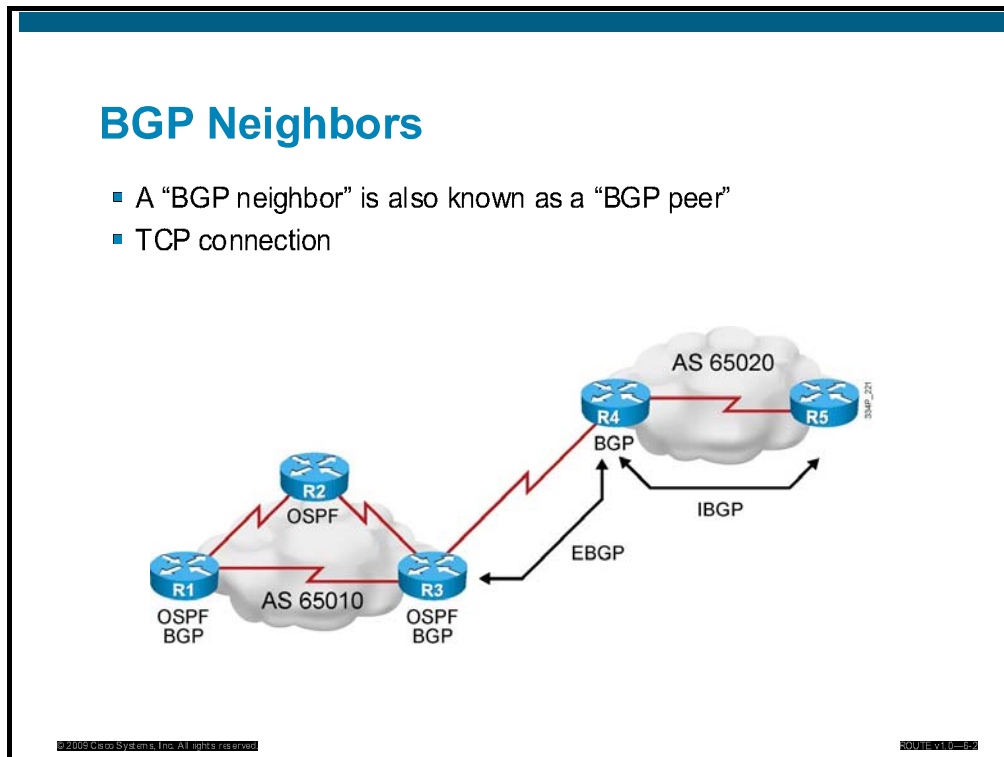
Objectives

Upon completing this lesson, you will be able to compare and contrast the requirements for establishing EBGP-to-IBGP neighbor relationships. This ability includes being able to meet these objectives:

- Define BGP neighbor relationships.
- Establish External BGP neighbor relationships.
- Establish Internal BGP neighbor relationships.

BGP Neighbor Relationships

This topic describes terms used for BGP routers and their relationships.



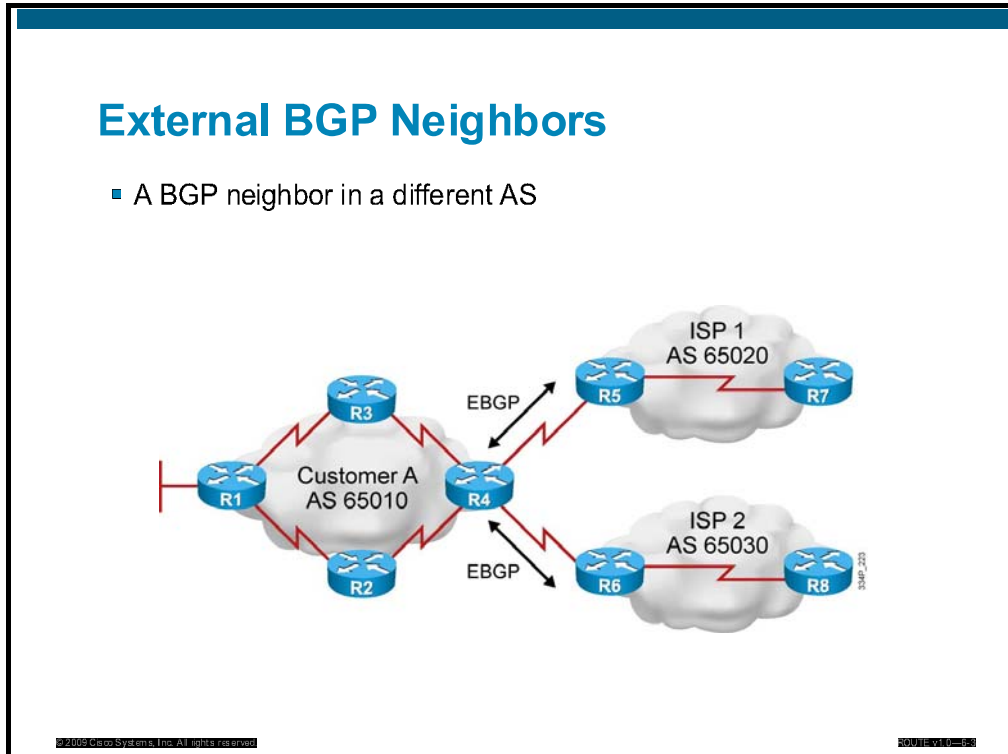
No one router can handle communications with all the routers that run BGP. There are tens of thousands of routers that run BGP and these routers are connected to the Internet, representing more than 21,000 autonomous systems (ASs).

A BGP router forms a direct neighbor relationship with a limited number of other BGP routers. Through these BGP neighbors, a BGP router learns of the paths that are available through the Internet to reach any advertised network. Any router that runs BGP is known as a “BGP speaker.” The term “BGP peer” has a specific meaning—it is a BGP speaker that is configured to form a neighbor relationship with another BGP speaker for the purpose of directly exchanging BGP routing information with each other. A BGP speaker has a limited number of BGP neighbors with which it peers and forms a TCP-based relationship.

BGP peers are also known as “BGP neighbors” and can be either internal or external to the AS. Internal BGP peer or IBGP peer typically forms the neighbor relationship inside the ISP network. External BGP peer or EBGP peer forms the neighbor relationship between the enterprise network and the ISP. External neighbors, as well as internal neighbors, require a TCP connection to be established. The processing of the BGP routes and its attributes is typically different, as is the connection type. External neighbors are directly connected. A network assigned to the link between neighbors, as well as the next-hop IP address, must be reachable for routing purposes.

Establishing External BGP Neighbor Relationships

This topic describes the requirements for establishing an External BGP (EBGP) neighbor relationship.




Recall that when BGP is running between routers in different autonomous systems, it is called EBGP. By default, routers running EBGP are directly connected to each other.

An EBGP neighbor is a router outside a home AS. An enterprise network can have a connection to one or several ISPs, and the ISPs themselves might be connected to several other ISPs as well. For each such connection between different autonomous systems, there is an EBGP session required between EBGP neighboring routers. EBGP neighbors are directly connected and they establish a TCP session prior exchanging BGP updates.

When multiple different autonomous systems are connected to each other and an enterprise network is connected to multiple ISPs, BGP runs between the ISPs and the enterprise network. In the figure above, the Customer A network is connected to two ISPs and an EBGP session is established between routers R3 and R5 as well as between routers R4 and R6. Routers establish neighbor relationships and exchange BGP routing updates with one another. In the figure above, Customer A routers learn the paths to the external autonomous systems from their respective EBGP neighbors.

Requirements for EBGP

- Different AS number
- Defined neighbors
- Reachability



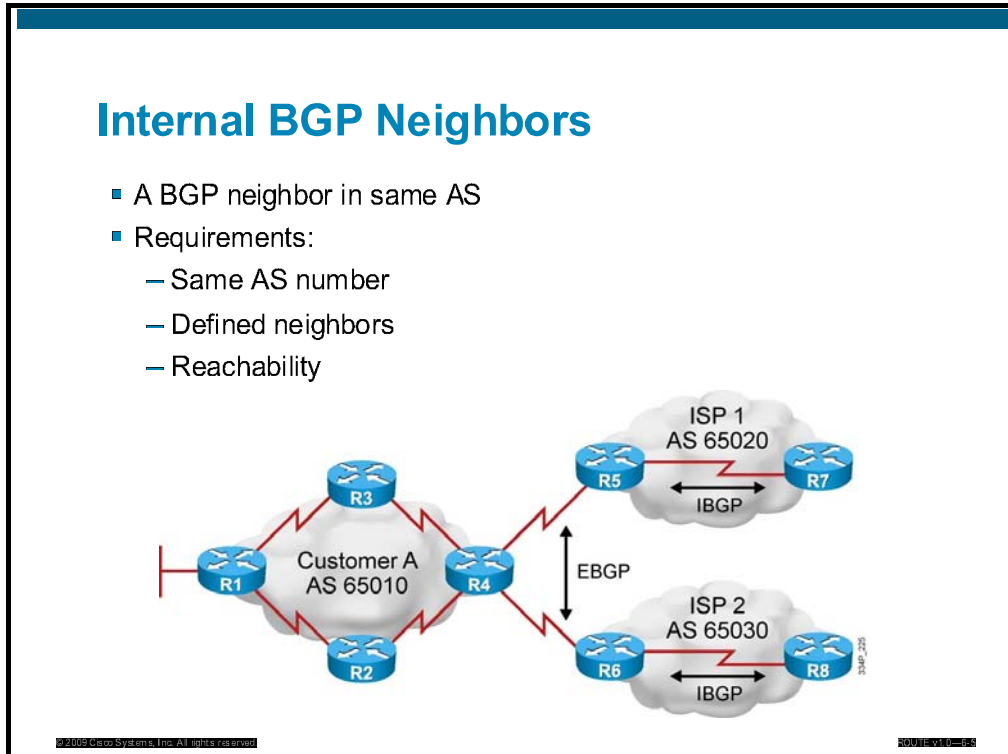
There are several different requirements for establishing an EBGP neighbor relationship:

- **Different AS number:** EBGP neighbors must reside in different autonomous systems to be able to form an EBGP relationship.
- **Defined neighbors:** A TCP session must be established prior to starting BGP routing update exchanges.
- **Reachability:** EBGP neighbors must be directly connected and IP addresses on that link must be reachable inside each AS.

IGP is the routing protocol that runs inside an AS. An IGP is not run between the EBGP neighbors residing in different autonomous systems. For two routers to exchange BGP routing updates, the TCP-reliable transport layer on each side must successfully pass the TCP three-way handshake before a BGP session can be established. Therefore, the IP address used in the BGP neighbor command must be reachable without using an IGP, which can be accomplished by pointing at an address that is reachable through a directly connected network or by using static routes to that IP address. Generally, the neighbor address that is used is the address on a network that is directly connected.

Establishing Internal BGP Neighbor Relationships

This topic describes the requirements for establishing an Internal BGP (IBGP) neighbor relationship.



Recall that BGP that runs between routers within the same AS is called IBGP. IBGP runs within an AS to exchange BGP information so that all BGP speakers have the same BGP routing information about outside autonomous systems.

When multiple routers in an AS are running BGP, they exchange BGP routing updates with one another. In the figure above, routers R3 and R4 learn the paths to the external autonomous systems from their respective EBGP neighbors (R5 and R6). If the link between R3 and R4 goes down, R3 must learn new routes to the external autonomous systems. Other BGP routers within AS 65010 that were using R3 to get to external networks must also be informed that the path through R3 is not available. Those BGP routers within AS 65010 need to have the alternate paths through router R4 in their BGP forwarding database. IBGP is typically established between the customer-edge (CE) routers inside the enterprise network autonomous system as well as between the routers inside each ISP.

There are several different requirements for establishing an IBGP neighbor relationship:

- **Same AS number:** IBGP neighbors must reside in the same AS to be able to form an IBGP relationship.
- **Defined neighbors:** A TCP session must be established between neighbors prior to exchanging BGP routing updates.
- **Reachability:** IBGP neighbors must be reachable, thus IGP typically runs inside an AS.

Summary

This topic summarizes the key points that were discussed in this lesson.

Summary

- Every BGP router establishes a neighbor relationship with other BGP neighbors.
- An External BGP (EBGP) neighbor relationship is established between the routers in different autonomous systems (AS).
- An Internal BGP (IBGP) neighbor relationship is established between routers in the same AS.

Lesson 4

Configuring and Verifying Basic BGP Operations

Overview

This lesson presents the commands and configuration examples to configure Border Gateway Protocol (BGP) properly.

After a successful configuration, BGP will be able to establish a neighbor relationship, set the next-hop address, set the source IP address of a BGP update, and announce networks to other BGP routers.

A router running BGP goes through several neighbor states through which BGP progresses to establish a BGP session, and offers hints for troubleshooting BGP because the session can be stuck in the stuck-in-active (SIA) or idle state. External BGP (EBGP) and Internal BGP (IBGP) configuration is explained, as well as authentication between neighbors.

This lesson also shows how to use the **show** and **debug** commands for troubleshooting BGP, providing a baseline for troubleshooting. A thorough understanding of this material is necessary to use BGP.

Objectives

Upon completing this lesson, you will be able to implement BGP operations for ISP connections in an enterprise network. This ability includes being able to meet these objectives:


- Identify specifications for implementing BGP.
- Establish internal and external BGP neighbor relationships.
- Determine how to shut down a BGP neighbor.
- Understand BGP configuration considerations.
- Identify BGP neighbor states.
- Implement authentication in BGP.
- Create an example of activating basic BGP.
- Configure BGP verification.

Specifications for Implementing BGP

This topic describes the types of information contained in a typical implementation plan.

Planning for BGP

- Define network requirements
- Define internal connectivity
- Define external connectivity to ISP
- Gather required parameters




The diagram illustrates a network topology for BGP planning. It features three routers labeled R1, R2, and R3. R1 and R2 are connected to a cloud, and R2 and R3 are connected to another cloud. The connections are represented by red lines with lightning bolt symbols. The routers are blue with white text. The clouds are white with a blue outline. The diagram is enclosed in a black border with a blue header bar at the top.

Prior to BGP configuration, a network administrator must define the network requirements including the internal connectivity for IBGP design and configuration, as well as the external connectivity to the ISP for EBGp design and configuration.

The next step is to gather all of the parameters needed to provide enough details for a network operator to start the BGP configuration.

Requirements for Basic BGP Configuration

- AS numbers
- Neighbors (IP addresses)
- Networks to be advertised



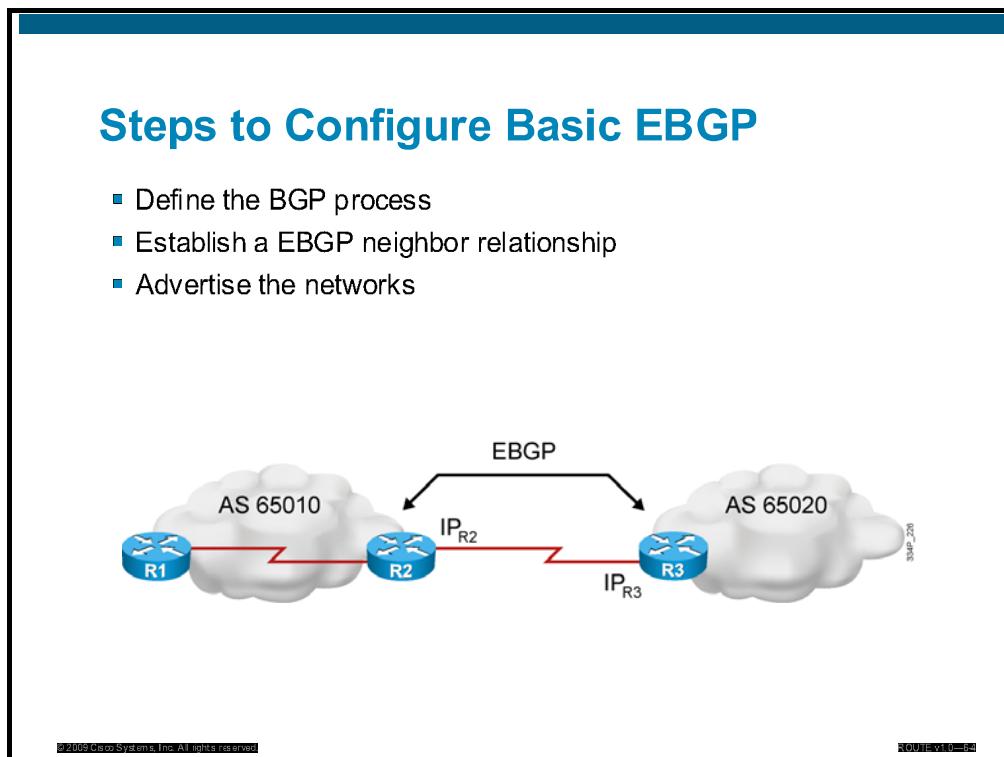
The requirements to configure basic BGP include the following details:

- Autonomous system (AS) numbers (your own and all remote AS numbers)
- All of the neighbors (peers) involved in BGP and IP addressing used among the BGP neighbors
- Networks that need to be advertised into BGP

A typical BGP configuration involves configuring BGP between a customer network and an ISP. This process is called External BGP (EBGP). Many times Internal BGP (IBGP) is required, as well as all the collected details for a complete configuration.

Establishing Internal and External BGP Neighbors

This topic describes the process of activating a BGP session for external and internal neighboring routers.




Basic EBGP configuration requires three main steps:

1. Define the BGP process.
2. Establish the neighbor relationship.
3. Advertise the networks into BGP.

In order to perform the steps above, information is needed about the neighbors, which AS numbers are used, which IP address is used as the IP address on a remote router (neighbor), and which network will be advertised.

Define BGP Process and Activate EBGP Session



R2 (config) #

```
router bgp 65010
```

- Define the BGP process locally with a local AS number.

R2 (config-router) #

```
neighbor 10.1.1.2 remote-as 65020
```

Activate EBGP session to the neighbor

- Remote router IP Address and AS number

The syntax of the basic BGP configuration commands is similar to the syntax for configuring internal routing protocols. However, there are significant differences in how BGP functions.

Use the **router bgp 65010** command to notify the router that any subsequent subcommands belong to this routing process. This command also identifies the local autonomous system (AS) in which this router belongs. AS 65010 is used as an example. The router needs to be informed of the AS number so it can determine whether the BGP neighbors to be configured next are IBGP or EBGP neighbors.

The command enters router configuration mode. The **router bgp** command alone cannot activate BGP on a router. At least one subcommand must be entered under the **router bgp** command to activate the BGP process on the router.

If the router is placed in AS 65010 and an attempt is then made to configure with a new **router bgp 65020** command, the router indicates that it is currently configured for AS 65010. To establish a connection to another AS, insert the AS number with a **neighbor** command in addition to the **router bgp** command so that the router can properly identify the relationship between the neighboring router and itself.

Note Only one instance of BGP can be configured on the router at a single time.

Use the **neighbor 10.1.1.2 remote-as 65020** command to activate an EBGP session for a neighboring router.

This command identifies a peer router with which the local router will establish a session and is mandatory for the establishment of each neighboring router relationship.


The address that is used in this command is the destination address for all BGP packets going to the neighboring router. For BGP to pass BGP routing information, this address must be reachable because BGP attempts to establish a TCP session and exchange BGP updates with the device at this IP address. An EBGP session is typically established between directly connected neighbors.

The AS number that is a part of this command is used to identify whether this neighbor is an EBGP neighbor or an IBGP neighbor. If the AS number is the same as the AS number for this router, that neighbor is an IBGP neighbor and the IP address listed in this neighbor command does not have to be directly connected. If the AS number is different from the AS number for this router, this neighbor is an EBGP neighbor and the address in this neighbor command must be directly connected by default. In the figure above, the AS is different and an EBGP neighbor relationship is established.

For more details about the **neighbor remote-as** and **router bgp** commands, please check the Cisco IOS IP Routing Protocols Command Reference on the following link:

http://www.cisco.com/en/US/docs/ios/iproute/command/reference/irp_book.html

Advertise Networks



Option 1:

R2 (config-router) #

```
network 10.2.2.0 mask 255.255.255.0
network 10.4.4.0 mask 255.255.255.0
```

- Configure the local networks to be advertised and include them in BGP

Option 2:

- Redistribution from IGP to BGP

Two options exist when advertising networks into the BGP. The first option is using the **network** command to define the networks required. The second option is the redistribution of IGP routes into the BGP routing process. Redistribution was described in a separate module. The command is explained in the following pages.

Use the **network 10.2.2.0 mask 255.255.255.0** command to permit BGP to advertise the 10.2.2.0 if it is present in the IP routing table.

The **network** command determines networks that the router originates. This concept is different from using the **network** command when configuring IGP. Unlike with IGP, the **network** command does not start BGP on specific interfaces; rather, it indicates to BGP which networks it should originate from this router.

The **mask** parameter indicates that BGP version 4 (BGP4) can handle subnetting and supernetting. The list of **network** commands must include all networks in the AS that are to be advertised, not just those that are locally connected to the router.

Prior to Cisco IOS Software Release 12.0, there was a limit of 200 network commands per BGP router. This limit has been removed. The resources of the router, such as the configured NVRAM or RAM, determine the maximum number of **network** commands that can be used.

The sole purpose of the **network** command is to notify BGP of the networks that are to be advertised. Without the **mask** option, this command announces only the classful network number. At least one subnet of the specified major network must be present in the IP routing table to allow BGP to start announcing the classful network as a BGP route.

However, if a network mask option is specified, an exact match to the network (both address and mask) must exist in the routing table. Before BGP announces a route, it checks to see whether it can reach it.

For more details about the **network** command, please check the Cisco IOS IP Routing Protocols Command Reference on the following link:

http://www.cisco.com/en/US/docs/ios/iproute/command/reference/irp_book.html

BGP network Command Details

R2(config-router) #

```
network 192.168.1.1 mask 255.255.255.0
```

- The router looks for 192.168.1.1/24 in the routing table, but cannot find it, so it will not announce anything.

R2(config-router) #

```
network 192.168.0.0 mask 255.255.0.0
```

- The router looks for 192.168.0.0/16 in the routing table.
- If the exact route is not in the table, you can add a static route to null0 so that the route can be announced.

R2(config-router) #

```
network 192.168.1.0
```

- The router looks for a C class 192.168.1.0 network in the routing table.

For example, if the command **network 192.168.1.1 mask 255.255.255.0** is misconfigured, BGP looks for exactly 192.168.1.1/24 in the routing table. It may find 192.168.1.0/24 or 192.168.1.1/32; however, it never finds 192.168.1.1/24. Because the routing table does not contain a specific match to the network, BGP does not announce the 192.168.1.1/24 network to any neighbors.

As another example, if **network 192.168.0.0 mask 255.255.0.0** is specified to advertise a classless interdomain routing (CIDR) block, BGP looks for 192.168.0.0/16 in the routing table. It may find 192.168.1.0/24 or 192.168.1.1/32; however, if it never finds 192.168.0.0/16, BGP does not announce the 192.168.0.0/16 network to any neighbors. In this case, the following static route can be configured toward the null interface so that BGP can find an exact match in the routing table:


```
ip route 192.168.0.0 255.255.0.0 null0
```

After finding an exact match in the routing table, BGP announces the 192.168.0.0/16 network to any neighbors.

If the **network 192.168.1.0** command is specified without the mask keyword, BGP looks for C class 192.168.1.0/24 in the routing table.

Note The BGP **auto-summary** router configuration command determines how BGP handles redistributed routes. With BGP summarization enabled (with **auto-summary**), all redistributed subnets are summarized to their classful boundaries in the BGP table. When disabled (with **no auto-summary**), all redistributed subnets are present in their original form in the BGP table, so only those subnets would be advertised. In Cisco IOS Software Release 12.2(8)T, the default behavior of the **auto-summary** command was changed to disabled (**no auto-summary**); prior to that the default was enabled (**auto-summary**).

Basic EBGP Configuration



Option 1

```
R2#  
!  
<output omitted>  
!  
router bgp 65010  
  neighbor 10.1.1.2 remote-as 65020  
  network 10.2.2.0 mask 255.255.255.0  
  network 10.4.4.0 mask 255.255.255.0  
!
```

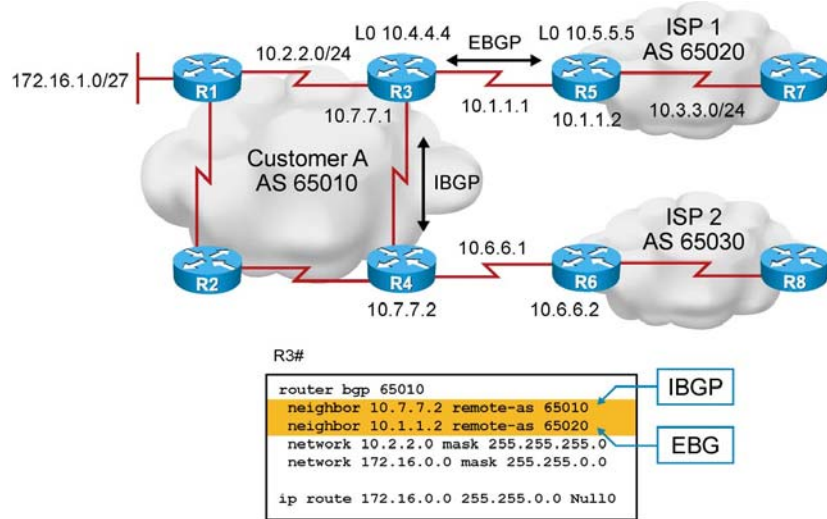
Option 2

```
R2#  
!  
<output omitted>  
router ospf 10  
  network 10.2.2.0 mask 255.255.255.0  
  network 10.4.4.0 mask 255.255.255.0  
!  
router bgp 65010  
  neighbor 10.1.1.2 remote-as 65020  
  redistribute ospf  
!
```

In option 1, router R2 is configured for EBGP configuration. The local AS number 65010 is used to enter BGP configuration mode. When defining the EBGP, the **neighbor remote-as** statement defines the relationship to the local AS number. The remote AS is different from the local AS because the neighbor is declared as an EBGP neighbor. Router R2 will establish a TCP session and start exchanging BGP updates. Router R2 is configured to advertise two local networks to its neighbor. The exact network entries (10.2.2.0/24 and 10.4.4.0/24) must exist inside the routing table in order to accept and advertise networks inside the BGP process.

In option 2, router R2 is configured for EBGP configuration, but its configuration to advertise the networks is different. The Open Shortest Path First (OSPF) protocol is used as an IGP on router R2, and networks 10.2.2.0 and 10.4.4.0 are advertised in OSPF. In order to advertise the networks into BGP, redistribution from OSPF to BGP is configured.

Basic IBGP and EBGP Configuration in the Customer A Network



Customer-edge routers (CE-routers) are connected to provider edge (PE) routers. EBGP is configured between CE and PE routers because enterprise networks and ISPs have different AS numbers. If enterprise networks are dual-homed or multihomed, then typically two CE-routers are configured for EBGP. In order to provide reliable connectivity to one or more ISPs, CE-routers must exchange BGP routes as well. An IBGP session must be configured between two CE-routers.

The figure above shows typical BGP configuration of CE-router R2. The EBGP neighbor relationship to router R5 and IBGP neighbor relationship to router R6 is configured. An IBGP session is established to the neighbor with the same AS number as its local AS number (AS 65010). An EBGP neighbor relationship is established with the neighbor with a different AS number (AS 65020) and local AS number (AS 65010). Network commands are used to advertise two networks. Because network 172.16.0.0 is not directly connected, a static route to the Null0 interface is configured in order to advertise network 172.16.0.0 into BGP.

Shutting Down a BGP Neighbor

This topic describes the process of administratively shutting down and reenabling a BGP neighbor.

Shutting Down a BGP Neighbor

```
R2(config-router)#  
neighbor 10.1.1.2 shutdown
```

- Administratively brings down a BGP neighbor
- Used for maintenance/policy changes to prevent route flapping

```
R2(config-router)#  
no neighbor 10.1.1.2 shutdown
```

- Reenables a BGP neighbor that has been administratively shut down

© 2009 Cisco Systems, Inc. All rights reserved. C91115-10-08-01

Use the **neighbor 10.1.1.2 shutdown** commands to administratively shut down and reenabling a BGP neighbor.

If major policy changes are implemented for a neighboring router and multiple parameters are changed, the neighboring router must be administratively shut down, the changes must be implemented, and then the neighboring router must be brought back up with the **no neighbor 10.1.1.2 shutdown** command.

Using the **neighbor 10.1.1.2 shutdown** command not only terminates the session but removes all associated routing information as well.

For more details about the **neighbor shutdown** command, please check the Cisco IOS IP Routing Protocols Command Reference on the following link:

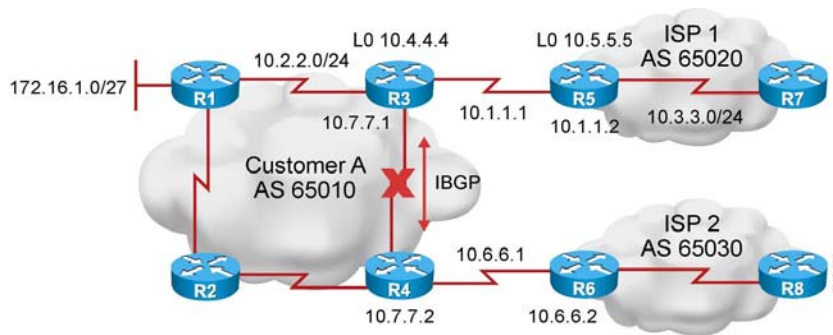
http://www.cisco.com/en/US/docs/ios/iproute/command/reference/irp_book.html

BGP Configuration Considerations

This topic describes what needs to be considered to correctly configure BGP.

IBGP Peering Issue

- An IBGP neighbor relationship is established.
- What happens if the link between R3 and R4 goes down?
- Which IP address should be used to establish an IBGP session?



To establish an IBGP session between routers R3 and R4, as shown in the figure above, which neighbor IP address should be used?

The problem is as follows: If router R3 uses the **neighbor 10.7.7.2 remote-as 65010** command, but router R4 is sending the BGP packets to router R3 via routers R2 and R1, the source IP address will be different.

When router R4 receives this BGP packet via routers R2 and R1, it will not recognize this BGP packet because a different IP address was not configured as a neighbor of router R4; therefore, the IBGP session between router R3 and router R4 cannot be established.

A solution to this problem is to establish the IBGP session using a loopback interface when there are multiple paths between the IBGP neighbors.

BGP Issues with Source IP Address

Create a BGP packet:

- The destination IP address defined by the neighbor statement
- The source IP address defined by the outbound interface

The source address of the received BGP packet is compared to list of neighbor statements:

- If a match is found in the list of neighbors, a relationship is established.
- If no match is found in the list of neighbors, the packet is ignored.

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-5-0


The BGP neighbor statement informs the router of the destination IP address for each update packet. The router must decide which IP address to use as the source IP address in the BGP routing update.

When a router creates a BGP packet for a neighbor, it checks the routing table for the destination network to reach that neighbor. The IP address of the outbound interface, as the routing table indicates, is used as the source IP address of the BGP packet.

When a BGP packet is received for a new BGP session, the source address of the packet is compared to the list of neighbor statements. This source IP address must match the address in the corresponding neighbor statement on the other router. Otherwise, the routers will not be BGP peers because they are not able to establish the BGP session.

IBGP Using Loopback Addresses

- A loopback interface can be used as the source and destination IP address of all BGP updates between neighbors.
- The **neighbor update-source** command is normally used only with IBGP neighbors.



```
R2#
router bgp 65020
neighbor 172.16.1.1 remote-as 65010
neighbor 3.3.3.3 remote-as 65020
neighbor 3.3.3.3 update-source Loopback0
!
router eigrp 1
network 10.0.0.0
network 2.0.0.0

R3#
router bgp 65020
neighbor 192.168.1.1 remote-as 65030
neighbor 2.2.2.2 remote-as 65020
neighbor 2.2.2.2 update-source Loopback0
!
router eigrp 1
network 10.0.0.0
network 3.0.0.0
```

Multiple paths can exist to reach each neighbor when peering with IBGP neighboring routers. If the BGP router is using a neighbor address that is assigned to a specific interface on another router, and that interface goes down, the router pointing to this address loses its BGP session with that neighbor.

If the router peers instead with the loopback interface of the other router, the loopback interface will always be available as long as the router itself does not fail. This peering arrangement adds resiliency to the IBGP sessions because the routers are not tied into a physical interface, which may go down for any number of reasons.

To peer with the loopback of another internal neighbor, the first router would point the neighbor statement at the loopback address of the other internal neighbor. Ensure that both routers have a route to the loopback address of the other neighbor in their routing table. Also ensure that both routers are announcing their loopback addresses into their local routing protocol.

The **update-source** option in the **neighbor** command overrides the default source IP address used for BGP packets. It is necessary to tell the router which IP address to use as the source address for all BGP packets if a loopback interface is to be used instead of the physical interface.

If the **update-source** option is not used in the **neighbor** command, an announcement packet going to a neighbor uses the IP address of the exiting interface as the source address for a packet.

When a router creates a packet, whether it is a routing update, a ping, or any other type of IP packet, the router does a lookup in the routing table for the destination address. The routing table lists the appropriate interface to get to the destination address. The address of this outbound interface is used as the source address of that packet by default.

Consider what would happen if a neighboring router uses the loopback interface address in its **neighbor** command for this router, but the **neighbor update-source** command is not used on this router. When the neighboring router receives an update packet and looks at the source address of the packet, it sees that it has no neighbor relationship with that source address, so it discards the packet.

BGP does not accept unsolicited updates; it must be aware of every neighboring router and have a neighbor statement for it.

In the figure above, router R2 has router R1 as an EBGP neighbor and router R3 as an IBGP neighbor. Routers R1 and R2 are directly connected. Router R3 is not directly connected to router R2, and in the figure above, Enhanced Interior Gateway Routing Protocol (EIGRP) is used in order to provide reachability between R2 and R3. Because the neighbor relationship between routers R2 and R3 is not tied to a physical interface, router R2 peers with the loopback interface on router R3 and uses its loopback address as the source IP address, and vice versa. If router R2 instead peered with 10.1.1.2 on router R3 and that interface went down, the BGP neighbor relationship would also go down.

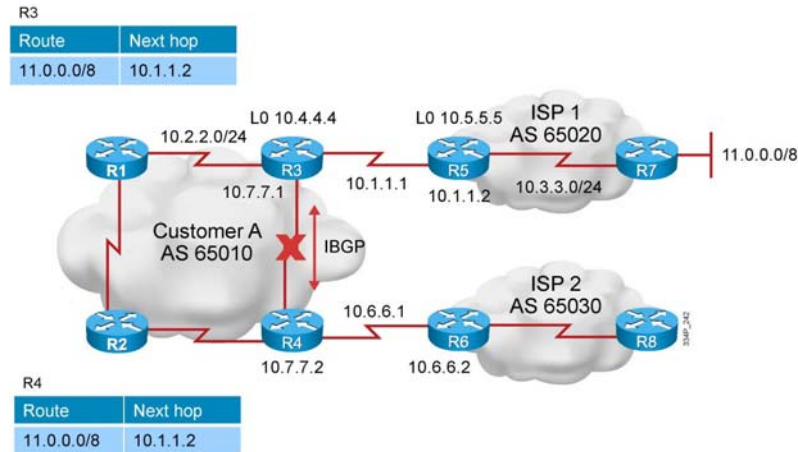
The **neighbor update-source** command should be used on both routers. If router R2 points to loopback address 3.3.3.3 of router R3, and router R3 points at loopback address 2.2.2.2 of router R2, and neither use the **neighbor update-source** command, the BGP session between these routers will not start.

For more details about **neighbor update-source** command, please check the Cisco IOS IP Routing Protocols Command Reference on the following link:

http://www.cisco.com/en/US/docs/ios/iproute/command/reference/irp_book.html

IBGP Next-Hop Behavior

- IBGP does not modify next hop.

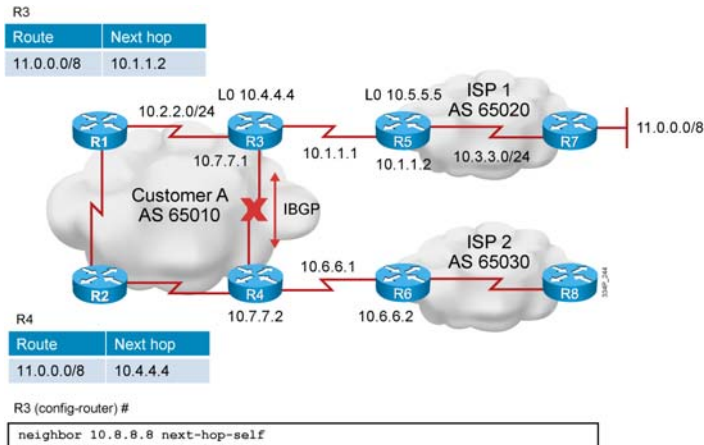


The way in which BGP establishes an IBGP relationship is very different from the way that IGP's behave. BGP is an AS-by-AS routing protocol and not a router-by-router routing protocol. Next hop is the IP address used to reach the next AS. The method that BGP uses to denote its next-hop address is also very different from the way that an IGP performs the same function. The default next hop is as follows:

- **EBGP:** IP address of the neighbor router sending the update
- **IBGP:** IP address advertised by EBGP should be carried in IBGP

In the figure above, the next hop for route 11.0.0.0/8 is always the IP address 10.1.1.2 advertised by the EBGP neighbor. In case a failure occurs on the link between router R3 and R4, and router R4 has no information regarding how to reach the address 10.1.1.2, it will not be able to forward packets toward destination network 11.0.0.0/8. This situation occurs even if the physical path exists (through routers R1 and R2 in the figure above).

BGP neighbor next-hop-self Command



- Forces all updates for neighbor R4 to be advertised with this router as the next hop—the same IP address as for the source of the BGP packet.

It is sometimes necessary to override the default next-hop behavior of a router and force it to advertise itself as the next-hop address for routes sent to a neighboring router.

The **neighbor next-hop-self** command forces BGP to use its own IP address as the next-hop address for each network that it advertises to its IBGP neighbor, rather than letting the protocol choose the next-hop address to use.

An internal protocol, such as Routing Information Protocol (RIP), EIGRP, or Open Shortest Path First Protocol (OSPF), always uses the source IP address of a routing update as the next-hop address for each network that is placed in the routing table. The **neighbor next-hop-self** command forces BGP use the source IP address of the update as the next-hop address for each advertised network.

Routers R3 and R4 are IBGP neighbors and router R3 will advertise all the updates toward router R4 with the next hop being the source IP address of the BGP packet going from router R3 to R4. The figure above shows that the next-hop address for the destination network 11.0.0.0 is now loopback interface 0 of router R3. It is reachable via routers R1 and R2 because all the routers in AS 65010 are running the IGP routing protocol providing reachability in the local autonomous system. Static routes can be used as well, but IGP is a more scalable solution.

For more details about **neighbor next-hop-self** command, please check the Cisco IOS IP Routing Protocols Command Reference on the following link:

http://www.cisco.com/en/US/docs/ios/iproute/command/reference/irp_book.html

Identifying BGP Neighbor States

This topic describes BGP neighbor states.

BGP States

When establishing a BGP session, BGP goes through the following states:

- 1 **Idle**: Router is searching the routing table to see whether a route exists to reach the neighbor.
- 2 **Connect**: Router found a route to the neighbor and has completed the three-way TCP handshake.
- 3 **Open sent**: Open message sent, with the parameters for the BGP session.
- 4 **Open confirm**: Router received an agreement on the parameters for establishing a session.
 - Alternatively, the router goes into **active** state if no response to open message
- 5 **Established**: Peering is established; routing begins.

After the TCP handshake is complete, the BGP application tries to set up a session with the neighbor. A number of steps must occur for the session to be established.

After the **neighbor** command is entered in BGP, BGP takes the IP address that is listed and checks the local routing table for a route to this address. At this point, BGP is in the “idle” state. If BGP does not find a route to the IP address, it stays in the idle state. If it finds a route, it goes to the “connect” state when the TCP handshaking synchronization-acknowledgment (SYN ACK) packet returns.

After the TCP connection is complete, BGP creates a BGP open packet and sends it to the neighbor. Once BGP sends this open packet, the BGP peering session changes to the “open sent” state. If there is no response for 5 seconds, the state changes to the “active” state.

If a response does come back in a timely manner, BGP goes to the “open confirm” state and starts scanning (evaluating) the routing table for the paths to send to the neighbor. When those paths have been found, BGP then goes to the “established” state and begins routing between the neighbors.

Note The states that two BGP routers are going through to establish a session can be observed using **debug** commands. In Cisco IOS Software Release 12.4, the **debug ip bgp ipv4 unicast** command can be used to see this process. In earlier Cisco IOS Software releases, the **debug ip bgp events** command gave similar output. Debugging uses up router resources and should be turned on only when necessary.

BGP Established and Idle States

- **Idle:** The router cannot find the address of the neighbor in the routing table.
 - **Solution:** Check for an IGP problem. Is the neighbor announcing the route?
- **Established:** Proper state for BGP operations.
 - Output of the **show ip bgp summary** command has a number in the state column indicating the number of routes learned from this neighbor.

© 2009 Cisco Systems, Inc. All rights reserved.

6-11-10-05-07

The idle state is an indication that the router does not know how to reach the IP address that is listed in the neighbor statement. The router is idle because of one of the following scenarios:

- It is waiting for a static route to that IP address or network to be configured.
- It is waiting for the local routing protocol (IGP) to learn about this network through an advertisement from another router.

The most common reason for a router to enter the idle state is that the neighbor is not announcing the IP address or the network toward which the neighbor statement of the router is pointing. Check these two conditions first to correct this problem:

1. Ensure that the neighbor announces the route in its local routing protocol (IGP).
2. Verify that an incorrect IP address has not been entered in the neighbor statement.

The established state is the desired state for the neighbor relationship. This state means that both routers have agreed to exchange BGP updates with one another and routing has begun.

Example: show ip bgp neighbors Command

```
R2#show ip bgp neighbors
BGP neighbor is 172.31.1.3, remote AS 65010, external link
BGP version 4, remote router ID 172.31.2.3
BGP state = Established, up for 00:19:10
Last read 00:00:10, last write 00:00:10, hold time is 180, keepalive
interval is 60 seconds
Neighbor capabilities:
  Route refresh: advertised and received(old & new)
  Address family IPv4 Unicast: advertised and received
Message statistics:
  InQ depth is 0
  OutQ depth is 0

      Sent      Rcvd
Opens:          7         7
Notifications: 0         0
Updates:       13        38
<output omitted>
```

Use the **show ip bgp neighbors** command to display information about the BGP connections to neighbors. In the figure above, the BGP state is established, which means that the neighbors have established a TCP connection and the two peers have agreed to use BGP to communicate.

For more details about **show ip bgp neighbors** command, please check the Cisco IOS IP Routing Protocols Command Reference on the following link:

http://www.cisco.com/en/US/docs/ios/iproute/command/reference/irp_book.html

BGP Active State Verification

Active: The router has sent an open packet and is waiting for a response.

- The state may cycle between active and idle.
- The neighbor may not know how to get back to this router because of the following reasons:
 - No route to the source IP address of the BGP open packet
 - The neighbor is peering with the wrong address
 - No neighbor statement for this router
 - The AS number is misconfigured

If the router is in the active state, it has found the IP address in the neighbor statement and has created and sent out a BGP open packet. However, the router has not received a response (open confirm packet).

One common problem in this case is that the neighbor may not have a return route to the source IP address. Ensure that the source IP address or network of the packets has been announced to the local routing protocol (IGP).

Another common problem associated with the active state occurs when a BGP router attempts to peer with another BGP router that does not have a neighbor statement peering back to the first router, or when the other router is peering with the wrong IP address on the first router. Check to ensure that the other router has a **neighbor** statement peering to the correct address of the router that is in the active state.

If the state toggles between the idle state and the active state, one of the most common problems is AS number misconfiguration.

Example: BGP Active State Verification

AS number misconfiguration:

- At the router with the wrong remote AS number:
 - %BGP-3-NOTIFICATION: sent to neighbor 172.31.1.3 2/2
(peer in wrong AS) 2 bytes FDFC
 - FFFF FFFF FFFF FFFF FFFF FFFF FFFF FFFF 002D 0104 FDFC
00B4 AC1F 0203 1002 0601 0400 0100 0102 0280 0002 0202 00
- At the remote router:
 - %BGP-3-NOTIFICATION: received from neighbor 172.31.1.1 2/2
(peer in wrong AS) 2 bytes FDFC

© 2009 Cisco Systems, Inc. All rights reserved.

20110110-1520

Example: BGP Active State Verification

If the AS number has been misconfigured, the following console message will be seen at the router with the wrong remote AS number configured in the neighbor statement:

```
%BGP-3-NOTIFICATION: sent to neighbor 172.31.1.3 2/2 (peer in  
wrong AS) 2 bytes FDE6  
FFFF FFFF FFFF FFFF FFFF FFFF FFFF FFFF 002D 0104 FDE6 00B4  
AC1F 0203 1002 0601 0400 0100 0102 0280 0002 0202 00
```

At the remote router, the following message will be displayed:

```
%BGP-3-NOTIFICATION: received from neighbor 172.31.1.1 2/2  
(peer in wrong AS) 2 bytes FDE6
```

Example: BGP Peering

```
R2#show ip bgp summary
BGP router identifier 10.1.1.1, local AS number 65010
BGP table version is 124, main routing table version 124
9 network entries using 1053 bytes of memory
14 path entries using 1144 bytes of memory
12/5 BGP path/bestpath attribute entries using 1488 bytes of memory
6 BGP AS-PATH entries using 144 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 3829 total bytes of memory
BGP activity 58/49 prefixes, 72/50 paths, scan interval 60 secs

Neighbor      V    AS MsgRcvd MsgSent  TblVer  InQ OutQ Up/Down  State/PfxRcd
10.1.0.2      4 65010    11     11    124    0   0 00:02:28      8
172.31.1.3   4 65020     0     0     0     0   0 never Active
172.31.11.4  4 65030    11    10    124    0   0 00:01:11      6

R2#
*Feb 26 14:49:28.211: %BGP-3-NOTIFICATION: received from neighbor 172.31.1.3 2/2
(peer in wrong AS) 2 bytes FDFC
```

Example: BGP Peering

The **show ip bgp summary** command is one way to verify the neighbor relationship. The figure above presents the output from this command. Some of the details of this command output are as follows:

- **BGP router ID:** The IP address that all other BGP speakers recognize as representing this router.
- **BGP table version:** The version increases in increments when the BGP table changes.
- **Main routing table version:** The last version of the BGP database that was injected into the main routing table.
- **Neighbor:** The IP address that is used in the neighbor statement with which this router has a relationship.
- **Version (V):** The version of BGP that this router is running with the listed neighbor.
- **AS:** The AS number of the listed neighbor.
- **Messages received (MsgRcvd):** The number of BGP messages that have been received from this neighbor.
- **Messages sent (MsgSent):** The number of BGP messages sent to this neighbor.
- **Table version (TblVer):** The BGP table version.
- **In queue (InQ):** The number of messages waiting to be processed from this neighbor.
- **Out queue (OutQ):** The number of messages queued and waiting to be sent to this neighbor. TCP flow control prevents this router from overwhelming a neighbor with a large update.

- **Up/Down:** The length of time that this neighbor has been in the current BGP state (established, active, or idle).
- **State (established, active, idle, open sent, open confirm, or idle [admin]):** The BGP state. A neighbor can be set to administratively shut down (admin state) by using the **neighbor shutdown** router configuration command.
- **Prefix received (PfxRcd):** When the session is in the established state, this value represents the number of BGP network entries received from the listed neighbor.

For more details about **show ip bgp summary** command, please check the Cisco IOS IP Routing Protocols Command Reference on the following link:

http://www.cisco.com/en/US/docs/ios/iproute/command/reference/irp_book.htm

BGP Authentication

This topic describes the configuration of Message Digest 5 (MD5) authentication on the BGP TCP connection between two routers.

BGP Neighbor Authentication

```
R2 (config-router) #
neighbor 10.1.1.2 password MYpasswordSTRING
```

- BGP authentication uses MD5
- Configure a key—password; router generates a message digest (is sent), or hash, of the key (is not sent) and the message
- Router generates and checks the MD5 digest of every segment sent on the TCP connection
- Router authenticates the source of each routing update packet that it receives

© 2009 Cisco Systems, Inc. All rights reserved. C91157-10-000

BGP neighbor authentication can be configured on a router so that the router authenticates the source of each routing update packet that it receives. This authentication is accomplished by the exchange of an authenticating key (sometimes referred to as a password) that is known to both the sending and the receiving router.

BGP supports MD5 neighbor authentication. MD5 sends a message digest (also called a “hash”) that is created using the key and a message. The message digest is then sent instead of the key. The key itself is not sent to prevent it from being read by an eavesdropper on the line while it is being transmitted.

To enable MD5 authentication on a TCP connection between two BGP peers, use the **neighbor {ip-address | peer-group-name} password string** router configuration command.

MD5 authentication can be configured between two BGP peers, which means that each segment sent on the TCP connection between the peers is verified. MD5 authentication must be configured with the same password on both BGP peers; otherwise, the connection between the peers will not be made. Configuring MD5 authentication causes Cisco IOS Software to generate and check the MD5 digest of every segment sent on the TCP connection.

Note If the authentication string is configured incorrectly, the BGP peering session will not be established. It is recommended that the authentication string be entered carefully and verify that the peering session is established after authentication is configured.

If a BGP peer group is specified by using the *peer-group-name* argument, all the members of the peer group will inherit the characteristic configured with this command.

If a router has a password configured for a neighbor, but the neighbor router does not, a message such as the following will appear on the console when the routers attempt to send BGP messages between themselves:

```
%TCP-6-BADAUTH: No MD5 digest from 10.1.0.2(179) to
10.1.0.1(20236)
```

Similarly, if the two routers have different passwords configured, a message such as the following will appear on the screen:


```
%TCP-6-BADAUTH: Invalid MD5 digest from 10.1.0.1(12293) to
10.1.0.2(179)
```

If the password or key used for MD5 authentication between two BGP peers is configured or changed, the local router will not tear down the existing session after the password is configured. The local router will attempt to maintain the peering session using the new password until the BGP hold-down timer expires. The default time period is 180 seconds. If the password is not entered or changed on the remote router before the hold-down timer expires, the session times out.

Note Configuring a new timer value for the hold-down timer will only take effect after the session has been reset. It is not possible to change the configuration of the hold-down timer to avoid resetting the BGP session.

For more details about **neighbor {ip-address | peer-group-name} password string** command, please check the Cisco IOS IP Routing Protocols Command Reference on the following link: http://www.cisco.com/en/US/docs/ios/iproute/command/reference/irp_book.html

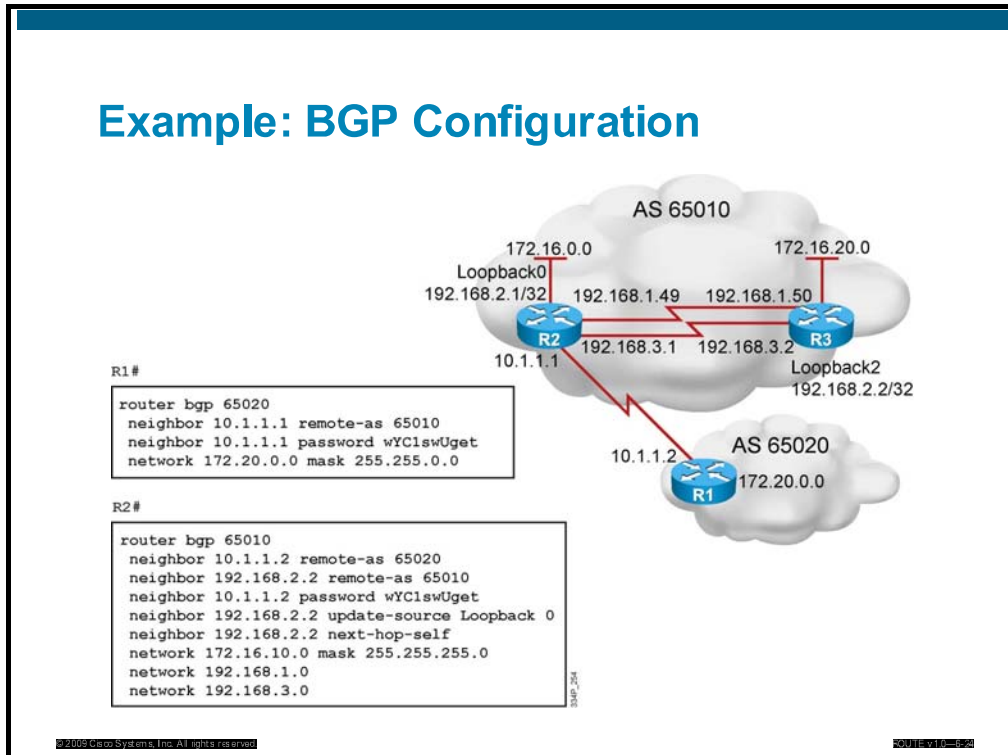
Example: BGP Neighbor Authentication



The example in the figure above configures MD5 authentication for the BGP peering session between routers R1 and R2. The same password must be configured on the remote peer before the hold-down timer expires.

Example of Activating Basic BGP

This topic describes the configuration of BGP operations in a single-homed environment.



The figure above shows another BGP example, which will take into account several BGP features and conclude the topic of configuring basic BGP configuration steps. Router R1 has one EBGP neighbor, router R2 has one EBGP and one IBGP neighbor, and router R3 has one IBGP neighbor. Showing the configuration of routers R1 and R2 will explain the basic configuration needed to create an EBGP and IBGP neighbor relationship.

The configuration for router R1 is as follows:

The first command under the **router bgp 65020** command establishes that router R1 has the following BGP neighbor:

- Router R2 in AS 65010

From the perspective of router R1, router R2 is an EBGP neighbor. The neighbor statement on router R1 for router R2 is pointing at the directly connected IP address to reach the EBGP neighbor, router R2. Authentication is enabled toward router R2 to authenticate the source of routing updates for the EBGP neighbor using password **wYClswUget**. The **network** command notifies BGP about which network to advertise. It contains a subnet of a class B address using the mask option.

The configuration for router R2 is as follows:

The first two commands under the **router bgp 65010** command establish that router R2 has the following two BGP neighbors:

- Router R1 in AS 65020
- Router R3 in AS 65010

From the perspective of router R2, router R1 is an EBGP neighbor and router R3 is an IBGP neighbor. The **neighbor** statement on router R2 for router R1 is pointing at the directly connected IP address to reach the EBGP neighbor, router R1. However, the **neighbor** statement on router R2 points to the loopback interface of router R3 because router R2 has multiple paths to reach router R3. If router R2 pointed at the 192.168.3.2 IP address of router R3 and that interface went down, router R2 would be unable to reestablish the BGP session until the link came back up. By pointing to the loopback interface of router R3 instead, the link stays established as long as any path to router R3 is available. Router R3 should also point to the loopback address of router R2 in its configuration.

Authentication is enabled toward router R1 to authenticate the source of routing updates for the EBGP neighbor using the correct password.

The **neighbor 192.168.2.2 update-source Loopback 0** command notifies router R2 to always use its loopback 0 address, 192.168.2.1, as the source IP address when sending an update to router R3 (192.168.2.2).

Router R2 also changes the next-hop address for networks that are reachable through it. The default next-hop setting for networks from AS 65020 is IP address 10.1.1.2. With the **next-hop-self** command, router R2 sets the next-hop address to the source IP address of the routing update, which is the router R2 loopback 0 interface, as set by the **update-source** command.

Three **network** commands notify BGP about which networks to advertise. The first is a subnet of a class B address using the mask option. The next two **network** statements are used for the two class C networks that connect router R2 and router R3. The default mask is 255.255.255.0, so it does not need to be included in the command.

BGP Verification

This topic describes the verification of BGP configuration.

Example: show ip bgp Command

```
R2#show ip bgp
BGP table version is 14, local router ID is 172.31.11.1
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal, r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network        Next Hop           Metric LocPrf Weight Path
*> 10.1.0.0/24    0.0.0.0             0      32768 i
* i              10.1.0.2            0      100    0 i
*> 10.1.1.0/24    0.0.0.0             0      32768 i
*>i10.1.2.0/24    10.1.0.2            0      100    0 i
*> 10.97.97.0/24  172.31.1.3          0      65020 65010 i
*                172.31.11.4         0      65030 65010 i
* i              172.31.11.4         0      100    0 65030 65010 i
*> 10.254.0.0/24  172.31.1.3          0      65020 i
*                172.31.11.4         0      65030 65020 i
* i              172.31.1.3          0      100    0 65020 i
r> 172.31.1.0/24  172.31.1.3          0      65020 i
r                172.31.11.4         0      65030 65020 i
r i              172.31.1.3          0      100    0 65020 i
*> 172.31.2.0/24  172.31.1.3          0      65020 i
<output omitted>
```

Displays networks from lowest to highest

Use the **show ip bgp** command to display the BGP topology database (BGP table).

The figure above is a partial sample output of the **show ip bgp** command. The status codes are shown at the beginning of each line of output, and the origin codes are shown at the end of each line. In this output, there is an asterisk (*) in most of the entries in the first column. This means that the next-hop address (in the fifth column) is valid. The next-hop address is not always the router that is directly connected to this router. Other options for the first column are as follows:

- An “s,” for suppressed, indicates that the specified routes are suppressed (usually because routes have been summarized and only the summary route is being sent).
- A “d,” for dampening, indicates that the route is being dampened (penalized) for going up and down too often. Although the route might be up right now, it is not advertised until the penalty has expired.
- An “h,” for history, indicates that the route is unavailable and is probably down; historic information about the route exists, but a best route does not exist.
- An “r,” for Routing Information Base (RIB) failure, indicates that the route was not installed in the RIB. The reason that the route is not installed can be displayed using the **show ip bgp rib-failure** command, as shown in the next figure.
- An “S,” for stale, indicates that the route is stale (this symbol is used in the nonstop forwarding-aware router).

The second column shows “>” when BGP has selected the path as the best path to a network.

The third column is either blank or shows “i.” If it is blank, BGP learned that route from an external peer. An “i” indicates that an IBGP neighbor advertised this path to the router.

The fourth column lists the networks that the router learned.

The Next Hop column lists all the next-hop addresses for each route. This column may contain the entry 0.0.0.0, which signifies that this router is the originator of the route.

The three columns to the left of the Path column list three BGP path attributes that are associated with the path: metric (multi-exit discriminator [MED]), local preference, and weight.

The column with the Path header may contain a sequence of autonomous systems in the path. From left to right, the first AS listed is the adjacent AS, from which this network was learned. The last number (the rightmost AS number) is the originating AS of this network. The AS numbers between these two numbers represent the exact path that a packet takes back to the originating AS. If the path column is blank, the route is from the current AS.

The last column signifies how this route was entered into BGP on the original router. If the last column has an “i” in it, the originating router probably used a network statement to introduce this network into BGP.

If the character is an “e,” the originating router learned this network from exterior gateway protocol (EGP), which is the historical predecessor to BGP. A question mark (?) signifies that BGP cannot absolutely verify the availability of this network because it is redistributed from an IGP into BGP.

Example: show ip bgp rib-failure Command

- Displays networks that are not installed in the RIB and the reason that they were not installed.

```
R2# show ip bgp rib-failure
Network      Next Hop      RIB-failure  RIB-NH Matches
172.31.1.0/24 172.31.1.3    Higher admin distance  n/a
172.31.11.0/24 172.31.11.4  Higher admin distance  n/a
```

Use the **show ip bgp rib-failure** command to display BGP routes that were not installed in the Routing Information Base (RIB) table and the reason that the routes were not installed.

The example in the figure above shows that the displayed routes were not installed because a route or routes with a better administrative distance already exist in the RIB.

For more details about **show ip bgp rib-failure** command, please check the Cisco IOS IP Routing Protocols Command Reference on the following link:

http://www.cisco.com/en/US/docs/ios/iproute/command/reference/rp_book.html

Clearing the BGP Session

- When policies change, the change takes effect immediately.
- The next time that a prefix or path is advertised or received, the new policy is used. This can take a long time for all networks.
- You must trigger an update for immediate action.
- Ways to trigger an update:
 - Hard reset
 - Soft reset
 - Route refresh

© 2009 Cisco Systems, Inc. All rights reserved.

6-80 (v1.0) 6-80

BGP can potentially handle huge volumes of routing information. When a policy configuration change occurs, the router cannot go through the huge table of BGP information and recalculate which entry is no longer valid in the local table; nor can the router determine which route or routes, already advertised, should be withdrawn from a neighbor.

There is an obvious risk that the first configuration change will be immediately followed by a second, which would cause the whole process to start all over again. To avoid such a problem, Cisco IOS Software applies changes only to those updates that are received or transmitted after the BGP policy configuration change has been performed. The new policy, enforced by the new filters, is applied only on routes that are received or sent after the change.

A network administrator who would like the policy change to be applied on all routes must trigger an update to force the router to let all routes pass through the new filter. If the filter is applied on outgoing information, the router has to resend the BGP table through the new filter. If the filter is applied on incoming information, the router needs its neighbor to resend its BGP table so that it passes through the new filters.

There are three ways to trigger an update: with a hard reset, soft reset, or route refresh.

Hard Reset of BGP Sessions

A BGP session makes the transition from established to idle; everything must be relearned.

R2#

```
clear ip bgp *
```

- Resets all BGP connections with this router.
- The entire BGP forwarding table is discarded.

R2#

```
clear ip bgp 10.1.1.2
```

- Resets only a single neighbor.
- Less severe than a **clear ip bgp *** command.

Resetting a session is a method of informing the neighbor or neighbors of a policy change. If BGP sessions are reset, all information received on those sessions is invalidated and removed from the BGP table. Also, the remote neighbor will detect a BGP session down state and will invalidate the routes that were received. After 30 to 60 seconds, the BGP sessions are reestablished automatically and the BGP table is exchanged again, but through the new filters. However, resetting the BGP session disrupts packet forwarding.

Both commands that are shown in the figure above cause a hard reset of the BGP neighbors that are involved. A hard reset means that the router issuing either of these commands will close the appropriate TCP connections, reestablish those TCP sessions as appropriate, and resend all information to each of the neighbors affected by the particular command that is used.

The **clear ip bgp *** command causes the BGP forwarding table on the router that issued this command to be completely deleted, and all networks must be relearned from every neighbor. If a router has multiple neighbors, this action is a very dramatic event. This command forces all neighbors to resend their entire tables simultaneously.

For example, consider a situation in which router R1 has eight neighbors and each neighbor has a full Internet table of about 32 MB in size. If router R1 issues the **clear ip bgp *** command, all eight routers resend their 32-MB tables at the same time. To hold all these updates, router R1 would need 256 MB of RAM. Router R1 would also need to be able to process all of this information. Processing 256 MB of updates would take a considerable number of CPU cycles for router R1, further delaying the routing of user data.

If the second command, **clear ip bgp 10.1.1.2**, is used instead, one neighbor is reset at a time. The impact is less severe on the router that is issuing this command; however, it takes longer to change policy for all of the neighbors, because each must be done individually rather than all at once as with the **clear ip bgp *** command. The **clear ip bgp** command still performs a hard reset and must reestablish the TCP session with the specified address, but this command affects only a single neighbor at a time.

For more details about **clear ip bgp** command, please check the Cisco IOS IP Routing Protocols Command Reference on the following link:

http://www.cisco.com/en/US/docs/ios/iproute/command/reference/irp_book.html

Soft Reset Outbound

R2#

```
clear ip bgp 10.1.1.2 soft out
```

- Routes learned from this neighbor are not lost.
- This router resends all BGP information to the neighbor without resetting the connection.
- This option is highly recommended when you are changing the outbound policy.
- The soft out option does not help if you are changing an inbound policy.

The **soft out** option of the **clear ip bgp** command causes BGP to do a soft reset for outbound updates. The router issuing the **clear ip bgp 10.1.1.2 soft out** command does not reset the BGP session; instead, the router creates a new update and sends the whole table to the specified neighbor.

This update includes withdrawal commands for the networks that the other neighbor will not see anymore based on the new outbound policy.

Note The **soft** keyword of this command is optional; **clear ip bgp out** does a soft reset for outbound updates.

Inbound Soft Reset

R2 (config-router) #

```
neighbor 10.1.1.2 soft-reconfiguration inbound
```

- This router stores all updates from this neighbor in case the inbound policy is changed.
- The command is memory intensive.

R2#

```
clear ip bgp 10.1.1.2 soft in
```

- Uses the stored information to generate new inbound updates.

© 2009 Cisco Systems, Inc. All rights reserved.

6-84

There are two ways to perform an inbound soft reconfiguration: using stored routing update information, as shown in the figure above, and dynamically, as shown in the next figure.

Enter the **neighbor** command that is shown in the figure above to inform BGP to save all updates that were learned from the neighbor specified. The BGP router retains an unfiltered table of what that neighbor has sent.

When the inbound policy is changed, use the **clear ip bgp** command shown in the figure above. The stored unfiltered table is used to generate new inbound updates; the new results are placed in the BGP forwarding database. Thus, if changes are made, the other side does not have to be forced to resend everything.

Route Refresh: Dynamic Inbound Soft Reset

R2#

```
clear ip bgp {*}|10.1.1.2} [soft in | in]
```

- Routes advertised to this neighbor are not withdrawn
- Does not store update information locally
- The connection remains established
- Introduced in Cisco IOS Software Release 12.0(2)S and 12.0(6)T

Cisco IOS Software Releases 12.0(2)S and 12.0(6)T introduced a BGP soft reset enhancement feature, also known as “route refresh,” which provides automatic support for dynamic soft reset of inbound BGP routing table updates, and is not dependent upon stored routing table update information. The **clear ip bgp soft in** command implements this feature. This method requires no preconfiguration and requires significantly less memory than the previous soft method for inbound routing table updates.

The **soft in** option generates new inbound updates without resetting the BGP session, but it can be memory intensive. BGP does not allow a router to force another BGP speaker to resend its entire table. If the inbound BGP policy is changed and a hard reset is not to be completed, configure the router to perform a soft reconfiguration.

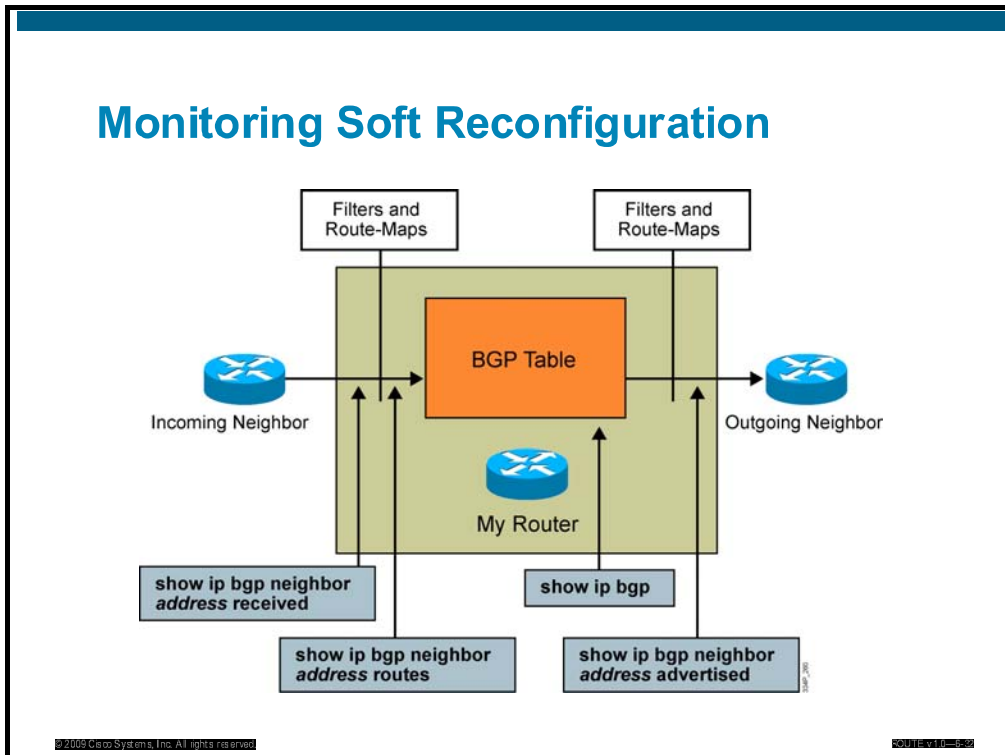
Note To determine whether a BGP router supports this route refresh capability, use the **show ip bgp neighbors** command. The following message is displayed in the output when the router supports the route refresh capability:

```
Received route refresh capability from peer
```

If all BGP routers support the route refresh capability, use the **clear ip bgp {*} | address | peer-group-name} in** command. It is not necessary to use the **soft** keyword because soft reset is automatically assumed when the route refresh capability is supported.

Note The **clear ip bgp soft** command performs a soft reconfiguration of both inbound and outbound updates.

Monitoring Soft Reconfiguration



When a BGP session is reset and soft reconfiguration is used, several commands exist to monitor the BGP routes that are received, sent, or filtered.

The following commands can be used:

- **show ip bgp neighbor address received**
- **show ip bgp neighbor address routes**
- **show ip bgp**
- **show ip bgp neighbor address advertised**

debug ip bgp updates Command

```
R1#debug ip bgp updates
Mobile router debugging is on for address family: IPv4 Unicast
R1#clear ip bgp 10.1.0.2
<output omitted>
*Feb 24 11:06:41.309: %BGP-5-ADJCHANGE: neighbor 10.1.0.2 Up
*Feb 24 11:06:41.309: BGP(0): 10.1.0.2 send UPDATE (format)
10.1.1.0/24, next 10.1.0.1, metric 0, path Local
*Feb 24 11:06:41.309: BGP(0): 10.1.0.2 send UPDATE (prepend, chgflags:
0x0) 10.1.0.0/24, next 10.1.0.1, metric 0, path Local
*Feb 24 11:06:41.309: BGP(0): 10.1.0.2 NEXT_HOP part 1 net
10.97.97.0/24, next 172.31.11.4
*Feb 24 11:06:41.309: BGP(0): 10.1.0.2 send UPDATE (format)
10.97.97.0/24, next 172.31.11.4, metric 0, path 65030 65010
*Feb 24 11:06:41.309: BGP(0): 10.1.0.2 NEXT_HOP part 1 net
172.31.22.0/24, next 172.31.11.4
*Feb 24 11:06:41.309: BGP(0): 10.1.0.2 send UPDATE (format)
172.31.22.0/24, next 172.31.11.4, metric 0, path 65030
<output omitted>
*Feb 24 11:06:41.349: BGP(0): 10.1.0.2 rcvd UPDATE w/ attr: nexthop
10.1.0.2, origin i, localpref 100, metric 0
*Feb 24 11:06:41.349: BGP(0): 10.1.0.2 rcvd 10.1.2.0/24
*Feb 24 11:06:41.349: BGP(0): 10.1.0.2 rcvd 10.1.0.0/24
```

The figure above shows a partial output from the **debug ip bgp updates** command on router R1 after the **clear ip bgp** command is issued to clear the BGP sessions with its IBGP neighbor 10.1.0.2.

After the neighbor adjacency is reestablished, router R1 creates and sends updates to 10.1.0.2. The first update highlighted in the figure, “10.1.1.0/24, next 10.1.0.1,” is an update about network 10.1.1.0/24, with a next hop of 10.1.0.1, which is the address of router R1.

The second update highlighted in the figure, “10.97.97.0/24, next 172.31.11.4,” is an update about network 10.97.97.0/24, with a next hop of 172.31.11.4, which is the address of one of the EBGP neighbors of router R1. The EBGP next-hop address is being carried into IBGP.

Router R1 later receives updates from 10.1.0.2. The update highlighted in the figure contains a path to two networks, 10.1.2.0/24 and 10.1.0.0/24. The attributes shown in this update are described in the next lesson.

Note Debugging uses up router resources and should be turned on only when necessary.

For more details about **debug ip bgp updates** command, please check the Cisco IOS Debug Command Reference on the following link:

http://www.cisco.com/en/US/docs/ios/debug/command/reference/db_book.html

Summary

This topic summarizes the key points that were discussed in this lesson.

Summary

- For a BGP configuration, the following must be defined: BGP requirements, BGP parameters, and connectivity.
- BGP is configured with the following basic BGP commands:
router bgp *autonomous-system*, neighbor *ip-address remote-as autonomous-system*, network *network-number [mask network-mask]*
- The **neighbor shutdown** command administratively shuts down a BGP neighbor.
- When creating a BGP packet, the neighbor statement defines the destination IP address and the outbound interface defines the source IP address.

Summary (cont.)

- When establishing a BGP session, the BGP goes through the following states: idle, connect, open sent, open confirm, and established.
- You can configure MD5 authentication between two BGP peers, which means that each segment sent on the TCP connection between the peers is verified.
- One EBGP neighbor exists in a single-homed environment.
- The **show** and **debug** commands are used to troubleshoot the BGP session.

Lesson 5

Lab 6-1 Debrief

Overview

In Lab 6-1, students configure Border Gateway Protocol (BGP) operations. Given a routed network, students simulate multihomed Internet service provider (ISP) connections and configure External BGP (EBGP) on the enterprise router to connect to the two ISPs.

After completing the lab, the instructor will lead the discussion about lab topology, tasks, verification, and checkpoints, as well as sample solutions and different alternatives. Students will present their implementation plan and solution.

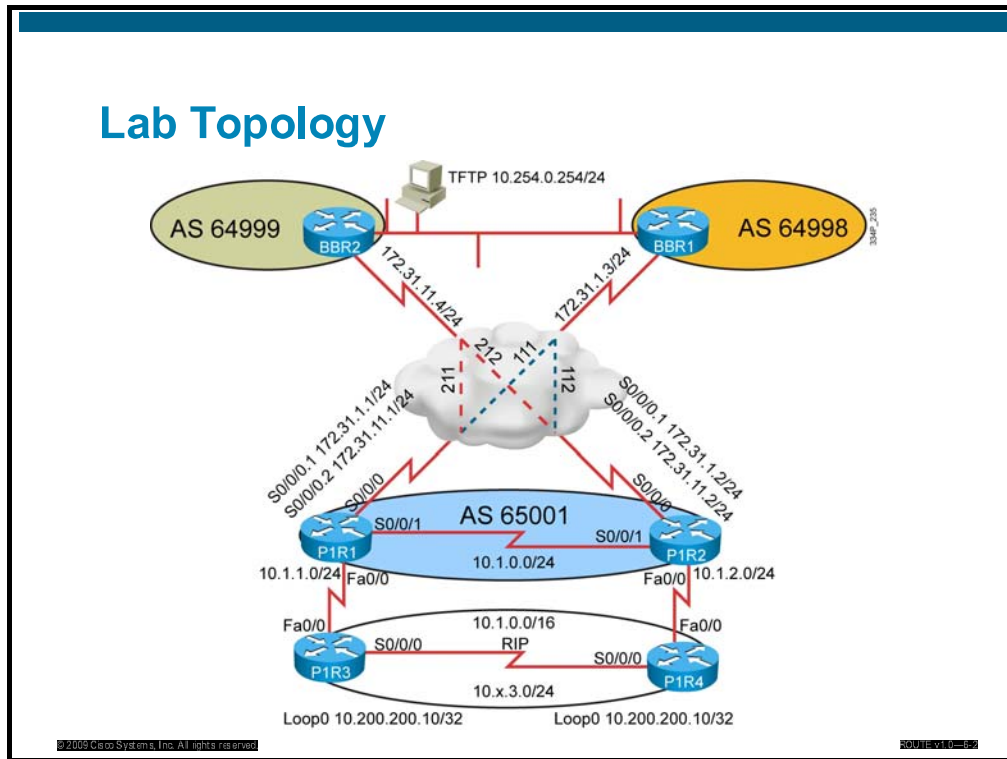
Objectives

Upon completing this lesson, you will be able to explain the lab topology, configure BGP operations, create checkpoints for configuration and verification, and find alternative solutions. This ability includes being able to meet these objectives:

- Compare your solution, findings, and action log against a set of checkpoints provided by the instructor and identify common and alternative solutions.
- Consolidate the lessons learned during the review discussions into a set of best practice methods and commands to aid you in future troubleshooting procedures.

Lab Overview and Verification

This topic describes the lab topology and key checkpoints used to create a solution and begin verification.



The figure above presents a logical lab topology used for configuration of a BGP operations lab. The topology uses four pod routers, which are members of an Internal BGP (IBGP) relationship, and two backbone routers, which are members of an EBGP relationship. The lab is preconfigured with addressing, including loopbacks and an IGP routing protocol.

Based on the topology, students can create a basic BGP operations configuration, including entering BGP processes, creating neighbors, advertising networks, as well as utilizing additional BGP options including synchronization, peer groups, and authentication.

Lab Review: What did you accomplish?

Task 1 – Cleaning Up

- What steps did you take to remove unnecessary configuration?
- How did you create the initial configuration?

Task 2 – Configuring BGP

- What is needed to configure basic BGP?
- What steps did you take to configure neighbors, advertise networks, and verify the configuration?

Task 3 – Configuring BGP authentication

- Which authentication type did you select?
- How can you verify MD5 authentication?

The lab consists of three tasks. Students need an implementation plan to fulfill the lab requirements.

In the first task, students need to perform lab cleanup, in which the previous configuration is erased and the lab is prepared with an initial configuration supporting Lab 6-1, “Configure BGP Operations.”

In the second task, students configure basic BGP steps from entering the BGP process to neighbor definition and network advertisements.

In the third task, students configure additional BGP configuration options including Message Digest 5 (MD5) authentication, and learn additional commands used for better reachability of next hop between IBGP and EBGP neighbors.

Verification

- Is your solution working?
- What method did you follow to verify the status of all BGP connections?
- Which commands did you use to verify the following information?
 - Display information about BGP and TCP connections to neighbors
 - Display the parameters and current state of the active routing protocol process
 - Display entries in the BGP routing table

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE-01-00-00

After each configuration task, verification takes place where students check their configurations. Several **show** commands are used during the verification process showing many parameters of the BGP process and neighbor relationship.

Verification includes use of the following commands:

- **show ip bgp summary**: Displays the status of all BGP connections.
- **show ip bgp neighbors**: Displays information about BGP and TCP connections to neighbors.
- **show ip protocols**: Displays the parameters and current state of the active routing protocol process.
- **show ip bgp**: Displays entries in the BGP routing table.

Check Points

- BGP process and AS number
- Neighbor relationship status
- Advertise networks
- BGP routing information base
- BGP routes
- BGP protocol
- Default route
- Authentication

During the configuration and verification phase, a network operator can deal with several checkpoints. After completing all configuration tasks, BGP configuration can be successfully completed or additional verification and troubleshooting is needed.

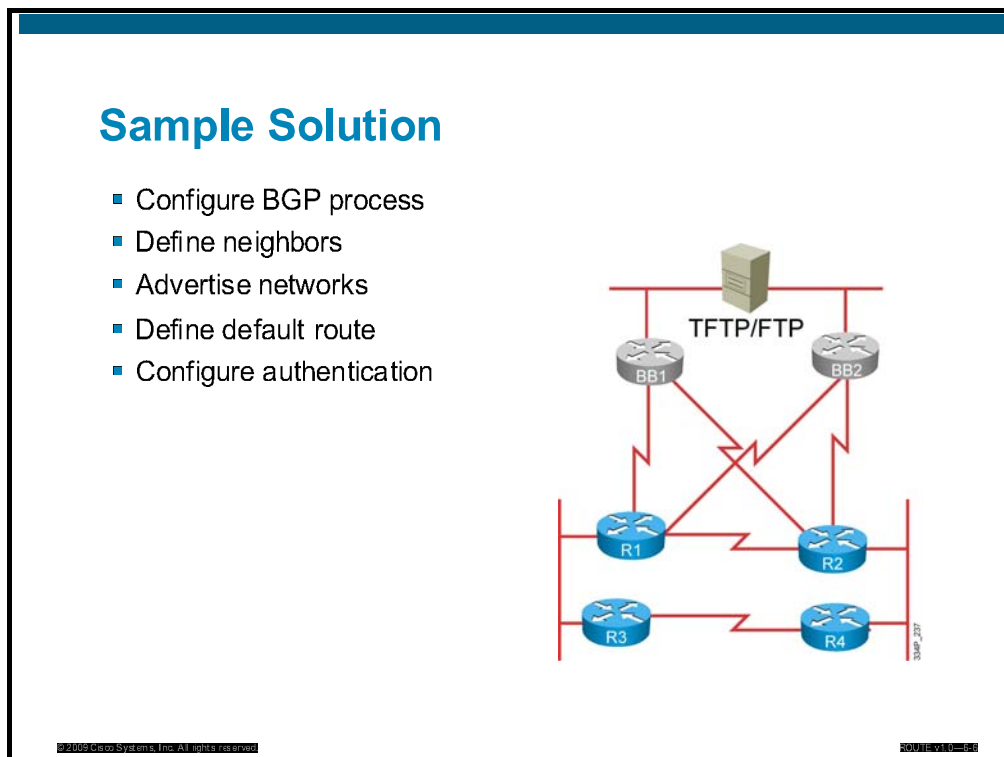
Optionally, a network operator can check the BGP configuration in different stages of the implementation using the checkpoints at each stage.

With different checkpoints, a network operator can check for proper configuration with the following steps:

- Check BGP process and autonomous system (AS) number
- Check neighbor relationship status
- Check which networks are advertised
- Check BGP routing information base
- Check BGP routes
- Check BGP protocol
- Check for default route
- Check authentication

Sample Solutions and Alternatives

This topic describes a sample solution and other alternatives.



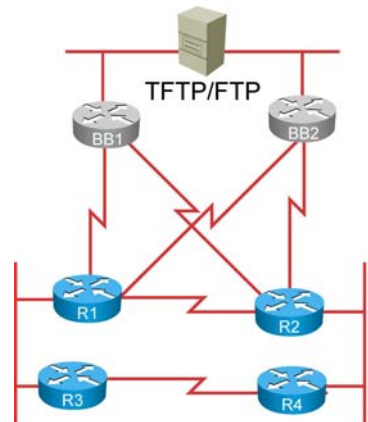
A sample solution includes configuration steps for each of the tasks. Different solutions are possible and the figure above points out the important tasks needed for a successful configuration.

To be able to successfully complete the lab configuration, use the following guidelines:

- Enter the BGP process using the correct AS number.
- Define neighbors using a directly connected interface (EBGP) or loopback interface (IBGP).
- Advertise the local networks into BGP using a network statement with the **mask** keyword.
- Define the default route inside IGP that is used inside the AS.
- Configure authentication on both neighbors.

Alternative Solutions

- BGP process
 - Different AS numbers and connectivity options
- EBGP/IBGP neighbors
 - **update-source** command
- Advertise networks
 - **next-hop-self**, synchronization



Same or similar results can be achieved by using different configuration steps.

When entering a BGP process, different AS numbers can be used between the neighbors. An EBGP or IBGP neighbor relationship is defined based on which AS number is used among the neighboring routers.

EBGP and IBGP relationships have different connectivity rules. An EBGP neighbor relationship is established between the directly connected neighbors and an IBGP neighbor relationship does not have this requirement. Loopbacks are used for IBGP neighbor relationships.

Networks are advertised with the next-hop attribute, which is processed in a different way based on an EBGP or IBGP neighbor relationship. The next-hop attribute can be changed, as well as the network advertisement between the EBGP and IBGP neighbors.

Q and A

- 1 Why is the AS number important?
- 2 Why is the neighbor IP address important?
- 3 Why is authentication necessary?

The AS number is important when creating neighbors because AS numbers define the EBGP or IBGP relationship.

An IP address used during neighbor configuration is important because EBGP neighbors require directly connected neighbors rather than IBGP neighbors where there is no need to be directly connected.

The synchronization rule says that a router will not advertise routes in BGP until it learns the routes in an IGP. It is safe to have the synchronization rule off only if all routers in the transit path in the AS are running full-mesh IBGP.

Authentication in BGP is used to verify that only trusted neighbors establish the neighbor relationship. If some other routers configured for BGP are added to the system, they cannot attract traffic and establish the neighbor relationship.

Summary

This topic summarizes the key points that were discussed in this lesson.

Summary

- Create a good implementation plan and define the BGP requirements before configuring BGP operations.
- Several solutions exist and alternative solutions give similar or very different results.

Using the BGP Attributes and Path Selection Process

Overview

Border Gateway Protocol (BGP) is used to perform policy-based routing. In order to manipulate the best paths chosen by BGP, a network administrator must understand the different attributes that BGP uses and how BGP selects the best path based on these attributes.

This lesson explains the BGP path selection process and the BGP attributes and their characteristics, and discusses configuration steps using different attributes for selecting a BGP path. Route filtering using route maps and prefix lists is described as well.

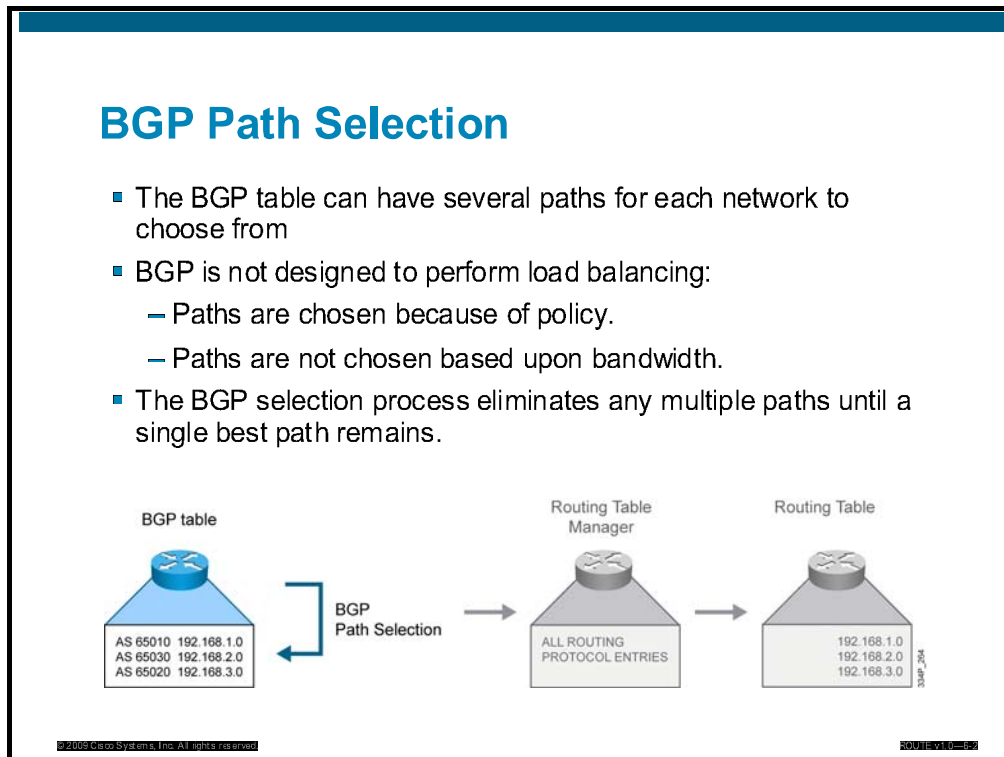
Objectives

Upon completing this lesson, you will be able to configure and verify BGP operations in a multihomed environment using the BGP attributes and route maps to control all BGP routes to and from the router. This ability includes being able to meet these objectives:

- Understand the BGP path selection process.
- Identify the characteristics of BGP attributes for path selection and path manipulation.
- Configure filtering of BGP routing updates.

BGP Path Selection

This topic describes the criteria for selecting a BGP path.



Routers often have several neighbors and receive routing updates from each neighbor. All routing updates enter the BGP forwarding table, and as a result, multiple paths may exist to reach a given network. The whole BGP forwarding table can be displayed using the **show ip bgp** command.


Next, paths for the network are evaluated to determine which path is best. Paths that are not the best are eliminated from the selection criteria but kept in the BGP forwarding table in case the best path becomes inaccessible. If one of the best paths is not accessible, then a new best path must be selected.

BGP is not designed to perform load balancing; paths are chosen based on policy, not based on bandwidth. The BGP selection process eliminates any multiple paths until a single best path remains.

Routing Table Manager

The best path is submitted to the routing table manager process.

- The best path is evaluated against the routes of other routing protocols for reaching that network.
- The route with the lowest administrative distance from the source will be installed in the routing table.



The best path is submitted to the routing table manager process and is evaluated against any other routing protocols that can also reach that network. The router usually runs BGP and one of the interior gateway protocols (IGP) (such as Routing Information Protocol [RIP], Enhanced Interior Gateway Routing Protocol [EIGRP], Open Shortest Path First [OSPF], and so on), which are sending candidates for the routing table to the routing table manager. The route from the source with the lowest administrative distance is installed in the routing table. The whole routing table can be displayed using the **show ip route** command.

Route Selection Decision Process

Consider only (synchronized) routes with no AS loops and a valid next hop. The next steps in the evaluation process are:

1	Prefer highest weight (local to router).
2	Prefer highest local preference (global within AS).
3	Prefer route originated by the local router (next hop = 0.0.0.0).
4	Prefer shortest AS path.
5	Prefer lowest origin code (IGP < EGP < incomplete).
6	Prefer lowest MED (exchanged between autonomous systems).
7	Prefer the EBGP path over the IBGP path.
8	Prefer the path through the closest IGP neighbor.
9	Prefer the oldest route for EBGP paths.
10	Prefer the path with the lowest neighbor BGP router ID.
11	Prefer the path with the lowest neighbor IP address.

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE-100-000

After BGP receives updates about different destinations from different autonomous systems, it chooses the single best path to reach a specific destination.

The decision process is based on the BGP attributes. When faced with multiple routes to the same destination, BGP chooses the best route for routing traffic toward the destination. BGP considers only synchronized routes with no autonomous system loops and a valid next hop. The following process summarizes how BGP chooses the best route on a Cisco router:

1. Prefer the route with the highest weight. (Recall that the weight is proprietary to Cisco and is local to the router only.)
2. If multiple routes have the same weight, prefer the route with the highest local preference. (Recall that the local preference is used within an autonomous system.)
3. If multiple routes have the same local preference, prefer the route that the local router originated. A locally originated route has a next hop of 0.0.0.0 in the BGP table.
4. If none of the routes were locally originated, prefer the route with the shortest autonomous system path.
5. If the autonomous system path length is the same, prefer the lowest origin code (IGP < EGP < incomplete).

6. If all origin codes are the same, prefer the path with the lowest multiexit discriminator (MED). (Recall that the MED is exchanged between autonomous systems.)

The MED comparison is made only if the neighboring autonomous system is the same for all routes considered, unless the **bgp always-compare-med** command is enabled.

Note The most recent Internet Engineering Task Force (IETF) decision regarding BGP MED assigns a value of infinity to the missing MED, making the route lacking the MED variable the least preferred. The default behavior of BGP routers running Cisco IOS Software is to treat routes without the MED attribute as having a MED of 0, making the route lacking the MED variable the most preferred. To configure the router to conform to the IETF standard, use the **bgp bestpath missing-as-worst** command.

7. If the routes have the same MED, prefer external paths (EBGP) to internal paths (Internal BGP, or IBGP).
8. If synchronization is disabled and only internal paths remain, prefer the path through the closest IGP neighbor. This step means that the router will prefer the shortest internal path within the autonomous system to reach the destination (the shortest path to the BGP next hop).
9. For EBGP paths, select the oldest route to minimize the effect of routes going up and down (flapping).
10. Prefer the route with the lowest neighbor BGP router ID value.
11. If the BGP router IDs are the same, prefer the router with the lowest neighbor IP address.

Only the best path is entered in the routing table and propagated to the BGP neighbors of the router.

Note The route selection process summarized here does not cover all cases but is sufficient for a basic understanding of how BGP selects routes.

For example, suppose there are seven paths to reach network 10.0.0.0. None of the paths have autonomous system loops and all paths have valid next-hop addresses, so all seven paths proceed to Step 1, which examines the weight of the paths.

All seven paths have a weight of 0, so all paths proceed to Step 2, which examines the local preference of the paths. Four of the paths have a local preference of 200, and the other three have local preferences of 100, 100, and 150.

The four with a local preference of 200 will continue the evaluation process in the next step. The other three will still be in the BGP forwarding table, but are currently disqualified as the best path.

BGP will continue the evaluation process until only a single best path remains. The single best path that remains will be submitted to the IP routing table as the best BGP path.

Path Selection with Multihomed Connection

An autonomous system rarely implements BGP with only one EBGP connection. This situation generally means that multiple paths exist for each network in the BGP forwarding database.

If only one path exists and it is loop-free and synchronized with the IGP for IBGP, and the next hop is reachable, the path is submitted to the IP routing table. There is no path selection taking place because there is only one path and manipulating it will derive no benefit.

Only the best path is put in the routing table and propagated to the router's BGP neighbors.

Without route manipulation, the most common reason for path selection is Step 4, prefer the shortest autonomous system path.

Step 1 looks at weight, which by default is set to 0 for routes that were not originated by this router.

Step 2 compares local preference which, by default, is set to 100 for all networks. Both of these steps have an effect only if the network administrator configures the weight or local preference to a nondefault value.

Step 3 looks at networks that are owned by this autonomous system. If one of the routes is injected into the BGP table by the local router, the local router prefers it to any routes received from other BGP routers.

Step 4 selects the path that has the fewest autonomous systems to cross. This is the most common reason a path is selected in BGP. If a network administrator does not like the path with the fewest autonomous systems, the administrator needs to manipulate the weight or local preference to change which outbound path BGP chooses.

Step 5 looks at how a network was introduced into BGP. This introduction is usually either with **network** statements (i for an origin code) or through redistribution (? for an origin code).

Step 6 looks at MED to judge where the neighbor autonomous system wants this autonomous system to send packets for a given network. Cisco sets the MED to 0 by default; therefore, MED does not participate in path selection unless the network administrator of the neighbor autonomous system manipulates the paths using MED.

If multiple paths have the same number of autonomous systems to traverse, the second most common decision point is Step 7, which states that an externally learned path from an EBGP neighbor is preferred over a path learned from an IBGP neighbor. A router in an autonomous system prefers to use the bandwidth of the Internet service provider (ISP) to reach a network rather than using internal bandwidth to reach an IBGP neighbor on the other side of its own autonomous system.

If the autonomous system path is equal and the router in an autonomous system has no EBGP neighbors for that network (only IBGP neighbors), it makes sense to take the quickest path to the nearest exit point. Step 8 looks for the closest IBGP neighbor; the IGP metric determines what "closest" means (for example, RIP uses hop count, and OSPF uses the least cost, based on bandwidth).

If the autonomous system path is equal and the costs via all IBGP neighbors are equal, or if all neighbors for this network are EBGP, the oldest path, Step 9, is the next common reason for selecting one path over another. EBGP neighbors rarely establish sessions at exactly the same time. One session is likely older than another, so the paths through that older neighbor are considered more stable because those paths have been up for a longer period of time.

If all of the above criteria are equal, the next most common decision is to choose the neighbor with the lowest BGP router ID, which is Step 10.


If the BGP router IDs are all the same (for example, if the paths are to the same BGP router), Step 11 states that the route with the lowest neighbor IP address is used.

Characteristics of BGP Attributes for Path Selection and Path Manipulation

BGP attributes inform the BGP routers receiving updates about how to treat the paths to the final network. This topic describes the characteristics of the BGP attributes that affect path selection.

Weight Attribute

- **Weight is an attribute that is proprietary to Cisco.**
- Weight is not sent to any BGP neighbors.
- It is local to the router only.
- Paths with the highest weight value are preferred.



The weight attribute is an attribute that Cisco defines for the path selection process. The weight is configured locally on a router and is not propagated to any other routers. This attribute applies when one router is used with multiple exit points out of an autonomous system, as opposed to the local preference attribute, which is used when two or more routers provide multiple exit points.

The weight can have a value from 0 to 65535. Paths that the router originates have a weight of 32768 by default, and other paths have a weight of 0 by default.

Routes with a higher weight are preferred when multiple routes exist to the same destination.

In the figure above, routers R2 and R4 learn about network 172.20.0.0 from autonomous system (AS) 65020 and propagate the update to router R1. Router R1 has two ways to reach 172.20.0.0, and it must decide which route to take—the path through router R2 or the path through router R4.

In the example, router R1 sets the weight of updates coming from router R2 to 200 and the weight of those coming from router R4 to 150. Because the weight for the route pointing to router R2 is higher than the weight for route pointing to router R4, router R1 uses router R2 as a next hop to reach 172.20.0.0.

Set Weight with Route Map

- **First BGP path selection criteria**
- Prefer the highest weight (local to router)
- BGP weight can be specified per neighbor by a complex criteria with route maps

One way of changing the path selection is to use weights. Weight is an attribute that is locally significant to a router. Weight is a property or parameter. Therefore, it is not seen on any neighboring routers. When designing BGP networks using weights, network administrators should set the weights on every router. If there is more than one path for the same network, a router will choose the one with the highest weight. The default value for weight is 0.

When a route map is applied, it is configured on the router. The route map can be arbitrarily complex and select routes based on various selection criteria, such as a network number or AS path. The selected routes can have some altered attributes. The route map can set the weight values of permitted routes. Selection can be done in several route-map statements, giving the opportunity to assign a certain weight value to some routes and another weight value to others. A route map can also completely filter out routes. AS path filters, prefix lists, or other BGP attributes are used to match routes, where any combination can be used. Routes that are not matched are discarded.

Using Route Maps for Path Selection

R2(config)#

```
route-map MY-Route-Map permit 10
```

- Enter route map configuration mode.

R2(config-route-map)#

```
match local-preference 150  
set weight 200
```

- Match on the BGP attribute.
- Set the new value for the BGP attribute.

R2(config-router)#

```
neighbor 10.0.0.1 route-map MY-Route-Map in|out
```

- Apply the route map to the incoming or outgoing updates.

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE-01-000009

Three steps are needed when using route maps for path selection.


In the first step, a route map must be created. Use the **route-map** command in global configuration mode to create a route map and to enter route map configuration mode.

In the second step a rule for changing the BGP attributes is created. Use the **match** command in route-map configuration mode to define which incoming or outgoing updates will be affected. Use the **set** command in route-map configuration mode to change selected BGP attributes.

The last step is required to apply a configured route map to incoming or outgoing routes. Use the **neighbor route-map** command in router configuration mode to apply the route map to BGP updates.

For more details about the **route-map**, **match**, **set**, and **neighbor route-map** commands, please check the Cisco IOS IP Routing Protocols Command Reference on the following link: http://www.cisco.com/en/US/docs/ios/iproute/command/reference/irp_book.html

Set Weight with Route Map Example




```
<output omitted>
!
router bgp 65040
 neighbor 10.0.0.1 route-map RM-SET-Weight in
!
route-map RM-SET-Weight permit 10
 match as-path 10
 set weight 150
!
route-map RM-SET-Weight permit 99
 set weight 100
!
ip as-path access-list 10 permit _65020$
```

In the example above, the routing policy dictates the selection of AS 65030 as the primary way out of AS 65040 for the traffic destined to any network originated by AS 65020. This limitation is achieved by placing a high weight of 150 on all incoming announcements from AS65030 (that is, from neighbor 10.0.0.1), which carry the information about the network originated in AS 65020.

Local Preference Attribute

- Used to select the outbound EBGP path
- Sent to IBGP neighbors only (and only within the AS)
- Stripped in the outgoing EBGP updates except in the EBGP updates with confederation peers
- The local preference attribute is **well known** and **discretionary**
- Default value = 100
- Paths with the highest local preference value are preferred



Local preference is a well-known discretionary attribute that provides information to routers in the autonomous system about the path that is preferred for exiting the autonomous system. A path with a higher local preference is preferred.

The local preference is an attribute that is configured on a router and exchanged among routers within the same autonomous system only. The default value for local preference on a Cisco router is 100.

To change the default local preference value of 100, use the **bgp default local-preference** command.

In the figure above, AS 65050 receives updates about network 172.16.0.0 from two neighbors. Each of the networks is advertising a different path to the destination. The local preference on router R1 for network 172.16.0.0 is set to 200, and the local preference on router R2 for network 172.16.0.0 is set to 150.

Because the local preference information is exchanged within AS 65050, all routers are aware of the exit point for network 172.16.0.0 out of AS 65050. In the figure, router R1 is configured with higher local preference than router R2 and all the traffic destined for network 172.16.0.0 will be sent to router R1 as an exit point from AS 65050.

For more details about the **route bgp default local-preference** command, please check the Cisco IOS IP Routing Protocols Command Reference on the following link:

http://www.cisco.com/en/US/docs/ios/iproute/command/reference/irp_book.html

Set Local Preference with Route Map

- **Second BGP path selection criteria**
- Prefer highest local preference (global within AS)
- Local preference can be set when
 - processing incoming route updates
 - doing redistribution
 - sending outgoing route updates
- BGP local preference can be specified per neighbor by complex criteria with route maps

Using BGP in autonomous systems with a single neighbor relationship usually does not require any advanced features. In multihomed situations however, it is important to ensure that the customer routers choose the correct link.


Local preference is the second-strongest criterion in the route selection process. If there are two or more paths available for the same network, a router will first compare weight, and if the weights are equal for all paths, the router will then compare the local preference attributes. The path with the highest local preference value will be preferred. The default value for local preference is 100.

Local preference is similar to weight because it is an attribute. It can be set once and can then be viewed on neighboring routers without having to reset it. This attribute has a default value of 100, which the router will apply to locally originated routes and updates that come in from external neighbors. Updates that come from internal neighbors already have the local preference attribute.

The local preference attribute is automatically stripped from outgoing updates to External Border Gateway Protocol (EBGP) sessions. This practice means that this attribute can be used only within a single AS to influence the route selection process.

BGP local preference can be specified per neighbor by complex criteria using route maps. AS path filters, prefix lists, and other BGP attributes are used to match routes, where any combination can be used. Routes that are not matched are discarded.

Set Local Preference with Route Map (Cont.)



```
R1#  
<output omitted>  
!  
router bgp 64520  
  neighbor 10.0.0.1 route-map RM-SET-LP in  
!  
route-map RM-SET-LP permit 10  
  set local-preference 150
```

In the example above, the routing policy dictates the selection of AS 65030 as the primary way out of the AS 65040. This limitation is achieved by placing a local preference that is higher than the default local preference of 150 on all incoming announcements from AS 65040; that is, from neighbor 10.0.0.1.

Updates entering AS 65040 from AS 65030 will be processed by all BGP routers inside AS 65040 and these updates will be preferred to other updates having the default local preference.

Set AS Path with Route Map

- **Fourth BGP path selection criteria**
- Prefer shorter AS paths (only length is compared)
- Influences the outbound path selection in a multihomed AS
- Manual manipulation of AS path length—AS path prepending
- AS path prepending can be specified per neighbor by complex criteria with route maps (AS path filters, prefix lists, or other BGP attributes that match the routes in any combination)

When connections to multiple providers are required, it is important that BGP selects the optimum route for traffic to use. The optimum, or best, route may not be what the network designer intended based on design criteria, administrative policies, or corporate mandate.


It is fairly easy for an AS to select the appropriate path for outgoing traffic. It is much more complicated to influence other autonomous systems to select the appropriate path for traffic that is returning to a specific AS. It is unlikely that the operator of an AS can request changes in router configurations in another AS. This limitation makes it virtually impossible to influence another AS to select the desired path based on the weight and local preference attributes, because both options would require configuration changes in the neighboring AS.

If no BGP path selection tools are configured on the route to influence the traffic flow, BGP will use the shortest AS path—the fourth option in the path selection process. If the AS path is not manually manipulated by some administrative means, the path going over the fewest number of autonomous systems is selected by the router regardless of available bandwidth.

However, if the AS that is attempting to influence the incoming traffic flow is sending out EBGp updates with a manipulated AS path attribute over that undesired path, the receiver of this update is less likely to select it as the best because the AS path now appears to be longer. AS path prepending potentially allows the customer to influence the route selection of its service providers. The AS path is extended with multiple copies of the AS number of the sender. There is no exact mechanism to calculate the required prepended AS path length.

The benefit of manipulating AS paths to influence route selection is that the configuration that is needed is done in the AS that is requesting a desired return path.

Set AS Path with Route Map (Cont.)



```
R1#  
-----  
<output omitted>  
!  
router bgp 65040  
 neighbor 172.16.1.1 route-map RM-SET-ASPath out  
!  
route-map RM-SET-ASPath permit 10  
 set as-path prepend 65040 65040 65040
```

AS paths can be manipulated by prepending AS numbers to existing AS paths. Normally, AS path prepending is performed on outgoing EBGP updates over the undesired return path. Because the AS paths sent over the undesired link become longer than the AS path sent over the preferred path, the undesired link is now less likely to be used as the return path.

The length of the AS path is extended because additional copies of the AS number of the sender are prepended to the AS path attribute. To avoid clashes with BGP loop prevention mechanisms, no other AS number, except that of the sending AS, should be prepended to the AS path attribute.


If another AS number is prepended in the AS path, the routers in the AS that has been prepended will reject the update because of BGP loop prevention mechanisms.

Prepending can be configured on a router for all routing updates that are sent to a neighbor or only on a subset of these updates.

In the example, all updates sent to neighbor 172.16.1.1 are prepended three times with AS 65040, the AS number of the sender, thus making that path less preferable for the returning traffic. Keep in mind that if any AS on the path also does AS path prepending, the policy may not work.

MED Attribute

- The paths with the lowest MED (also called the metric) value are the most desirable.
- MED is used to advertise an exit path to be used by EBGp neighbors to reach networks owned by this AS.
- The MED attribute is **optional** and **nontransitive**.



The MED attribute, also called the metric, is an optional nontransitive attribute.

The MED is an indication to EBGp neighbors about the preferred path into an AS. The MED attribute is a dynamic way to influence another AS about which path that it should choose to reach a certain route when multiple entry points into an AS exist. A lower metric is preferred.

Unlike local preference, the MED is exchanged between autonomous systems. The MED is sent to EBGp peers. Those routers propagate the MED within their AS, and the routers within the AS use the MED but do not pass it on to the next AS. When the same update is passed on to another AS, the metric is set back to the default of 0. To change this value, use the **default-metric number** command under the BGP process. All routes that are advertised to an EBGp neighbor are set to the value specified using this command.

MED influences inbound traffic to an AS, and local preference influences outbound traffic from an AS.

By default, a router compares the MED attribute only for paths from neighbors in the same AS.

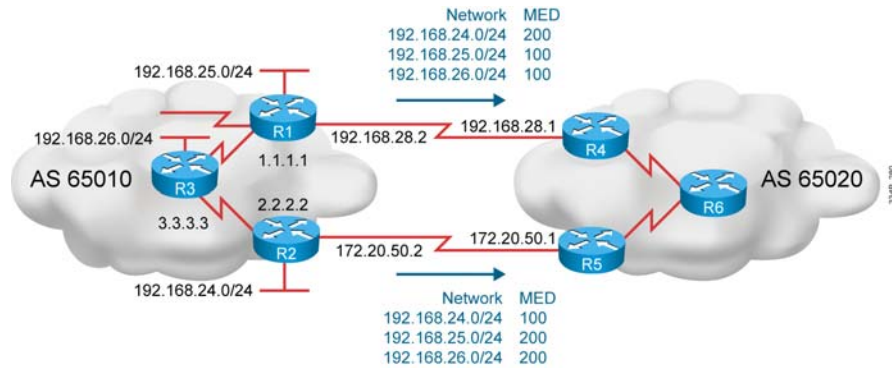
Note The MED attribute means that BGP is the only protocol that can affect how routes are sent into an AS.

In the figure above, the router R2 MED attribute is set to 150, and the router R3 MED attribute is set to 200. When router R1 receives updates from routers R2 and R3, it picks router R2 as the best next hop because its MED of 150 is less than the router R3 MED of 200.

For more details about the **default-metric number** command, please check the Cisco IOS IP Routing Protocols Command Reference on the following link:

http://www.cisco.com/en/US/docs/ios/iproute/command/reference/irp_book.html

Set MED with Route Map



This figure is used in the following configurations to demonstrate how to manipulate inbound traffic using route maps to change the BGP MED attribute. The intention of these route maps is to designate router R1 as the preferred path to reach networks 192.168.25.0/24 and 192.168.26.0/24 and to designate router R2 as the preferred path to reach network 192.168.24.0/24. The other networks should still be reachable through each router in case of a link or router failure.

Route Map for Router R1

R1#

```
router bgp 65010
neighbor 2.2.2.2 remote-as 65010
neighbor 3.3.3.3 remote-as 65010
neighbor 2.2.2.2 update-source loopback0
neighbor 3.3.3.3 update-source loopback0
neighbor 192.168.28.1 remote-as 65020
neighbor 192.168.28.1 route-map med_65020 out
!
access-list 66 permit 192.168.25.0 0.0.0.255
access-list 66 permit 192.168.26.0 0.0.0.255
!
route-map med_65020 permit 10
match ip address 66
set metric 100
!
route-map med_65020 permit 100
set metric 200
```

The MED is set outbound when a router is advertising to an EBGp neighbor. In the configuration example for router R1, a route map named “med_65020” is linked to neighbor 192.168.28.1 as an outbound route map.

When router R1 sends an update to neighbor 192.168.28.1 (router R4), it processes the outbound update through route map med_65020 and uses a set statement to change any values that are specified, as long as the preceding match statement is met in that section of the route map.

The first line of the route map is a permit statement with a sequence number of 10 for the route map med_65020; this defines the first route-map statement. The match condition for this statement checks all networks that are permitted by access list 66. The first line of access list 66 permits any networks that start with the first three octets of 192.168.25.0, and the second line of access list 66 permits networks that start with the first three octets of 192.168.26.0.

All networks that are permitted by either of these lines are set to a MED of 100. All other networks are denied by this access list (there is an implicit “deny all” at the end of all access lists), so those networks are not set to a MED of 100; their MED is not changed. These other networks must proceed to the next route map statement in the med_65020 route map.

The second statement of the route map is a permit statement with a sequence number of 100 for the route map med_65020. The route map does not have any match statements, just a set metric 200 command. This is a “permit all” statement for route maps.

Because the network administrator does not specify a match condition for this portion of the route map, all networks being processed through this section of the route map (sequence number 100) are permitted, and are set to a MED of 200. If the network administrator did not set the MED to 200, by default it would have been a MED of 0. Because 0 is less than 100, the routes with a MED of 0 would have been the preferred paths to the networks in AS 65010.

Route Map for Router R2

R2#

```
router bgp 65010
neighbor 1.1.1.1 remote-as 65010
neighbor 3.3.3.3 remote-as 65010
neighbor 1.1.1.1 update-source loopback0
neighbor 3.3.3.3 update-source loopback0
neighbor 172.20.50.1 remote-as 65020
neighbor 172.20.50.1 route-map med_65020 out
!
access-list 66 permit 192.168.24.0.0 0.0.0.255
!
route-map med_65020 permit 10
match ip address 66
set metric 100
!
route-map med_65020 permit 100
set metric 200
```

Similarly, in the configuration example for router R2, a route map named “med_65020” is linked to neighbor 172.20.50.1 (router R5) as an outbound route map.

Before router R2 sends an update to neighbor 172.20.50.1, it will process the outbound update through route map med_65020 and use a set statement to change any values that are specified, as long as the preceding match statement is met in that section of the route map.

The first line of the route map is a permit statement with a sequence number of 10 for the route map med_65020, which defines the first route-map statement. The match condition for that line checks all networks that are permitted by access list 66. Access list 66 on router R2 permits any networks that start with the first three octets of 192.168.24.0.

Any networks that are permitted by this line are set to a MED of 100. All other networks are denied by this access list, and are not set to a MED of 100. These other networks must proceed to the next route map statement in the med_65020 route map.

The second statement of the route map is a permit statement with a sequence number of 100 for the route map med_65020, but it does not have any match statements, just a **set** metric 200 command. This is a permit all statement for route maps. Because the network administrator does not specify a match condition for this portion of the route map, all networks being processed through this second statement of the route-map are permitted, but are set to a MED of 200.

If the network administrator did not set the MED to 200, by default it would have been set to a MED of 0. Because 0 is less than 100, the routes with a MED of 0 would have been the preferred paths to the networks in AS 65010.

On router R6, there are multiple paths to reach each network from AS 65010. These paths all have valid next-hop addresses, have synchronization disabled, and are loop-free. All networks have a weight of 0 and a local preference of 100, so Steps 1 and 2 do not determine the best path.

None of the routes were originated by this router or any router in AS 65020; all networks came from AS 65010, so Step 3 does not apply. All networks have an AS path of one AS (65010) and were introduced into BGP with network statements (“i” is the origin code), so Steps 4 and 5 are equal.

Step 6 states that BGP chooses the lowest MED if all preceding steps are equal or do not apply.

For network 192.168.24.0, the next hop of 172.20.50.2 has a lower MED than the next hop of 192.168.28.2. Therefore, for network 192.168.24.0, the path through 172.20.50.2 is the preferred path. For networks 192.168.25.0 and 192.168.26.0, the next hop of 192.168.28.2 has a lower MED of 100 compared to the MED of 200 through the next hop of 172.20.50.2; therefore, 192.168.28.2 is the preferred path for those networks.

Filtering of BGP Routing Updates

BGP is receiving a high number of routing updates. In order to optimize the BGP configuration, route filtering must be applied. This topic describes the steps needed to configure filtering of routing updates.

Steps to Configure BGP Route Filtering Using IP Prefix Lists

- Define traffic filtering requirements:
 - Filtering updates
 - Controlling redistribution
- Configure matching statements using:
 - mask filtering, ge, le
- Apply a prefix list to filter inbound or outbound updates

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-5-10

When planning filter configuration using prefix lists, the following steps are required and an implementation plan must be created for these steps:

- Define the traffic filtering requirements:
 - Filtering updates
 - Controlling redistribution
- Configure the matching statements:
 - Use mask filtering, ge, le
- Apply a prefix list to filter inbound or outbound updates

Configuring Filtering of BGP Routing Updates

R2 (config) #


```
ip prefix-list ANY-8to24-NET permit 0.0.0.0/0 ge 8 le 24
```

- Configure a matching statement to match all networks with the mask from /8 to /24.

R2 (config-router) #

```
neighbor 172.16.1.2 prefix-list ANY-8to24-NET in
```

- Applies an inbound prefix list filter to prevent distribution of subnets /8 to /24. The prefix list is applied to incoming advertisements from the 172.16.1.2 neighbor.



Prefix-list entries where the **ge** and **le** option are specified match any prefix within the address range that are specified using the **ge** and **le** parameters. In the example above, the prefix list named “ANY-8to25-NET” is configured to match routes from any network that has a mask length from 8 to 24 bits. The 0.0.0.0 network/length combination does not match a specific network, but is used to define any network. The combination of **ge 8 le 24** parameters specifies that any network with a mask length between 8 and 24 is a match.

A prefix list must be applied to inbound or outbound updates. In the figure above, the prefix list ANY-8to25-NET is applied to the inbound traffic to prevent distribution of BGP neighbor information. The prefix list is applied to incoming advertisements from the BGP neighbor 172.16.1.2, which prevents distribution of routes from any network that has a mask length from 8 to 24 bits.

For more details about the **neighbor prefix-list** command, please check the Cisco IOS IP Routing Protocols Command Reference on the following link:


http://www.cisco.com/en/US/docs/ios/iproute/command/reference/irp_book.html

Verify Filtering of BGP Routing Updates

```
R2#show ip prefix-list detail ANY-8to24-NET
```

```
ip prefix-list ANY-8to24-NET:  
Description: test-list  
count: 1, range entries: 1, sequences: 10 - 10, refcount: 3  
seq 10 permit 0.0.0.0/0 ge 8 le 24 (hit count: 0, refcount: 1)
```

- Configure a matching statement to match all networks with the mask from /8 to /24.



To display information about a prefix list or prefix list entries, use the **show ip prefix-list** command. The **show ip prefix-list** command output shows the count of packets being matched. The **clear ip prefix-list** command is used to clear the prefix list hit counters. The hit count is a value indicating the number of matches to a specific prefix list entry.

For more details about the **show ip prefix-list** and **clear ip prefix-list** commands, please check the Cisco IOS IP Routing Protocols Command Reference on the following link:

http://www.cisco.com/en/US/docs/ios/proute/command/reference/irp_book.html

Steps to Configure Route Filtering with a Route Map

- Define the route map
 - Define match statements
 - Define set statements
- Define route filtering using a route map

© 2009 Cisco Systems, Inc. All rights reserved.

2009-10-10-15:24

When planning filter configuration using route maps, the following steps are required and an implementation plan must be created for these steps:

- Define the route map
 - Define the match statements
 - Define the set statements
- Define the route filtering using route-maps

Using Route Maps for Filtering Routing Updates

R1 (config-router) #

```
neighbor 172.16.1.2 route-map RouteFilter in
```

- Applies a route map RouteFilter to incoming BGP updates from neighbor 172.16.1.2
- Filtering can be applied to outgoing updates.
- Prefixes that are not permitted by the route map are discarded.
- Route maps can also change the BGP attributes of incoming or outgoing updates.

© 2009 Cisco Systems, Inc. All rights reserved.

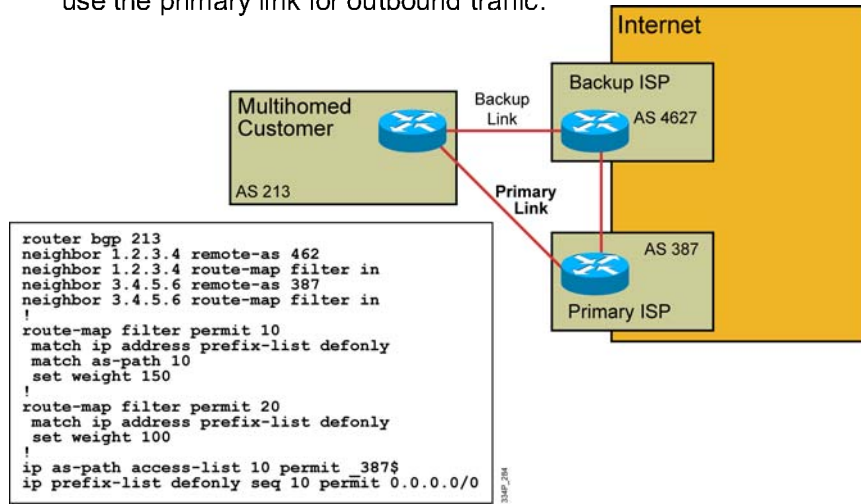
ROUTE-5-29

To apply a route map to filter incoming or outgoing BGP routes, use the **neighbor route-map** command in router configuration mode. The routes that are permitted can have their attributes set or changed by the set clauses in the **route-map** command. Setting attributes on routes is useful when attempting to influence route selection. Route map permit statements can be written to change the attributes of the permitted routes or to leave the attributes unchanged. When BGP performs route selection, the attribute values indicate that one route is preferred more than the other.

Similar route filtering could be performed for the OSPF or EIGRP routing process. In this case, the **distribute-list in** or **distribute-list out** commands are used together with the **route-map** command, which defines which updates are dropped and which updates are accepted.

Using Route Maps as BGP Filters

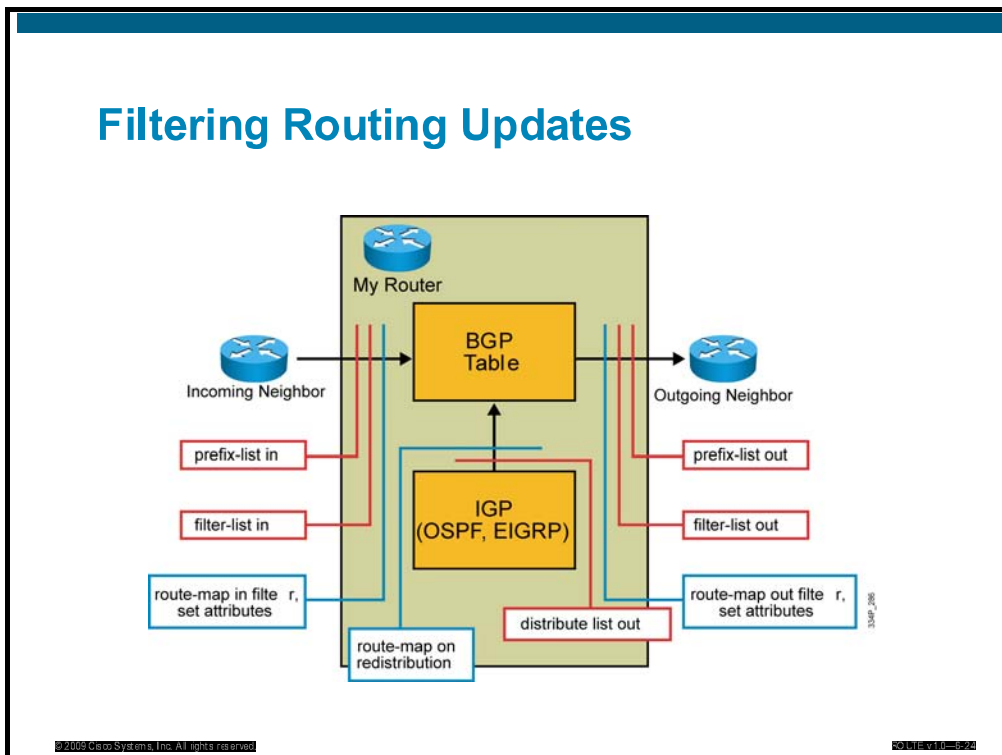
- Requirement: The customer will accept only a default route and use the primary link for outbound traffic.



In this example, the customer will accept only a default route from the ISPs, and will use the link to AS 387 as the primary link for outbound traffic.

The customer router in AS 213 is configured for BGP with two neighbors using the **neighbors remote-as** command. Both neighbors are configured with the **neighbor route-map** command to filter the incoming routing update traffic according to the route map named “filter.” Route-map filter allows a default route into the customer network, and with careful setting of the BGP weight attribute, the primary link becomes preferred. A higher value for weight is preferred and the default route coming from the ISP in AS 387 gets a weight value of 150 and the secondary default route gets a weight value of 100.

Filtering Routing Updates



As an option, filter lists, prefix lists, and route maps can be applied on either incoming or outgoing information, or in any combination in BGP. The incoming prefix list, the incoming filter list, and the incoming route map must all permit the routes that are received from a neighbor before they will be accepted into the BGP table. Outgoing routes must pass the outgoing filter list, the outgoing prefix list, and the outgoing route map before they will be transmitted to the neighbor.

When a router is configured to redistribute routing information from IGP into BGP, the routes must successfully pass any prefix list or route map applied to the redistribution process before the route is injected into the BGP table.

Summary

This topic summarizes the key points that were discussed in this lesson.

Summary

- After BGP receives updates about multiple destinations from different autonomous systems, it follows a multiple-step process for selecting the best route to reach a destination; the best route is a candidate for the routing table.
- BGP metrics are called path attributes and describe the paths to reach each network.
- BGP is receiving a high number of routing updates. In order to optimize the BGP configuration, route filtering with prefix-lists must be applied
- Route maps are used to set selected attributes for selected routes to control the outbound EBGp path selection.

Summary (cont.)

- The local preference attribute is a well-known discretionary attribute that provides an indication to routers in the AS about which path is preferred to exit the AS.
- The weight attribute is an attribute that Cisco defines for the path selection process; routes with a higher weight are preferred when multiple routes exist to the same destination.

Lesson 7

Lab 6-2 Debrief

Overview

In Lab 6-2, students manipulate the BGP path selection with route maps using local preference and MED. Students use different BGP path selection methods and compare these methods in a routed network.

After completing the lab, the instructor will lead a discussion about the lab topology, tasks, verification, and checkpoints and discuss a sample solution and different alternatives. Students will present their implementation plans and solutions.

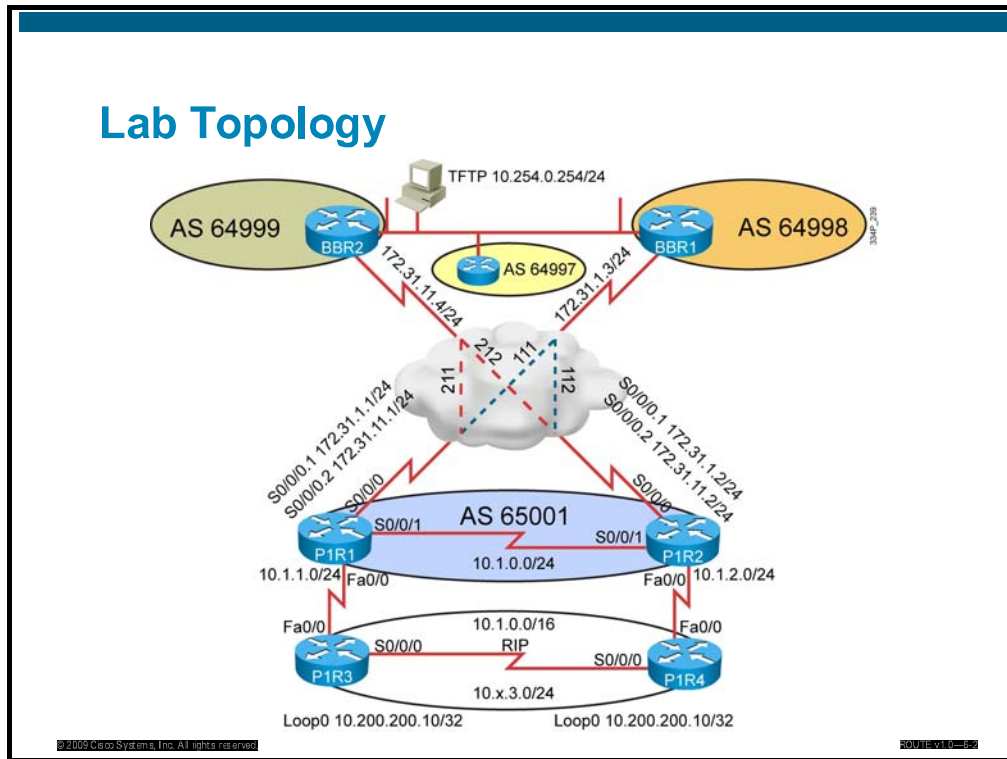
Objectives

Upon completing this lesson, you will be able to explain the lab topology, configure the BGP path manipulation using a local preference and the MED BGP attributes, create check points for configuration, verification, and find alternative solutions. This ability includes being able to meet these objectives:

- Identify the implementation and verification tasks to establish BGP adjacency in the network and manipulate the BGP path using local preference and MED BGP attributes.
- Present the sample solution and identify possible alternative solutions .

Lab Overview and Verification

This topic describes the lab topology and key checkpoints used to create a solution and to start the verification.



This figure presents the lab topology used for the configuration of BGP path manipulation lab. The topology uses four pod routers, which are members of a full Internal BGP (IBGP) relationship and two backbone routers, which are members of an External BGP (EBGP) relationship. The lab is preconfigured with addressing including loopbacks and an IGP routing protocol.

Based on the topology, students can create a BGP path manipulation configuration including a full IBGP process, advertising next hop, synchronization, and path manipulation configuration.

Lab Review: What did you accomplish?

- Task 1: Configure fully meshed IBGP
 - What steps did you take to configure fully meshed IBGP?
 - How did you deal with loopback interfaces?
- Task 2: Use the MED and local preference (LP) with route maps for BGP path manipulation
 - What is the difference between a MED and an LP implementation?
 - In what order did you execute your plan and why?
 - How did you apply a new policy to the BGP neighbor?

In the first task, students are asked to clean up previous router configurations and establish a full-mesh IBGP. When configuring the full-mesh IBGP, next hop and loopback reachability must be achieved. Because of the full-mesh IBGP, there is no need for synchronization.

In the second task, MED and local preference are used to manipulate the path. Each time the policy is changed there is a need to perform a soft reconfiguration to see the results of the new configuration.

Verification

- Is your solution working?
- What method did you follow to verify that BGP path manipulation is working correctly?
- Which commands did you use to verify the proper operation of different BGP policies?

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE-01-002-00

A common approach to verify the proper BGP path manipulation would be to take the following steps:

- Verify that all IBGP peering sessions are operational. The following commands can be used for this verification:
 - **show ip bgp summary**: Displays the status of all BGP connections
 - **show ip bgp neighbors**: Displays information about the BGP and TCP connections to neighbors
- Verify what the MED and local preference values are and which routes are inside the BGP table. The following command can be used for this verification:
 - **show ip bgp**: Displays entries in the BGP routing table
- Verify which BGP routes were entered in the routing table. The following command can be used for this verification:
 - **show ip route**: Displays entries in the IP routing table

Check Points

- Full-mesh BGP
- Check IP routing table
- Loopbacks and next hop
- Path used to reach the remote networks
- Why a different path is used
- Soft reconfiguration

During the configuration and verification phase, the network operator uses several checkpoints. After completing all of the configuration tasks, the BGP configuration can be successfully completed or additional verification and troubleshooting will be needed.

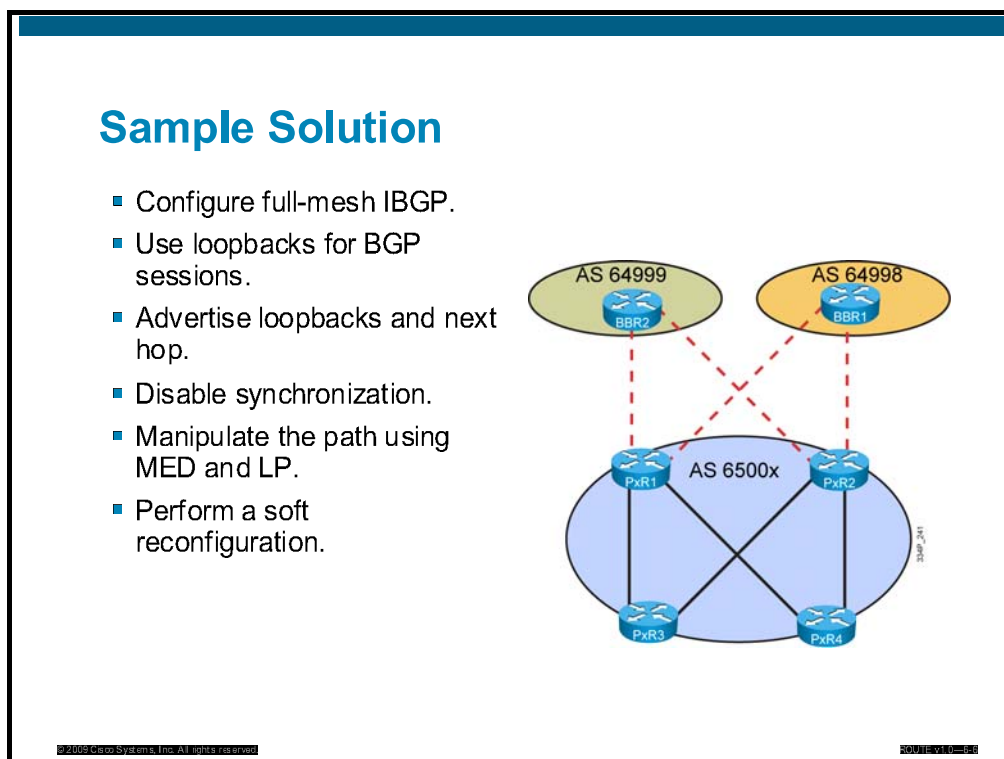
Optionally, the network operator can check the BGP configuration in different stages of the implementation using checkpoints to verify each stage.

With different checkpoints, the network operator can check for proper configuration using the following steps:

- Step 1** Check the full-mesh BGP.
- Step 2** Check the IP routing table.
- Step 3** Check the loopbacks and next hop.
- Step 4** Check the path used to reach the remote networks.
- Step 5** Check why a different path is used.
- Step 6** Check the soft reconfiguration.

Sample Solutions and Alternatives

This topic describes sample solution and other alternatives.




This sample solution includes the configuration steps necessary for each of the tasks. Different solutions are possible. However, this figure points out the important tasks required for successful configuration.

To be able to complete the lab configuration successfully, complete the following steps:

- Step 1** Configure a full-mesh IBGP.
- Step 2** Use loopbacks for the BGP sessions.
- Step 3** Advertise the loopbacks and next hop.
- Step 4** Disable synchronization.
- Step 5** Manipulate the path using MED and local preference.
- Step 6** Perform a soft reconfiguration.

Alternative Solutions

- IBGP process—full mesh or partial mesh
- Use loopbacks or other interfaces for BGP session
- Install routes into the IP routing table or advertise next hop
- BGP synchronization ON or OFF
- Using other BGP attributes (beside MED or LP) for path manipulation
- How to trigger an update of BGP table after policy change



The same or similar results can be achieved by using different configuration steps.

When configuring an IBGP neighbor relationship, it can be either a full-mesh IBGP or partial mesh. There is no mechanism in IBGP to detect whether there is an update loop. For that reason, the IBGP neighbor never sends updates to another IBGP neighbor. If there is a need for all of the IBGP neighbors to be aware of all the BGP updates, then there is a need for a full-mesh IBGP. The IBGP next-hop address is reachable by running an IGP routing protocol inside an autonomous system (AS).

Loopbacks are always on if the router is operational, and it is very convenient to establish IBGP sessions between the loopback interfaces. Loopbacks must be advertised inside an IP routing table as well as other next-hop IP addresses.

When a BGP is advertising networks to other peers, the route must be present in the IP routing table or synchronization must be turned *off*. If running a full-mesh IBGP, there is no need for synchronization to be turned *on*.

MED and local preferences can be used to manipulate the path selection. MED is used to influence the inbound EBGP path selection. Local preference can be used to influence the outbound EBGP path selection. There are other BGP attributes that can be used to influence the path selection: AS Path, weight, and others.

When the BGP policy changes, there is a need to refresh the BGP and IP routing table. Resetting the BGP session is not an elegant solution. A soft reset is available to trigger an update of the BGP table.

Q and A

What is the advantage of full-mesh IBGP?

Why is the next hop IP address important?

When is synchronization needed?

What are different reconfiguration options?

What is the difference between MED and LP?

© 2009 Cisco Systems, Inc. All rights reserved.

6-136-010000

Because IBGP peers are not sending BGP updates to other IBGP neighbors, within an AS there is a need to run a full-mesh IBGP in order to have a complete BGP table.

Next hop is used when forwarding the packets toward their destination. Every BGP update has information about the remote networks together with the next-hop address used for forwarding the packets. The next-hop address must be reachable in order to forward the packets.

The synchronization rule says that a router will not advertise routes in BGP until it learns them from an IGP. It is safe to have it off only if all routers in the transit path within the AS are running a full-mesh IBGP.

To update the BGP table, the operator can clear the BGP session or trigger the neighbor to resend the complete BGP table again. This last option is called a soft reconfiguration and requires more resources. At the same time, it is not as aggressive as a BGP session staying up all of the time.

MED is used to manipulate the inbound path to the AS and local preference is used to manipulate the outbound path from the AS.

Summary

This topic summarizes the key points that were discussed in this lesson.

Summary

- Create a good implementation plan and define the BGP requirements before configuring a BGP path manipulation.
- Several solutions exist and alternative solutions give similar or very different results.

E-Learning Training on IPv6 and Routing for Branch Offices and Remote Workers

Overview

The *Implementing Cisco IP Routing (ROUTE) v1.0* instructor-led training (ILT) course is a comprehensive learning experience. Among the learning tools available are e-learning modules that complement the classroom instructor-led content and help complete your experience with self-paced materials and demonstrations. Upon accessing the e-learning content, you will be able to reinforce the knowledge acquired in class, and witness real-life scenarios demonstrated in real routers and switches. The content structure is flexible, and you can navigate it at your own pace, at the time of your choosing, and at the depth you desire according to your level of experience.

This lesson presents the topics of IPv6 and remote access connectivity. A total of two e-learning modules are reviewed: “Implementing IPv6” and “Implementing Routing Facilities for Branch Offices and Mobile Workers.” The content includes demonstrations that cover critical topics in both areas.

Objectives

Upon completing this lesson, you will be able to discuss the e-learning modules that pertain to additional topics in IPv6 and remote access connectivity. This ability includes being able to meet these objectives:

- Describe the content of the “Implementing IPv6” e-learning module.
- Describe the content of the “Implementing Routing Facilities for Branch Offices and Mobile Workers” e-learning module.
- Understand the structure and the process of accessing and using e-learning content.

Preview of E-Learning Modules “Implementing IPv6” and “Routing Facilities for Branch Offices and Mobile Workers”

This topic describes two e-learning products that pertain to additional topics in IPv6 and remote access connectivity.

CCNP E-Learning

- Complement and enhance your classroom experience.
- Reinforce concepts and their application.
- Learn at your own pace.
- Review advanced topics.
- Experience real-life scenarios through directed demonstrations.

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE-101-001

One of the key ideas behind the design of the Cisco CCNP® curriculum is the understanding that there is no one "best" method of learning for every student. Some students prefer individual labs, while others prefer one-on-one tutoring, hands-on sessions, self-paced computer-assisted instruction, direct discovery learning or cooperative learning, and other methods.

E-learning solutions offer new approaches that complement and enhance classroom-based learning. Designed with flexibility and learning effectiveness in mind, the “Implementing IPv6” and “Implementing Routing Facilities for Branch Offices and Mobile Workers” e-learning modules are based on knowledge you acquired during your *Implementing Cisco IP Routing (ROUTE) v1.0* instructor-led training course.

Both modules include new concepts and describe these concepts starting with the fundamentals and going into more depth through directed demonstrations.

Implementing IPv6

Lesson	Description
IPv6 Addressing and Unicast	Learn the fundamentals of IPv6 technologies, focusing on static and stateless address planning and configuration, and unicast connectivity in point-to-point and multipoint links.
IPv6 with RIPng, OSPFv3, EIGRP and Redistribution	Demonstrate router configuration in scenarios of IPv6 routing, including the integration and redistribution of multiple routing protocols such as RIPng, OSPFv3, EIGRP, and BGP.
IPv6 Transition Techniques	Understand the transition path to IPv6, using techniques such as manual IPv6 tunnels, 6to4 tunnels, and ISATAP tunnels.
NAT and PAT with IPv6	Implement static and dynamic NAT-PT.

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-50

The “Implementing IPv6” e-learning module covers topics that enhance Module 6, Connection of an Enterprise Network to an ISP Network, of the *Implementing Cisco IP Routing (ROUTE) v1.0* course. It starts with an introductory lesson that discusses the benefits of the protocol and structural and operational fundamentals, such as addressing. Directed demonstrations will guide subsequent lessons on IPv6 routing and routing protocols, IPv4-to-IPv6 transition techniques including IPv6, 6to4 and Intra-Site Automatic Tunnel Addressing Protocol (ISATAP) tunnels and Network Address Translation (NAT) as well as Port Address Translation (PAT) in IPv6 scenarios. IPv6 is supported by RIP named RIP next generation (RIPng), OSPF version3 (OSPFv3) and EIGRP routing protocols.

Implementing Routing Facilities for Branch Offices and Mobile Workers

Lesson	Description
Analyzing Branch Office Designs and Planning for Branch Office Installations	Watch a discussion on design fundamentals and trends for the branch office. Overview the design scenarios that will be demonstrated in the next lesson.
Directed Demo on How to Implement Special Facilities for Branch Offices	Demonstrate router configurations in scenarios of branch office connectivity, including verification of existing services, routing around IPSec tunnels using GRE, and other routing facilities.
Lab Debrief	Gather conclusions, sample configurations, and alternative methods related to the demonstrations in the previous lesson.

© 2009 Cisco Systems, Inc. All rights reserved.

6-142-01-00000

The “Implementing Routing Facilities for Branch Offices and Mobile Workers” e-learning module also covers topics that enhance Module 6, Connection of an Enterprise Network to an ISP Network, of the *Implementing Cisco IP Routing (ROUTE) v1.0* course. It is structured in two major topics: branch office scenarios and mobile worker scenarios.

Branch office designs are analyzed through an introductory lesson that discusses trends and factors to consider in the implementation plan of new installations and upgrades. The directed demonstrations focus on scenarios involving backup connectivity, alternative routing across IPv6 tunnels using Generic Routing Encapsulation (GRE), load sharing, and other situation.

Implementing Routing Facilities for Branch Offices and Mobile Workers (cont.)

Lesson	Description
Analyzing Mobile Workers Designs and Planning for Mobile Workers Installations	Review design fundamentals and trends for the mobile worker. Overview the design scenarios that will be demonstrated in the next lesson.
Directed Demo: Implement Special Facilities for Mobile Workers.	Demonstrate router configurations in scenarios of mobile worker connectivity, including verification of existing services, IP address, and routing facilities provision for mobile workers.
Lab Debrief	Gather conclusions, sample configurations, and alternative methods related to the demonstrations in the previous lesson.

© 2009 Cisco Systems, Inc. All rights reserved.

2009-10-08

The mobile worker topic is approached in a similar way. An introductory lesson will describe trends and design considerations. This lesson is followed by a directed demonstration of remote access connectivity for mobile users, including addressing and routing considerations around the use of IP Security (IPsec) tunnels as a connectivity option.

Where to find E-Learning Modules



The “E-Learning Training on IPv6 and Routing for Branch Offices and Remote Workers” e-learning modules are available on a CD as part of your classroom materials. Please contact your instructor with questions about finding and accessing the CD.

The objectives of the “E-Learning Training on IPv6 and Routing for Branch Offices and Remote Workers” e-learning modules are part of your CCNP certification exam, and as such, candidates for the CCNP certification should review these objectives.

E-Learning Module Structure

The screenshot displays a three-panel interface for an e-learning module. The top-left panel, titled "Addressing and Topology Specifics", shows a network diagram with three routers (R1, R2, R3) and their interconnections. R1 and R2 are connected via a link with IP 10.2.7.0/24 (OSPF). R1 and R3 are connected via a link with IP 10.3.7.0/24 (EIGRP). R2 and R3 are connected via a link with IP 192.168.1.0/24. R1 is connected to a network with IP 192.168.2.0/24, and R3 is connected to a network with IP 192.168.253.0/24. The top-right panel, titled "Debrief : Alternative Configuration", shows a configuration snippet for route-maps and redistribution. A speech bubble next to the configuration says "Keep it simple!". The bottom panel shows a console output for R1, displaying policy routing statistics and a trace command for 192.168.254.1. The console output shows that the route is traced to 192.168.2.2 in 12 msec.

```
Addressing and Topology Specifics
```

```
Debrief : Alternative Configuration
```

```
route-map TAGS deny 10
match tag 1000
route-map TAGS permit 20
set tag 1000
....
router ospf 1
redistribute eigrp 1 metric 4 route-map TAGS
....
router eigrp 1
redistribute ospf 1 metric 1000000 0 255
1 800 route-map TAGS
....
```

```
R1 R2 ISP
Policy routing matches: 0 packets, 0 bytes
route-map SETTAG, permit, sequence 20
Match clauses:
Set clauses:
Policy routing matches: 0 packets, 0 bytes
R1#trace 192.168.254.1

Type escape sequence to abort.
Tracing the route to 192.168.254.1

 0  1 192.168.2.2 12 msec * 12 msec
R1#
R1#
```

Utilizing a combination of lecture sessions, animated content, lab demonstrations, and assessments, each e-learning module presents a hierarchical structure of lessons and topics. You can navigate through that structure by using an intuitive GUI that includes playback controls, slide selection, and lesson and topic selection. Using these tools, you will be able to navigate the content at your own pace.

The three-panel screen presented in the directed demonstrations allows you to focus on the device console demonstrations, while having the command syntax and the topology diagram available for verification and additional information.

Summary

This topic summarizes the key points that were discussed in this lesson.

Summary

- The content of the “Implementing IPv6” e-learning module was described.
- The content of the “Implementing Routing Facilities for Branch Offices and Mobile Workers” e-learning module was described.
- The structure and process of accessing and using e-Learning content was reviewed.

Module Summary

This topic summarizes the key points that were discussed in this module.

Module Summary

- BGP is a path-vector routing protocol that allows routing policy decisions at the AS level to be enforced.
- BGP is a policy-based routing protocol and controls traffic flow using multiple BGP path attributes.
- BGP forms EBGP relationships with external neighbors and IBGP with internal neighbors. All routers in the transit path within an AS must run fully meshed IBGP.
- When BGP is properly configured, it will establish a neighbor relationship and announce the networks with a next-hop and source IP address to other BGP routers.
- BGP performs a multistep process when selecting the best path to reach a destination.
- BGP can manipulate path selection to affect the inbound and outbound traffic policies of an AS using route maps and BGP attributes.

The Internet has proven to be a valuable tool to many companies, resulting in multiple redundant connections to many different Internet service providers (ISPs). The function of Border Gateway Protocol (BGP) is to provide alternatives to using default routes to control path selections.

References

For additional information, refer to these resources:

- RFCs 1772, 1773, 1774, 1930, 1966, 1997, 1998, 2042, 2385, 2439, 2545, 2547, 2796, 2858, 2918, 3065, 3107, 3392, 4223, and 4271.
- RFC 1518, *An Architecture for IP Address Allocation with CIDR*.
- RFC 1519, *Classless Inter-Domain Routing (CIDR): An Address Assignment and Aggregation Strategy*.
- RFC 2050, *Internet Registry IP Allocation Guidelines*.

Module Self-Check

Use the questions here to review what you learned in this module. The correct answers and solutions are found in the Module Self-Check Answer Key.

- Q1) Which type of connectivity is typically expected to and from an enterprise network? (Source: Planning the Enterprise-to-ISP Connection)
- A) one-way
 - B) two-way
 - C) unidirectional
 - D) connectivity from the clients to the Internet

- Q2) What is not a requirement of enterprise network-to-ISP connectivity? (Source: Planning the Enterprise-to-ISP Connection)
- A) public IP address space
 - B) link type and bandwidth availability
 - C) AS routing policy
 - D) connection redundancy

- Q3) List five routing update exchange options for enterprise network-to-ISP connectivity. (Source: Planning the Enterprise-to-ISP Connection)

- Q4) Which one of these statements is not the reason for using BGP as a routing update exchange mechanism? (Source: Planning the Enterprise-to-ISP Connection)
- A) A customer deploys BGP to announce its public networks.
 - B) A BGP is typically used for inter-AS routing.
 - C) Customer routers are connected to SPs PE routers.
 - D) Customer network implementation requires a complete Internet routing table.

- Q5) What are four enterprise network-to-ISP connection options? (Source: Planning the Enterprise-to-ISP Connection)

- Q6) Which is a characteristic of dual-multihomed ISP connectivity? (Source: Planning the Enterprise-to-ISP Connection)
- A) a connection to two or more different ISPs with two links per ISP
 - B) a connection to multiple ISPs with one link per ISP
 - C) the default route points to each ISP from an enterprise network
 - D) each ISP announces a default route with a different metric to the enterprise network

- Q7) What are three common ways to perform multihoming? (Choose three.) (Source: Planning the Enterprise-to-ISP Connection)
- A) Each ISP passes only a default route to the AS.
 - B) Each ISP passes a default route and provider-owned specific routes to the AS.
 - C) Each ISP passes selected provider-owned routes but no default routes to the AS.
 - D) Each ISP passes all routes to the AS.
- Q8) Which statement is true about the autonomous system (AS)? (Source: Considering the Advantages of Using BGP)
- A) The AS is a collection of networks under single administrative domain.
 - B) The AS is a collection of networks that belong to one enterprise network.
 - C) The AS requires IGP protocol to exchange routing information between autonomous systems.
 - D) EBGP neighbors must be configured within the same AS.
- Q9) What are the two typical reasons for multihoming? (Choose two.) (Source: Considering the Advantages of Using BGP)
- A) To increase the reliability of the connection to the Internet.
 - B) To increase the performance of the connection.
 - C) To increase the bandwidth of the connection.
 - D) To simplify the IGP protocol configuration.
- Q10) What is a drawback of having all of your connections to a single ISP? (Source: Considering the Advantages of Using BGP)
- A) It has redundancy with the multiple connections.
 - B) Connectivity issues in that single ISP can cause your autonomous system to lose connectivity to the Internet.
 - C) It is not tied into the routing policy of a single connection.
 - D) It has more paths to the same networks for better policy manipulation.
- Q11) An enterprise network is configured as a multihoming network and only receives a default route from each ISP. (Source: Considering the Advantages of Using BGP)
- A) true
 - B) false
- Q12) Which two conditions are valid reasons to run BGP in an AS? (Choose two.) (Source: Considering the Advantages of Using BGP)
- A) The AS has only a single connection to another AS.
 - B) Path and packet flow manipulation is required in this AS.
 - C) You have a limited understanding of BGP routing and route filtering.
 - D) The AS is an ISP.
- Q13) Which routing method best describes BGP? (Source: Considering the Advantages of Using BGP)
- A) distance vector
 - B) link-state
 - C) path-vector
 - D) hybrid of link-state and distance vector

- Q14) Which protocol does BGP use? (Source: Considering the Advantages of Using BGP)
- A) IP protocol number 88
 - B) IP protocol number 89
 - C) UDP port 520
 - D) TCP port 179
- Q15) Which four message types are defined by BGP? (Source: Considering the Advantages of Using BGP)
- _____
- _____
- _____
- _____
- Q16) Which two terms refer to routers that are configured to exchange BGP information with one another? (Choose two.) (Source: Comparing the Functions and Uses of EBGP and IBGP)
- A) BGP peer
 - B) BGP speaker
 - C) BGP router
 - D) BGP neighbor
- Q17) By default, what are two conditions for routers to be EBGP neighbors? (Choose two.) (Source: Comparing the Functions and Uses of EBGP and IBGP)
- A) routers must be in the same AS
 - B) routers must be in different autonomous systems
 - C) routers are running an IGP between them to establish an adjacency
 - D) routers are directly connected
- Q18) What are three ways to form an adjacency between IBGP neighbors by default? (Choose three.) (Source: Comparing the Functions and Uses of EBGP and IBGP)
- A) The neighbors can be directly connected.
 - B) The neighbors can be reachable from one another by static routes.
 - C) The neighbors can be reachable from one another by a dynamic internal routing protocol.
 - D) The neighbors can be in different autonomous systems.
- Q19) Which statement about IBGP is true? (Source: Comparing the Functions and Uses of EBGP and IBGP)
- A) Routes learned via IBGP are never sent to EBGP peers.
 - B) All the routers between IBGP neighbors must not be running IGP instead of BGP.
 - C) Routes learned via IBGP are never propagated to other IBGP peers.
 - D) Routes are never learned via IBGP.

Q20) Test your understanding of BGP terminology by matching statements with terms. Write the letter of the statement in front of the term that the statement describes. A statement can describe more than one term. Each term can match multiple statements, but choose only the statement that best describes the term. (Source: Explaining EBGP and IBGP)

Statement:

- A) These routers never advertise BGP routing information to IBGP neighbors.
- B) This is a set of BGP routers that are explicitly configured to exchange BGP information.
- C) This is a set of BGP routers that have a neighbor relationship with BGP routers in the same and different autonomous systems.
- D) This is a set of BGP routers that, by default, must be directly connected and must be in different autonomous systems.

Term:

- _____ 1. IBGP neighbors
- _____ 2. BGP neighbors
- _____ 3. EBGP neighbors
- _____ 4. BGP routers

Q21) What does BGP use during the best path selection process? (Source: Explaining EBGP and IBGP)

- A) speed
- B) AS routing policy
- C) bandwidth plus delay
- D) number of routers to reach a destination network

Q22) Which parameters are required for a basic BGP configuration? (Source: Configuring and Verifying Basic BGP Operations)

Q23) Which command indicates to a BGP router whether an IP address belongs to an IBGP or an EBGP neighbor? (Source: Configuring and Verifying Basic BGP Operations)

- A) **neighbor 10.1.1.1 shutdown**
- B) **neighbor 10.1.1.1 update-source Locpback0**
- C) **neighbor 10.1.1.1 remote-as 65010**
- D) **neighbor 10.1.1.1 next-hop-self**

Q24) Which command sets the source IP address of a BGP update to be the IP address of a specific interface? (Source: Configuring and Verifying Basic BGP Operations)

- A) **neighbor 10.2.2.2 shutdown**
- B) **neighbor 10.2.2.2 update-source Locpback0**
- C) **neighbor 10.2.2.2 remote-as 65020**
- D) **neighbor 10.2.2.2 next-hop-self**

- Q25) What is the result of using this command: **router bgp 65010**? (Source: Configuring and Verifying Basic BGP Operations)
- A) The BGP process starts in the router.
 - B) The BGP process starts on the interface.
 - C) The neighboring router AS is defined.
 - D) The router enters into BGP configuration mode with AS number 65010 used locally.
- Q26) The **network** command that is used in the router BGP process identifies the interfaces from which to advertise BGP updates. (Source: Configuring and Verifying Basic BGP Operations)
- A) true
 - B) false
- Q27) What is the result of using this command: **neighbor 10.3.3.3 shutdown**? (Source: Configuring and Verifying Basic BGP Operations)
- A) BGP neighbor 10.3.3.3 is administratively disabled.
 - B) It prevents BGP neighbor 10.3.3.3 from receiving BGP updates.
 - C) The BGP process on the neighbor 10.3.3.3 is disabled.
 - D) The command also requires the AS number in order to shut down the neighbor.
- Q28) Which command must be used if an EBGP neighbor is not directly connected? (Source: Configuring and Verifying Basic BGP Operations)
- A) **neighbor 10.4.4.4 ebgp-multihop**
 - B) **neighbor 10.4.4.4 update-source Locpback0**
 - C) **neighbor 10.4.4.4 remote-as 65020**
 - D) **neighbor 10.4.4.4 next-hop-self**
- Q29) What are two benefits of using BGP peer groups? (Choose two.) (Source: Configuring and Verifying Basic BGP Operations)
- A) Peer groups simplify the BGP configuration.
 - B) Peer groups make BGP updates more efficient.
 - C) Peer groups are able to find a more optimal way out of the AS.
 - D) Peer groups change BGP attributes for more BGP neighbors at a time.
- Q30) What states does BGP go through during the establishment of a BGP session? (Source: Configuring and Verifying Basic BGP Operations)
- _____
- _____
- _____
- Q31) Which state indicates that the router does not have a path to the neighbor IP address? (Source: Configuring and Verifying Basic BGP Operations)
- A) active
 - B) idle
 - C) established
 - D) open confirm

- Q32) If you configure or change the password or key used for MD5 authentication between two BGP peers, the local router will not tear down the existing session after you configure the password. (Source: Configuring and Verifying Basic BGP Operations)
- A) true
 - B) false
- Q33) Which **clear ip bgp** command is the least intrusive for resetting a BGP session after changing outbound policy for neighbor 10.5.5.5? (Source: Configuring and Verifying Basic BGP Operations)
- A) **clear ip bgp ***
 - B) **clear ip bgp 10.5.5.5 soft out**
 - C) **clear ip bgp 10.5.5.5**
 - D) **clear ip bgp 10.5.5.5 soft in**
- Q34) Place the BGP selection criteria in order from the first step to the last step evaluated to select the BGP path that is submitted to the IP routing table. (Source: Selecting an Outbound EBGP Path)
- A) _____ prefer the path with the lowest neighbor BGP router ID
 - B) _____ prefer the lowest MED
 - C) _____ prefer the shortest AS path
 - D) _____ prefer the oldest route for EBGP paths
 - E) _____ prefer the lowest origin code (IGP < EGP < incomplete)
 - F) _____ prefer the highest weight
 - G) _____ prefer the path through the closest IGP neighbor
 - H) _____ prefer the highest local preference
 - I) _____ prefer the route that was originated by the local router
 - J) _____ prefer an EBGP path over IBGP path
 - K) _____ prefer the lowest neighbor IP address
- Q35) Which description applies to the AS path attribute? (Source: Selecting an Outbound EBGP Path)
- A) well-known mandatory
 - B) well-known discretionary
 - C) optional transitive
 - D) optional nontransitive
- Q36) Which description applies to the next-hop attribute? (Source: Selecting an Outbound EBGP Path)
- A) well-known mandatory
 - B) well-known discretionary
 - C) optional transitive
 - D) optional nontransitive
- Q37) Which description applies to the origin attribute? (Source: Selecting an Outbound EBGP Path)
- A) well-known mandatory
 - B) well-known discretionary
 - C) optional transitive
 - D) optional nontransitive

- Q38) Which description applies to the local preference attribute? (Source: Selecting an Outbound EBGPath)
- A) well-known mandatory
 - B) well-known discretionary
 - C) optional transitive
 - D) optional nontransitive
- Q39) Which description applies to the MED attribute? (Source: Selecting an Outbound EBGPath)
- A) well-known mandatory
 - B) well-known discretionary
 - C) optional transitive
 - D) optional nontransitive
- Q40) Which description applies to the weight attribute? (Source: Selecting an Outbound EBGPath)
- A) well-known mandatory
 - B) well-known discretionary
 - C) optional transitive
 - D) proprietary to Cisco and not advertised to other BGP routers
- Q41) Which two statements are true regarding local preference? (Choose two.) (Source: Selecting an Outbound EBGPath)
- A) The higher value for local preference is preferred.
 - B) Local preference is used only between EBG neighbors.
 - C) The lower value for local preference is preferred.
 - D) Local preference is used only between IBGP neighbors.
- Q42) Which two statements are true regarding weight? (Choose two.) (Source: Selecting an Outbound EBGPath)
- A) The lower value for weight is preferred.
 - B) The higher value for weight is preferred.
 - C) Weight is used only between IBGP neighbors.
 - D) Weight is used only locally inside the router.
- Q43) Which two statements are true regarding the MED? (Choose two.) (Source: Manipulating Inbound EBGPaths Selection)
- A) The higher value for the MED is preferred.
 - B) The lower value for the MED is preferred.
 - C) The MED is exchanged between autonomous systems.
 - D) The MED is local to an AS.
- Q44) Which two statements are true regarding the AS Path? (Choose two.) (Manipulating Inbound EBGPaths Selection)
- A) The shorter AS path is preferred.
 - B) The longer AS path is preferred.
 - C) The AS path is prepended and exchanged between autonomous systems.
 - D) The AS path is local to an AS.

- Q45) Which command changes the MED for all routes? (Source: Manipulating Inbound EBG P Paths Selection)
- A) **bgp med** *number*
 - B) **default-metric** *number*
 - C) **bgp default-metric** *number*
 - D) **set med** *number*
- Q46) The MED is used to decide how to enter an AS from neighboring autonomous systems, when multiple paths exist between two autonomous systems. (Source: Manipulating Inbound EBG P Paths Selection)
- A) true
 - B) false
- Q47) The MED is set inbound when a router is receiving router updates from an EBG P neighbor. (Source: Manipulating Inbound EBG P Paths Selection)
- A) true
 - B) false
- Q48) The length of the AS path is extended because additional copies of the AS number of the sender are prepended to the original AS path attribute. To avoid clashes with BGP loop prevention mechanisms, no other AS number, except that of the neighboring AS, should be prepended to the AS path attribute. (Source: Manipulating Inbound EBG P Paths selection)
- A) true
 - B) false

Module Self-Check Answer Key

- Q1) B
- Q2) C
- Q3) Static routes, Common IGP, MPLS VPNs, Circuit Emulation, BGP
- Q4) C
- Q5) Single-homed, Dual-homed, Multihomed, Dual-multihomed
- Q6) A
- Q7) A, B, D
- Q8) A
- Q9) A, B
- Q10) B
- Q11) A
- Q12) B, D
- Q13) C
- Q14) D
- Q15) Open, Keepalive, Update, Notification
- Q16) A, D
- Q17) B, D
- Q18) A, B, C
- Q19) C
- Q20) 1-A
2-C
3-D
4-B
- Q21) B
- Q22) neighbors (peers) involved, AS numbers used, IP Addresses used and Networks, which need to be advertised
- Q23) C
- Q24) B
- Q25) D
- Q26) B
- Q27) A
- Q28) A
- Q29) A, B
- Q30) Idle, Connect, Open sent, Open confirm, Established
- Q31) B
- Q32) A

- Q33) B
- Q34) 1-F
2-H
3-I
4-C
5-E
6-B
7-J
8-G
9-D
10-A
11-K
- Q35) A
- Q36) A
- Q37) A
- Q38) B
- Q39) D
- Q40) D
- Q41) A, D
- Q42) B, D
- Q43) B, C
- Q44) A, C
- Q45) B
- Q46) A
- Q47) B
- Q48) B