MSPRP

Maintaining Cisco Service Provider Routing Protocols

Volume 1

Version 1.0

Student Guide

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.**

Students, this letter describes important course evaluation access information!

Welcome to Cisco Systems Learning. Through the Cisco Learning Partner Program, Cisco Systems is committed to bringing you the highest-quality training in the industry. Cisco learning products are designed to advance your professional goals and give you the expertise you need to build and maintain strategic networks.

Cisco relies on customer feedback to guide business decisions; therefore, your valuable input will help shape future Cisco course curricula, products, and training offerings. We would appreciate a few minutes of your time to complete a brief Cisco online course evaluation of your instructor and the course materials in this student kit. On the final day of class, your instructor will provide you with a URL directing you to a short post-course evaluation. If there is no Internet access in the classroom, please complete the evaluation within the next 48 hours or as soon as you can access the web.

On behalf of Cisco, thank you for choosing Cisco Learning Partners for your Internet technology training.

Sincerely,

Cisco Systems Learning

Table of Contents

Volume 1

Course Introduction

Overview

The *Maintaining Cisco Service Provider Routing Protocols* (MSPRP) course is a five-day instructor-led training (ILT) course. It is intended to help prepare individuals to pass the Cisco Certified Network Professional: Service Provider Operations (CCNP® SP Operations) certification exams. This course and certification will also prepare and validate learner proficiency in monitoring and troubleshooting interior gateway protocols (IGPs) and Border Gateway Protocol (BGP) on the service provider core Cisco IP Next-Generation Network (NGN) network.

This course is designed to provide learners with the knowledge needed to provide support in a service-provider environment using an IGP such as Open Shortest Path First (OSPF) or Intermediate System-to-Intermediate System (IS-IS) and BGP. It also provides learners with an understanding of advanced routing policies using route maps with Cisco IOS Software and Routing Policy Language (RPL) with Cisco IOS XR Software. The course also focuses on monitoring and troubleshooting these technologies in service provider environments. The course is intended to help prepare students for the CCNP SP Operations certification exam.

Learner Skills and Knowledge

This subtopic lists the skills and knowledge that learners must possess to benefit fully from the course. The subtopic also includes recommended Cisco learning offerings that learners should first complete to benefit fully from this course.

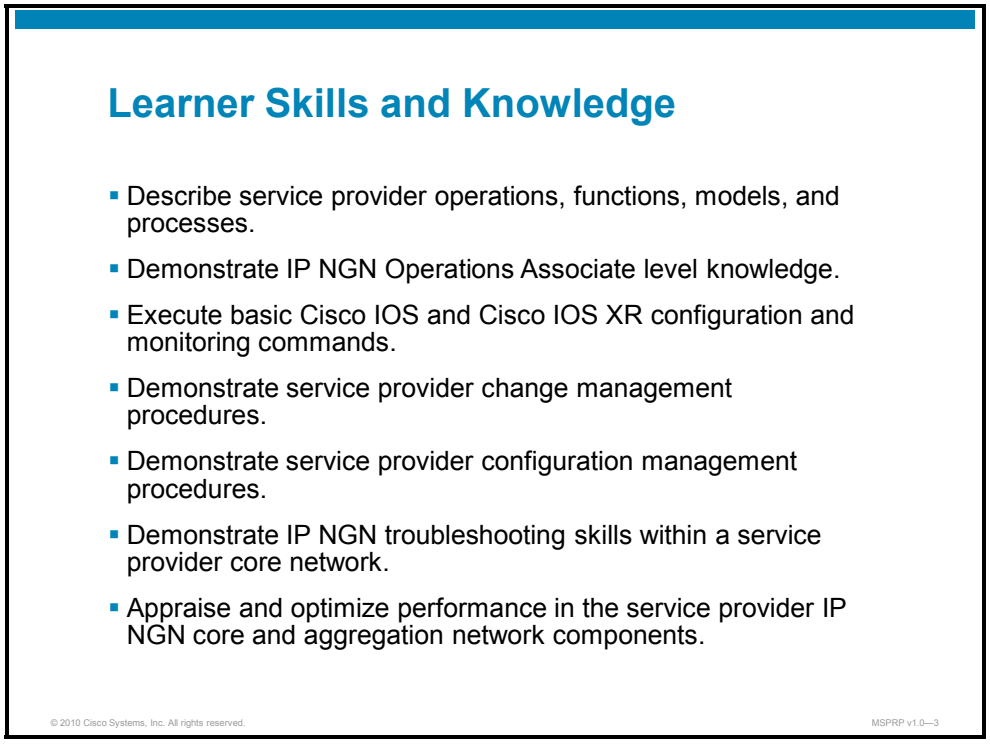

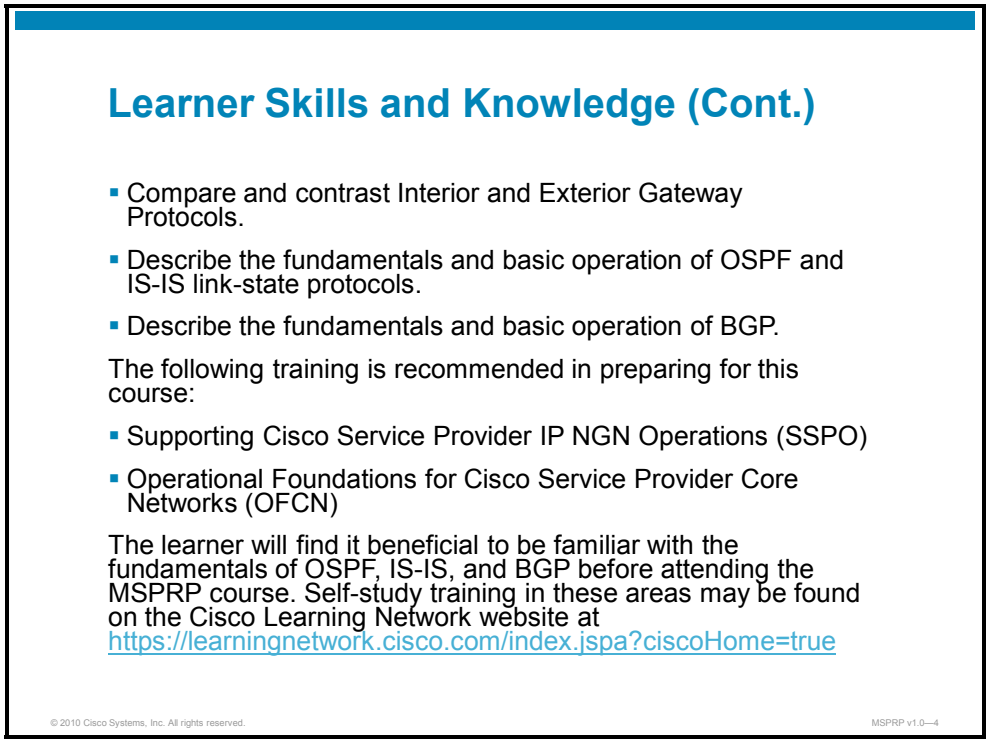

Course Goal and Objectives

This topic describes the course goal and objectives.

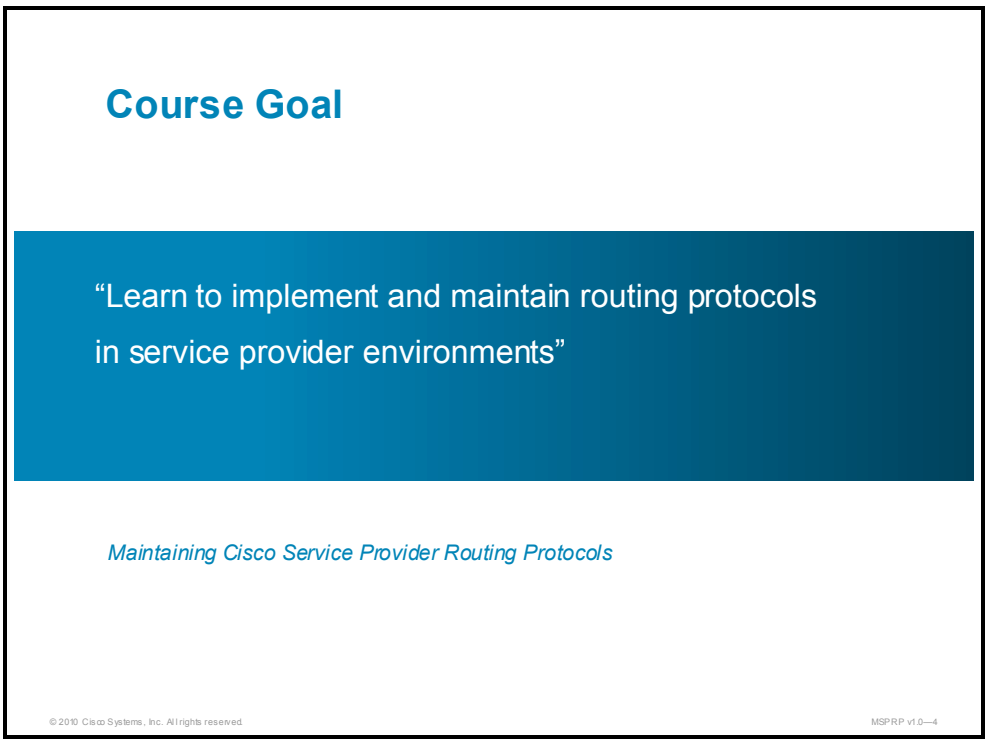

Upon completing this course, you will be able to meet these objectives:

- Identify the typical routing requirements in service provider networks and list the routing solutions and describe the change, performance, and fault management procedures
- Change the routing configuration based on designs and implementation templates as well as Information Technology Infrastructure Library (ITIL®)-based change management processes and procedures
- Gather and interpret performance statistics for routing protocols using various tools such as **show** commands and Simple Network Management Protocol (SNMP)-based monitoring tools
- Perform fault management for routing protocols using various tools, such as **show** commands and monitoring tools

Course Flow

This topic presents the suggested flow of the course materials.

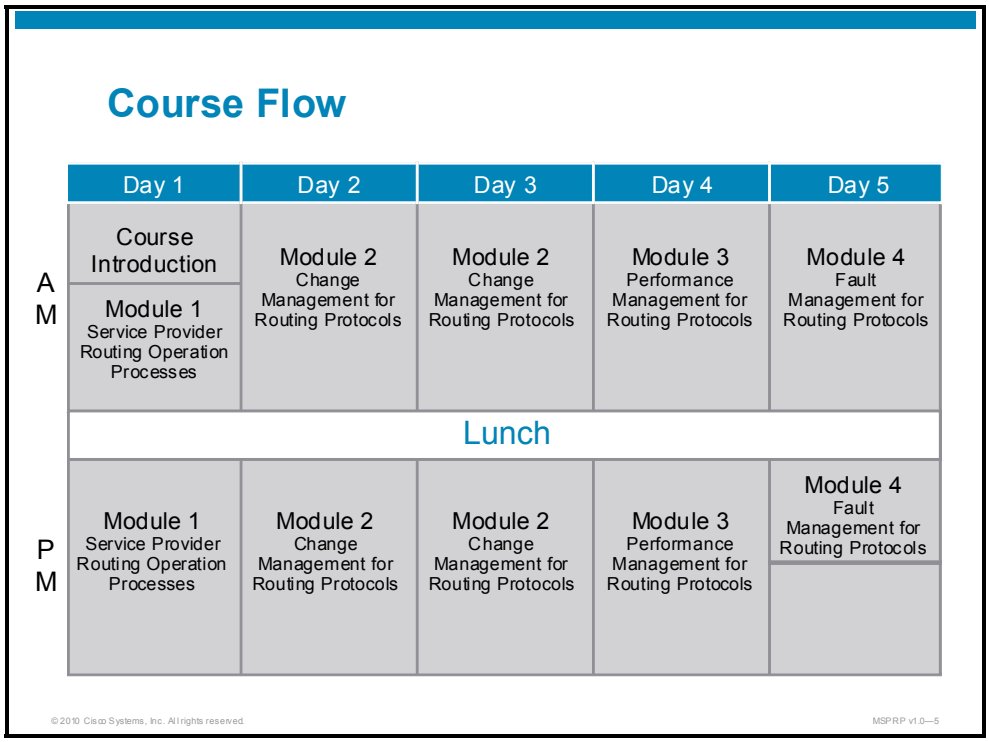

The schedule reflects the recommended structure for this course. This structure allows enough time for the instructor to present the course information and for you to work through the lab activities. The exact timing of the subject materials and labs depends on the pace of your specific class.

Additional References

This topic presents the Cisco icons and symbols that are used in this course, as well as information on where to find additional technical references.

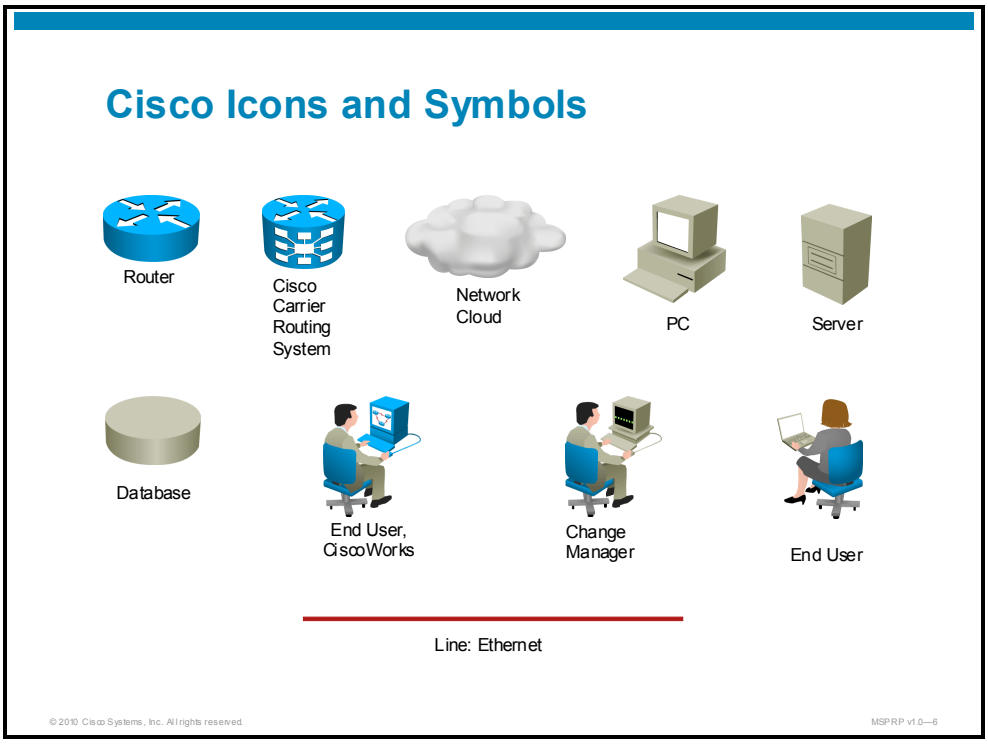

Cisco Glossary of Terms

For additional information on Cisco terminology, refer to the *Cisco Internetworking Terms and Acronyms* glossary of terms at http://www.cisco.com/univercd/cc/td/doc/cisintwk/ita/index.htm.

Your Training Curriculum

This topic presents the training curriculum for this course.

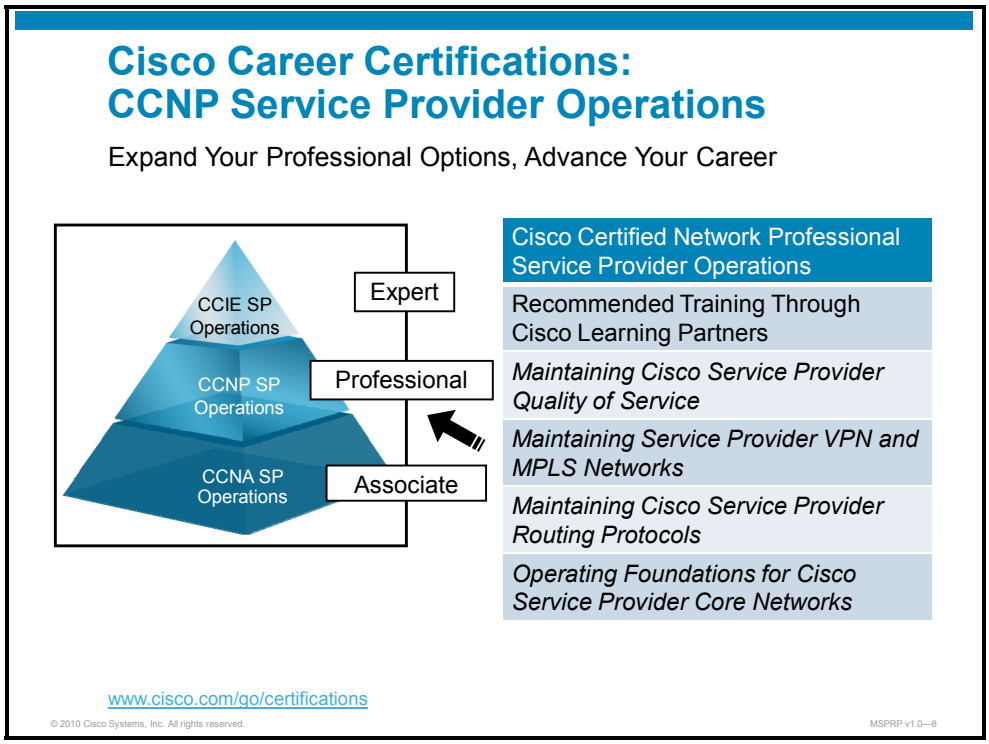

You are encouraged to join the Cisco Certification Community, a discussion forum open to anyone holding a valid Cisco Career Certification (such as Cisco CCIE®, CCNA®, CCDA®, CCNP, CCDP®, CCIP®, CCVP™, or CCSP™). It provides a gathering place for Cisco certified professionals to share questions, suggestions, and information about Cisco Career Certification programs and other certification-related topics. For more information, visit http://www.cisco.com/go/certifications.

Module 1

Service Provider Routing Operation Processes

Overview

This module identifies service provider routing requirements, solutions and processes for change, performance, and fault management. The module covers a broad range of issues concerning routing protocols that are used in service provider environments. The main function of interior gateway protocols (IGPs) and Border Gateway Protocol (BGP) in service provider network will be discussed. The module also presents the mechanisms that are available in combination with routing protocols to filter routing information or to implement desired routing policies. As you will see, you also need tools in a service provider environment to successfully manage and monitor the network devices and services. The end of the module provides a brief description of the Information Technology Infrastructure Library® (ITIL) standard that is widely used in enterprise and service provider environments.

Module Objectives

Upon completing this module, you will be able to identify the typical routing requirements in service provider networks. You will also be able to list the routing solutions and describe the change, performance, and fault management procedures. This ability includes being able to meet these objectives:

- Identify the main characteristics of routing protocols that are used in the service provider environments
- Identify the main Cisco IOS and Cisco IOS XR Software tools that are used in service provider environments in combination with routing protocols
- Identify the main external tools that are used in service provider environments in combination with routing protocols
- Use ITIL®-based processes for change, performance, and fault management of routing protocols in service provider environments

Lesson 1

Understanding Service Provider Routing Protocols

Overview

This lesson provides the main characteristics of routing protocols that are used in service provider environments. The lesson describes how a service provider ensures IP connectivity within the Internet—to end customers and to other service providers. The lesson explains the need for exchanging Internet routing information via Border Gateway Protocol (BGP). On the other hand, interior gateway protocols (IGPs) are responsible for providing IP connectivity within each autonomous system (AS). The most commonly used IGP protocols in service provider networks are Open Shortest Path First (OSPF) and Intermediate System-to-Intermediate System (IS-IS). Both are briefly described, as is the BGP routing protocol.

Objectives

Upon completing this lesson, you will be able to identify the main characteristics of routing protocols that are used in service provider environments. This ability includes being able to meet these objectives:

- Describe the characteristics and requirements for routing protocols in service provider environments
- Describe the characteristics of OSPF in service provider environments
- Describe the characteristics of IS-IS in service provider environments
- Describe the characteristics of BGP in service provider environments

Overview of Routing Protocols

This topic describes the characteristics and requirements for routing protocols in service provider environments.

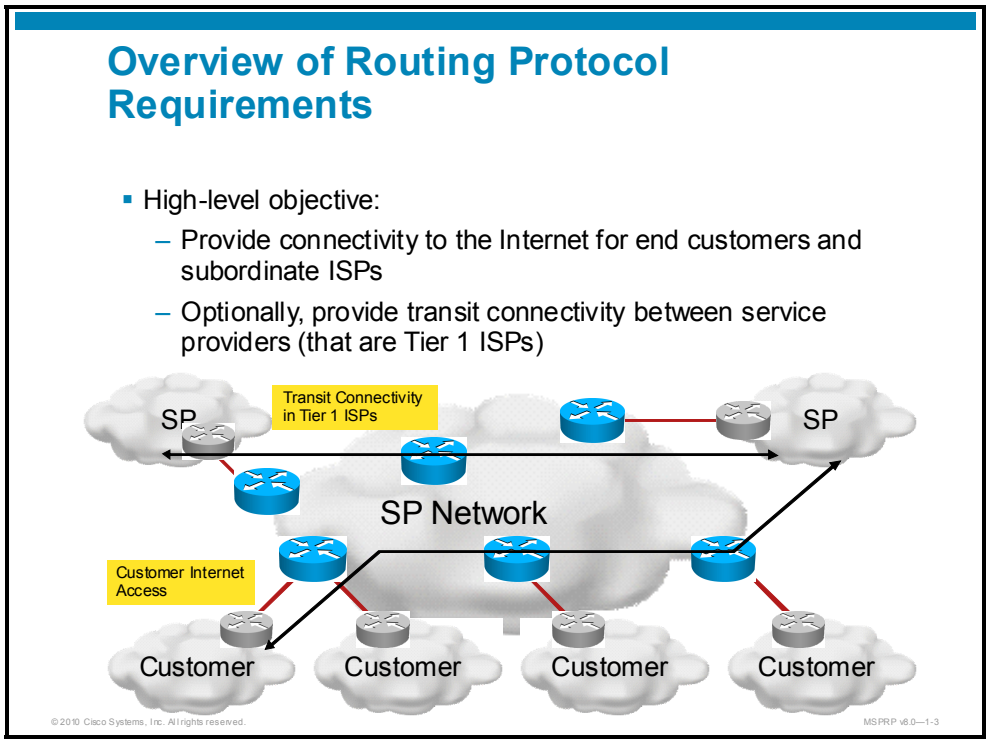

The course covers a broad range of issues concerning routing protocols that are used in service provider environments. Depending on the type of service provider, there are many different connectivity requirements, which can be summarized as follows:

- Provide Internet connectivity to end customers or subordinate ISPs
- Provide transit connectivity to peering ISPs (Tier 1) and subordinate ISPs (Tier 2)

In addition, you must consider local routing within the service provider network to ensure reachability for all local addresses.

Overview of Routing Protocols

- Interior Gateway Protocol (IGP) to exchange local routing information
- Border Gateway Protocol (BGP) to exchange external routing information

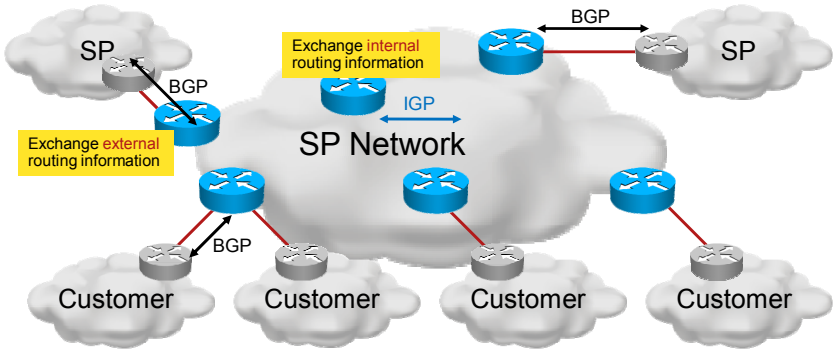

An ISP uses BGP to exchange Internet routing information with other ISPs and with those customers that require it. BGP can be configured:

- Only to propagate a default route (for example, to customers)
- By a portion of the complete Internet routing table (for example, multihomed customers)
- By the entire Internet routing table (for example, multihomed customers, subordinate ISPs)

An IGP is used to provide connectivity within an AS. The most important function of an IGP is to provide reachability to BGP neighbors and BGP next-hop addresses.

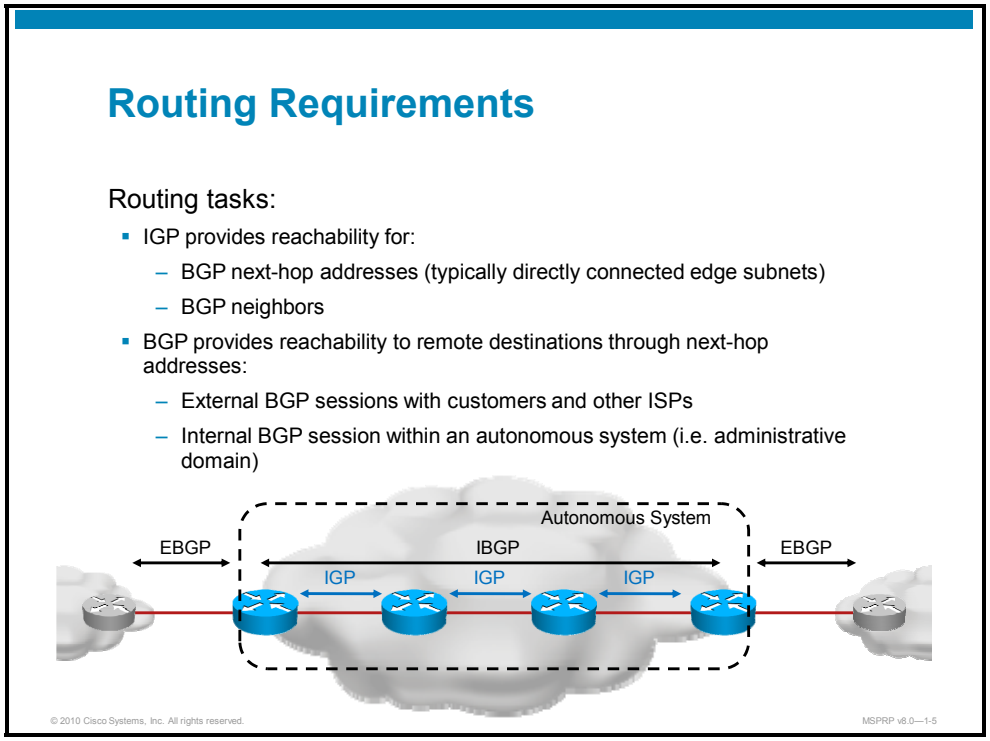

There are two characteristics of BGP that require the assistance of an IGP:

- BGP next-hop addresses do not change as the BGP routes are propagated through an AS.
- Internal Border Gateway Protocol (IBGP) neighbors can be several hops away from each other (External Border Gateway Protocol, or EBGP, neighbors are typically reachable through a directly connected edge subnet).

An IGP is therefore needed to propagate the following:

- Next-hop addresses (IP addresses of EBGP neighbors) throughout the AS
- IP addresses of internal BGP neighbors (typically loopback addresses)

The simplified illustration shows the three components in a service provider routing environment:

- EBGP sessions with other autonomous systems to exchange the Internet routing information
- Internal BGP sessions to carry external routing information across the service provider AS to all routers that require it
- IGP to provide the reachability of next-hop and neighbor addresses

Routing Example

Part 1: BGP

1. R1 receives an external BGP update: 209.165.201.0/24; the next hop is 192.168.200.2.
2. R4 receives an internal BGP update:
   - By default, the next-hop address does not change.
   - Optionally, BGP on R1 can be configured to change the next-hop address to its own address (typically a loopback address).
3. R4 forwards the update and changes the next-hop address to 192.168.11.1.

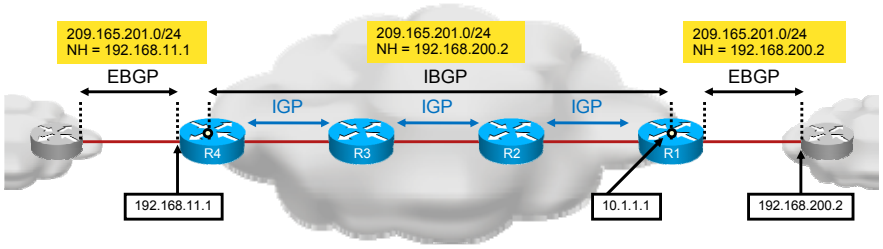

The figure illustrates a sample network in which a route is received from an AS over an EBGP session. The route is then forwarded to all other internal routers over an Internal Border Gateway Protocol (IBGP) session. Egress routers will forward the route to other external neighbors.

To change the behavior, you might, for example, configure router R1 with the **neighbor R4 next-hop-self** command. Then R1 would change the next-hop attribute to its loopback address, R4. The loopback address is used as the source address for IBGP sessions.

After R4 sends the update out to an external neighbor, it will change the next-hop address to its own IP address used for the EBGP session with the external neighbor. This is the default behavior.

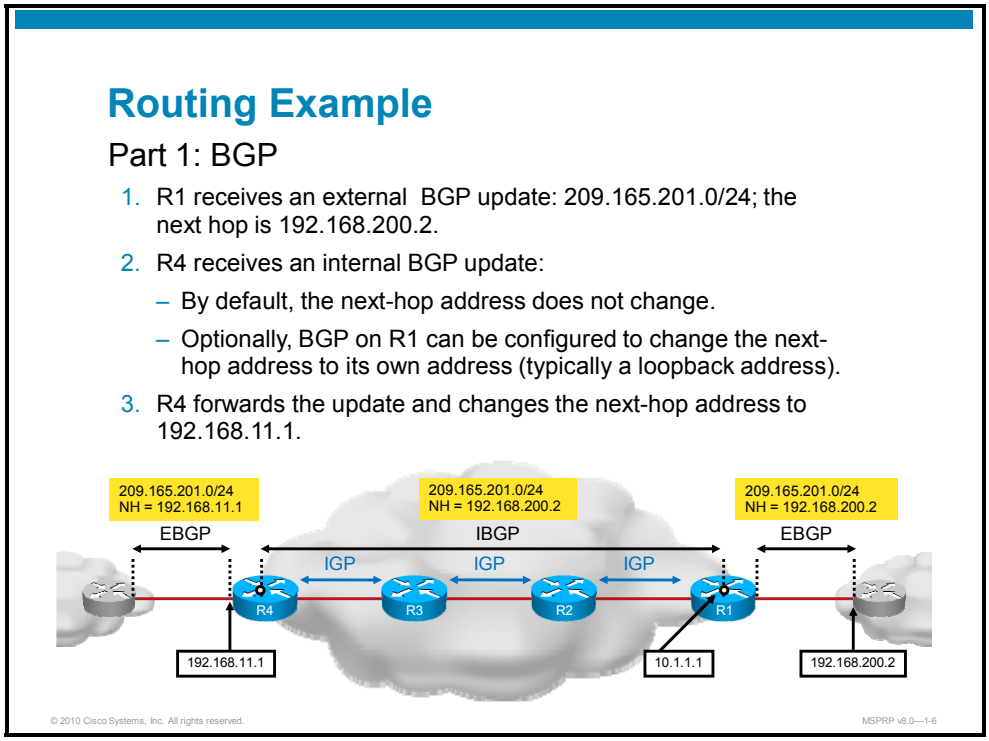

The second illustration shows the required functionality of an IGP in order to support the BGP functionality.

The IGP is propagating two important addresses throughout the AS:

- The IP address of the external neighbor, which is also the next-hop address carried within BGP updates coming from the external neighbor. Note that this part is optional if you use the alternative method (using the next-hop-self feature).
- The loopback IP address of the ingress BGP router that is used for all IBGP sessions.

If the same service provider network is also used to provide other services, such as Multiprotocol Label Switching (MPLS)-based virtual private networks (VPNs) or MPLS Traffic Engineering (MPLS TE), you must *not* summarize the next-hop IP addresses in the backbone. Doing so would break MPLS label-switched paths (LSPs).

Routing Example (Cont.)

Part 3: Routing Table

End-to-end connectivity is provided thorough recursive routing table lookups (optimized for Cisco Express Forwarding):

- BGP for end prefixes
- IGP for BGP next-hop reachability

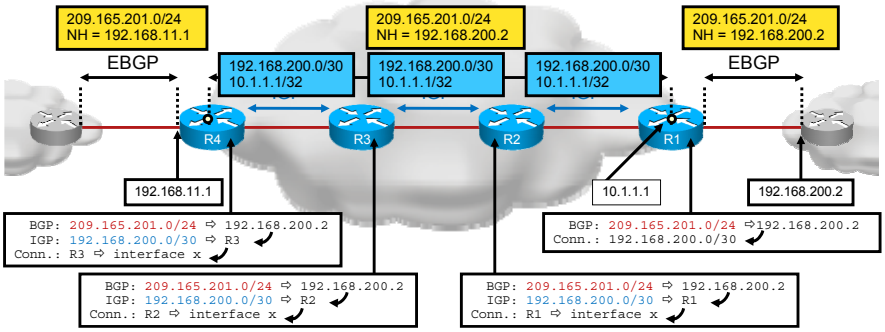

The figure illustrates the final result of BGP and IGP routing information propagation as reflected by the routing table. A recursive set of routes can be observed in the routing tables of all routers:

- EBGP routes point to BGP next-hop addresses. Router R1 receives multiple external routes, which all use the same next-hop address that is the address of the EBGP peer.
- BGP next-hop addresses point to these peers:
  - Directly connected external peer on ingress edge routers (router R1 in the example)
  - Nonadjacent addresses reachable via the IGP (routers R2, R3, and R4 in the example, which require reachability to the BGP next-hop address via the IGP)
- IGP peers are reachable through an attached link.

Note For performance reasons, routers do not perform recursive lookup when forwarding packets. Cisco Express Forwarding is used to optimize the forwarding table for performance.

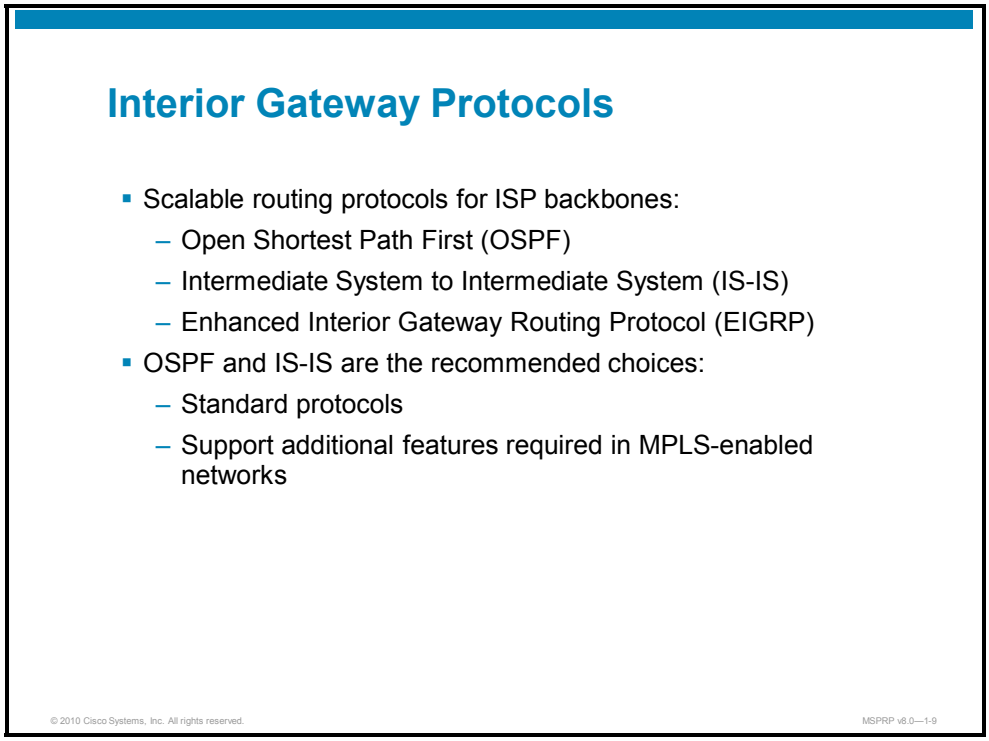

There are, in general, three routing protocols that satisfy the main service provider requirements for an IGP:

- Scalability
- Performance

There are three commonly used IGPs:

- OSPF
- IS-IS
- EIGRP

Most service providers today use either OSPF or IS-IS for two reasons:

- EIGRP is a Cisco proprietary protocol and may hinder interoperability with devices from other vendors.
- MPLS TE requires the help of a link-state protocol (EIGRP is a distance-vector protocol).

Overview of OSPF

This topic describes the characteristics of OSPF in service provider environments.

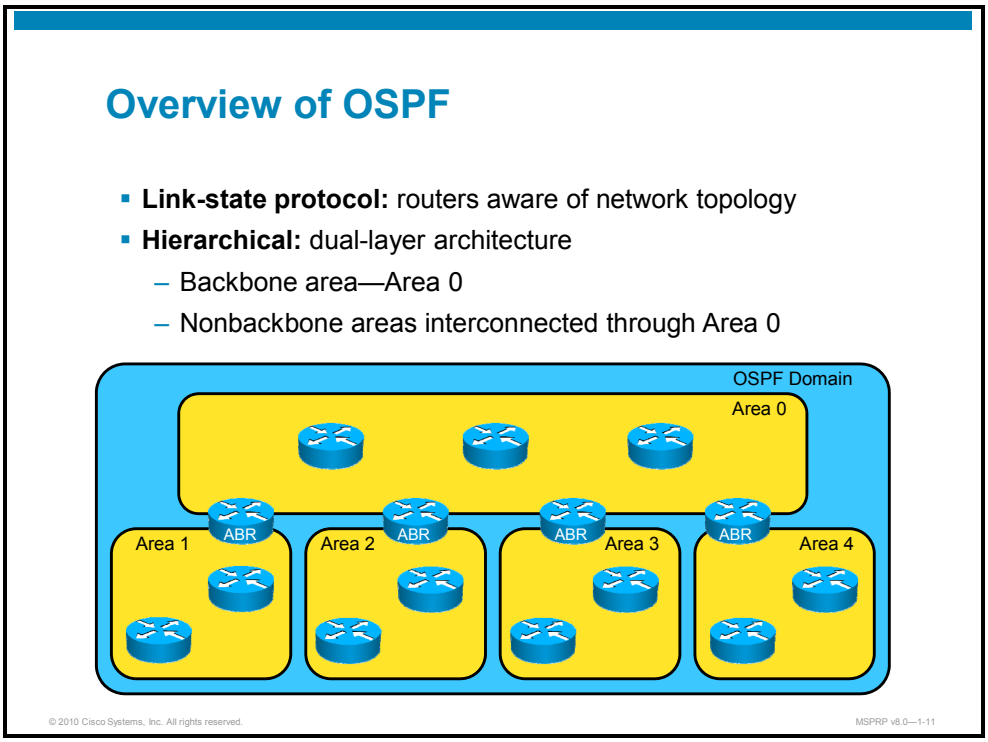

The OSPF protocol was developed due to a need in the Internet community to introduce a high-functionality, nonproprietary IGP for the TCP/IP protocol family. The discussion of the creation of a common interoperable IGP for the Internet started in 1988 and did not get formalized until 1991. At that time, the OSPF working group requested that OSPF be considered for advancement to Draft Internet Standard.

The OSPF protocol is based on link-state technology, which is a departure from the Bellman-Ford vector-based algorithms that are used in traditional Internet routing protocols such as Routing Information Protocol (RIP). OSPF has introduced new concepts such as authentication of routing updates, variable-length subnet masks (VLSMs), and route summarization.

OSPF uses a link-state algorithm in order to build and calculate the shortest path to all known destinations. The algorithm by itself is quite complicated. The following is a high-level, simplified way of looking at the various steps of the algorithm:

- Upon initialization or due to any change in routing information, a router will generate a link-state advertisement. This advertisement will represent the collection of all link states on that router.
- All routers will exchange link states using flooding. Each router that receives a link-state update should store a copy in its link-state database and then propagate the update to other routers.
- After the database of each router is completed, the router will calculate a shortest path tree to all destinations. The router uses the Dijkstra algorithm to calculate the shortest path tree. The destinations, the associated costs, and the next hops to reach those destinations will form the IP routing table.
- If no changes in the OSPF network occur, such as a change to the cost of a link or the addition or deletion of a network, OSPF should be quiet.

OSPF uses flooding to exchange link-state updates between routers. Any change in routing information is flooded to all routers in the network. Areas are introduced to put a boundary on the growth of link-state updates. Flooding and calculation of the Dijkstra algorithm on a router is limited to links within an area. All routers within an area have the exact link-state database. Routers that belong to multiple areas, and connect these areas to the backbone area, are called Area Border Routers (ABRs). ABRs must therefore maintain information describing the backbone areas and any other attached areas.

ABRs also provide mechanisms for aggregating routes and cutting down on the unnecessary propagation of subnet information.

OSPF Characteristics

- Each router has topology information for all areas in which it resides
- Many design choices exist:
  - Different types of nonbackbone areas
  - Different types of OSPF adjacencies
  - Routing information exchanged through different types of link-state advertisements (LSAs)

OSPF defines a hierarchy of two levels, in which Area 0 is the top level. Area 0 interconnects other areas, which can be of different types. Depending on the route origin and area of origin, different types of link-state advertisements (LSAs) will be generated and different rules will apply to different LSAs when crossing area borders.

Additionally, OSPF adjacencies have several modes of operation, depending on link types and configuration.

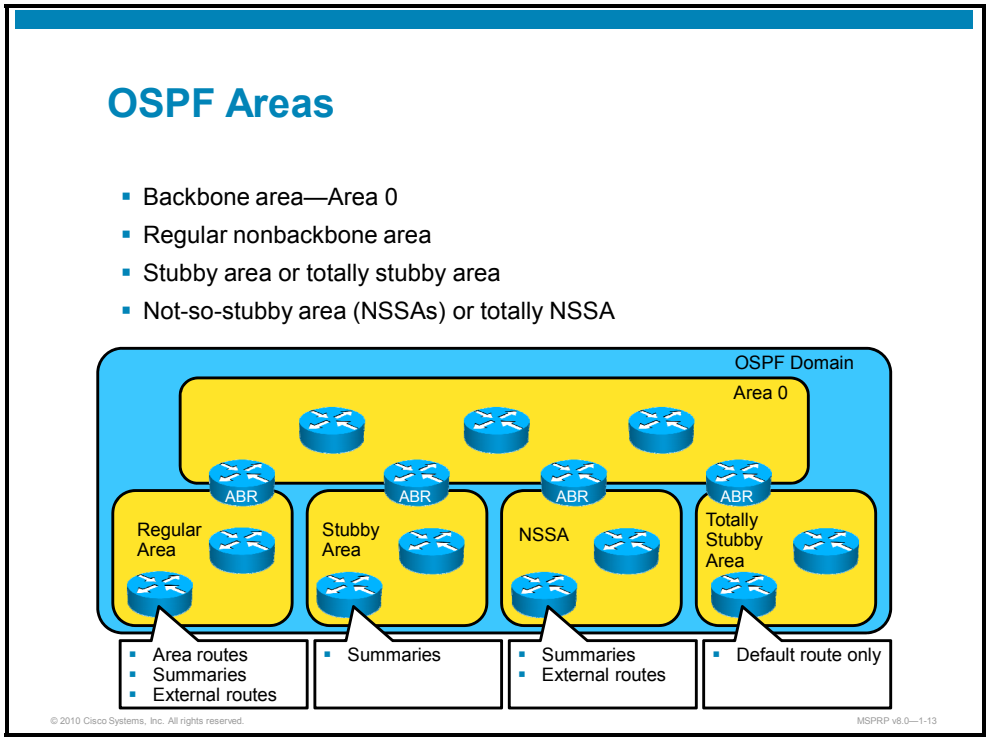

In general, there are six different types of areas:

- Backbone area or Area 0, which typically carries all the routing information
- Regular nonbackbone areas
- Stubby areas that do not originate or receive any external routes
- Totally stubby areas that do not originate (redistribution for other protocols) or receive any external routes or summaries (only the default route)
- Not-so-stubby areas (NSSAs) that can originate external routes (redistribution for other protocols) but do not receive them from other OSPF areas.
- NSSAs that can originate external routes (redistribution for other protocols), but do not receive them from other areas or receive summaries (only the default route).

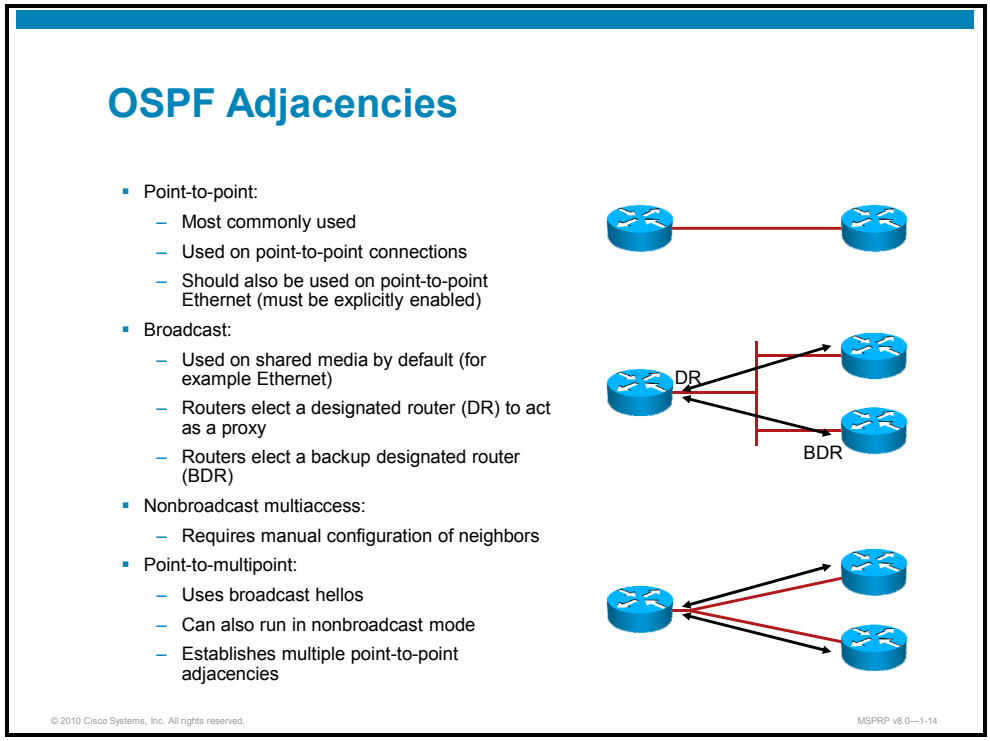

OSPF uses IP multicast addresses 224.0.0.5 and 224.0.0.6 to communicate between OSPF neighbors. Depending on the type of media and topology, you can select one of the following adjacency modes:

- Point-to-point adjacency is used for point-to-point links.
- Broadcast mode is used for adjacencies on shared media such as Ethernet. A designated router (DR) is elected to act as proxy between other routers. A backup designated router (BDR) is also elected to take over the DR functionality in case the primary DR fails.
- Nonbroadcast multiaccess mode requires manual configuration of neighbors and was sometimes used on multiaccess media such as point-to-multipoint Frame Relay or ATM links.
- Point-to-multipoint is also designed for multiaccess media, but automates the finding of neighbors and establishes a set of point-to-point neighbor relationships.

Most core links in modern service provider networks are point-to-point, even if Ethernet is used. Use the point-to-point mode on such links to avoid introducing DR election and virtual nodes into the Shortest Path First (SPF) calculation.

On shared media such as Ethernet, you can retain the default broadcast mode.

OSPF Adjacencies (Cont.)

- Hello packets with multicast address 224.0.0.5 are used to find neighbors
- DR and BDR communication uses 224.0.0.6
- Neighbors must agree on the following:
  - Area ID
  - Authentication method and password
  - Hello and dead intervals
  - Flags defining a stub area
  - IP address and subnet mask
  - Interface MTU
- Router IDs are used to uniquely identify a router

OSPF discovers neighbors using hello packets on multicast address 224.0.0.5, while DR and BDR routers communicate using multicast address 224.0.0.6. For OSPF to work correctly, neighbors must agree on the area ID, authentication method and password, hello and dead intervals, flags defining a stub area, IP address and subnet mask, and interface maximum transmission unit (MTU). OSPF uses the router ID to uniquely identify a router.

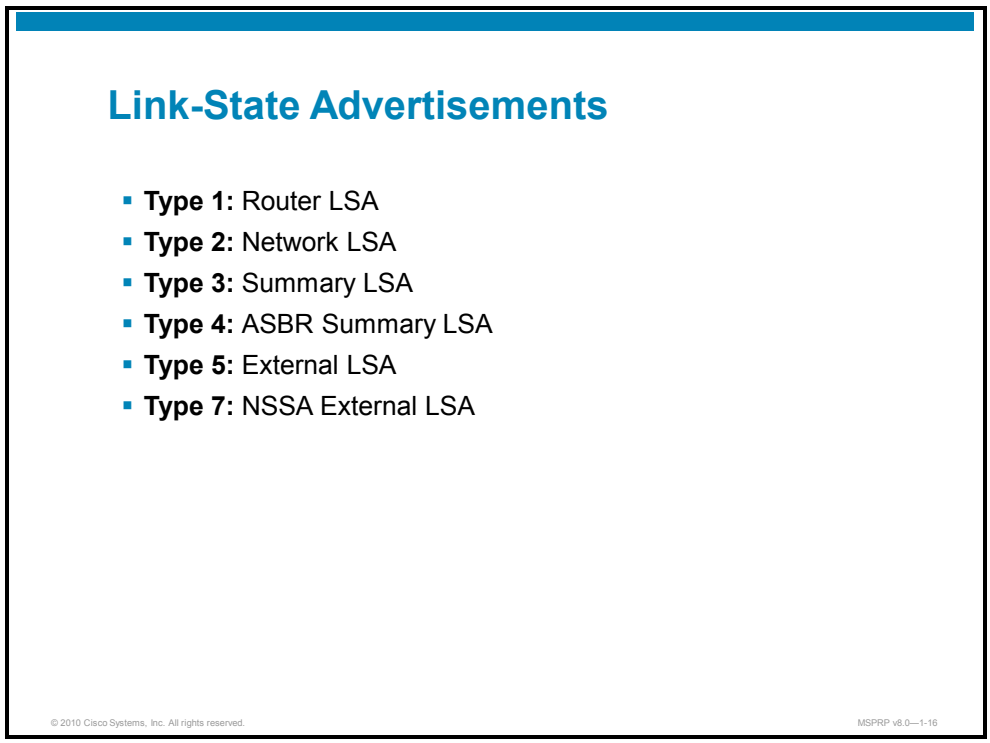

The figure shows only the most commonly used OSPF link-state advertisements (LSAs). There are other LSAs that are not used or are related to MPLS (that is, opaque LSAs), which are not discussed in this course.

Link-State Advertisements (Cont.)

Type 1: Router LSA

- Represents a node (router) in the topology
- Includes:
  - Links to other nodes (routers)
  - Costs of links
  - Leaf networks attached to the node
- Is only propagated within an area
- Becomes a type 3 Summary LSA when crossing an area border

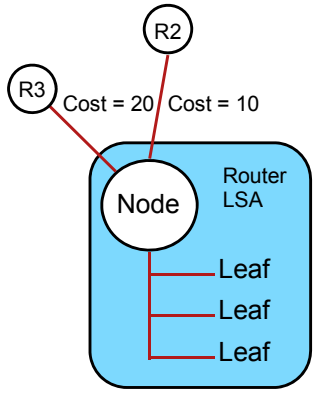

An OSPF topology consists of nodes and links between them. The Dijkstra algorithm is then used to convert the topology into a shortest path tree that is based on the cost of individual links.

OSPF type 1 LSAs or router LSAs are used to propagate the information about nodes and links (including costs). A router LSA also contains leaf networks that are attached to the node.

Link-State Advertisements (Cont.)

Type 2: Network LSA

- Represents a shared network in the topology (router)
- Designated router that acts as virtual node interconnecting all other nodes
- Only propagated within an area
- Becomes a type 3 summary LSA when crossing an area border

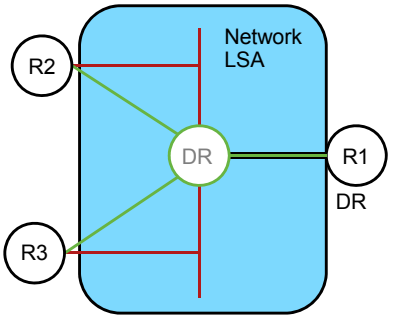

Type 2 LSAs, or network LSAs, are used to represent shared media for which the DR becomes a virtual node in the topology. The virtual node is connected to other routers using “virtual” links. Connecting the node in this way converts a shared topology into a set of point-to-point links between nodes, thus allowing the SPF algorithm to be used.

Link-State Advertisements (Cont.)

Type 3: Summary LSA

- Area Border Routers (ABRs):
  - Convert type 1 and type 2 LSAs into type 3 LSAs
  - Forward them into other areas
- Hierarchical OSPF implementation scales better:
  - A non-backbone router only sees the topology of its area
  - ABRs see topologies of all attached areas
  - Other backbone routers only see the topology of Area 0

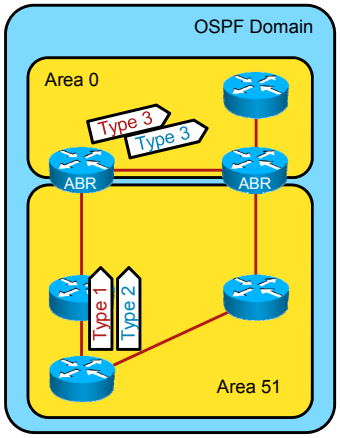

Type 3 summary LSAs are used to carry routing information from one area into another by converting type 1 and type 2 LSAs.

This conversion improves the scalability of OSPF; individual areas do not see the topologies of other areas, thus making SPF calculations more efficient.

ABRs are responsible for creating summary LSAs and injecting them into Area 0 and other areas.

Link-State Advertisements (Cont.)

Other LSAs

- LSA Type 4: ASBR Summary LSA
  - Provides reachability for ASBRs
  - Optimizes routing towards external destinations (type 5 E1)
- LSA Type 5: External LSA
  - Carries external (redistributed) routing information
  - Has two subtypes—E1 and E2
- LSA Type 7: External LSA
  - Is similar to type 5
  - Is generated in NSSA
  - Is converted to type 5 on ABRs

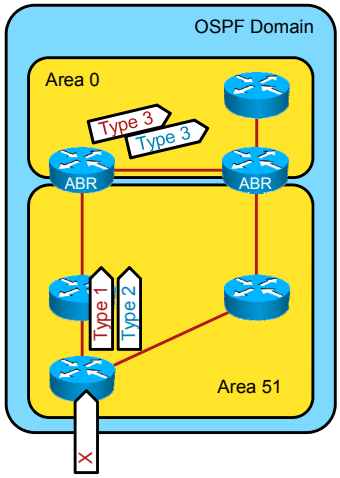

Type 4 LSAs, or Autonomous System Boundary Router (ASBR) summary LSAs, preserve the cost to an ASBR. This behavior allows distant routers in other areas to choose the appropriate external path when type 5 external LSAs are used.

Type 5 external LSAs are generated when redistribution from other protocols is used. There are two subtypes, E1 and E2. E2 is the default for redistributed routes. The main difference between the two subtypes is that the E2 subtype only carries the external cost (cost determined at the time of redistribution), while the E1 subtype also includes the internal OSPF cost.

Type 7 external LSAs resemble type 5 LSAs, except that they are generated in an NSSA. They are converted to type 5 LSAs on ABRs.

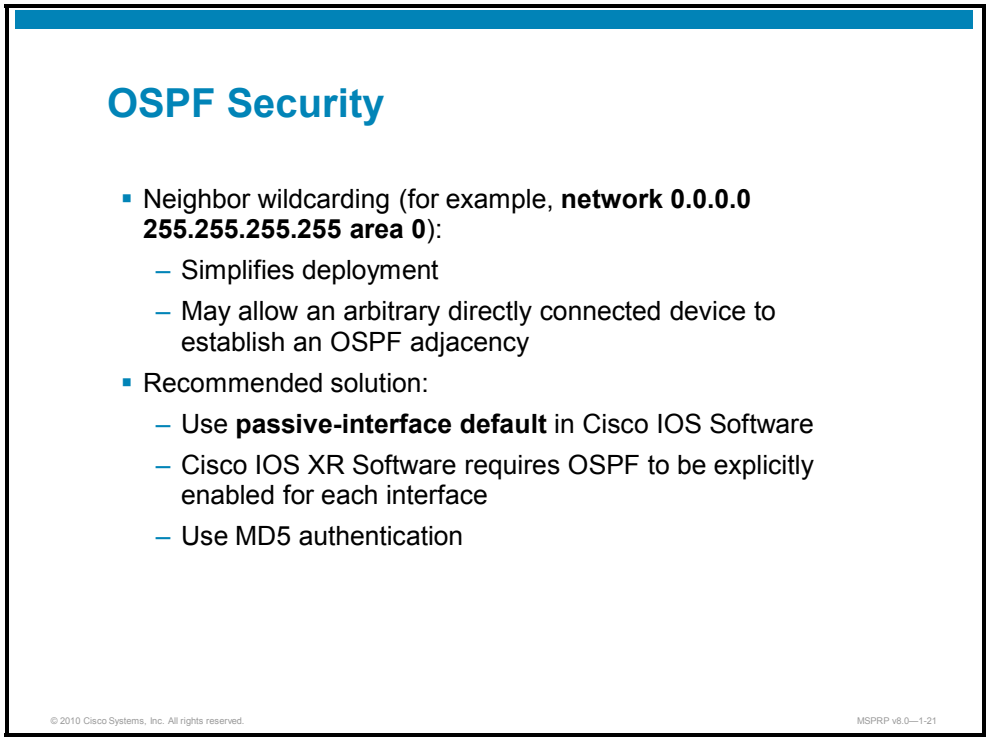

OSPF is prone to bad configurations that allow someone to accidentally or intentionally establish an adjacency and poison the routing table. The most common mistake is to use the wildcard configuration that enables the sending of hello packets on all interfaces:

```
network 0.0.0.0 255.255.255.255 area 0
```

For security reasons, it is recommended that you only configure networks for interfaces that are supposed to be used for OSPF. Additionally, you can use the **passive-interface default** command to disable the sending of hello packets on all interfaces and specifically enable interfaces on which you wish to enable OSPF.

Cisco IOS XR Software is designed to improve the default security posture of OSPF—interfaces must be specifically enabled for OSPF.

Additionally, you should authenticate OSPF sessions using Message Digest 5 (MD5) hashing. OSPF can use the MD5 one-way hash function in combination with a password to authenticate an LSA (MD5 is better that cleartext authentication mechanisms, because it does not put passwords into packets).

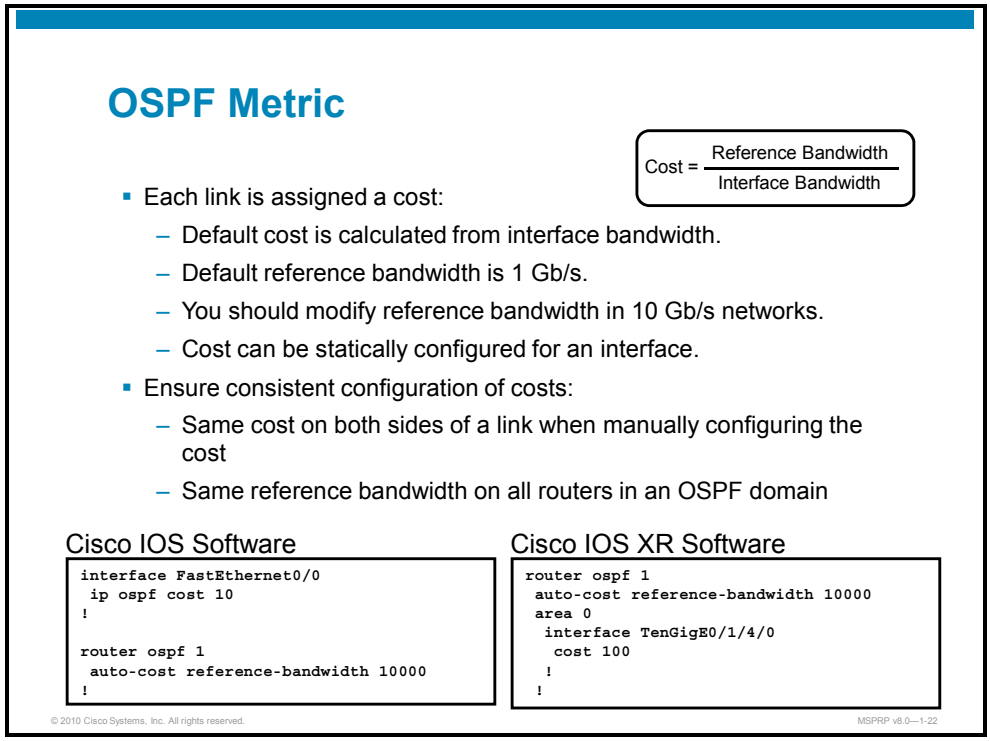

Each link is assigned a cost that is propagated in type 1 LSAs. Default cost is calculated from interface bandwidth and the reference bandwidth, which defaults to 1 Gb/s (1000 Mb/s). The figure illustrates the formula that is used to calculate costs for individual links that are based on default or configured link bandwidth (not the actual link speed).

The default reference bandwidth is only useful in networks where there are no links faster than 1 Gb/s. In faster environments, different links may be given the same cost. For example, 10-Gb/s and 1-Gb/s links would both be assigned cost 1.

Alternatively, cost can be statically configured on a per-link basis. Make sure that you configure costs consistently (the same cost on both routers).

Also make sure that you configure the same reference bandwidth domain-wide.

Route Selection Process

- Order of preference:
  1. Intra-area route
  2. Interarea route
  3. External type 1
  4. External type 2
  5. Lowest cost route
- Note: OSPF cost manipulation may fail when comparing routes of different types.

When OSPF receives multiple LSAs for the same destination, it must compare the paths to determine which path to insert into the forwarding table. The figure shows that the cost of the route is more or less ignored if the two paths are of different types.

Typical OSPF Designs

- Single-area design:
  - All routers in Area 0
  - Simple routing design
  - Mostly point-to-point adjacencies
  - Optimal routing decisions
  - Scalability limited to a few hundred routers in the network
- Multiarea design:
  - Regular areas or NSSA typically used
  - Scales to thousands of routers in the network
  - Mostly point-to-point adjacencies
  - More complex routing design
  - May result in suboptimal routing (such as dual-attached areas)
  - Less practical in MPLS-enabled networks

Most modern service provider networks use MPLS with some of the MPLS-based solutions. When implementing MPLS-based VPNs and traffic engineering, it is important to consider the interaction between MPLS Label Distribution Protocol (LDP) and an IGP. If summarization is used for addresses for which LDP is used to generate label-switched paths (LSPs), it will break those LSPs and consequently break MPLS VPNs.

From a design and implementation perspective, it is preferable to implement OSPF using one area (Area 0). The limitation of this approach is scalability, which is mostly influenced by the number of nodes (routers) in an area.

In large service provider environments, you may be forced to use a hierarchical design. The characteristics and limitations of the hierarchical approach must be considered when designing MPLS solutions.

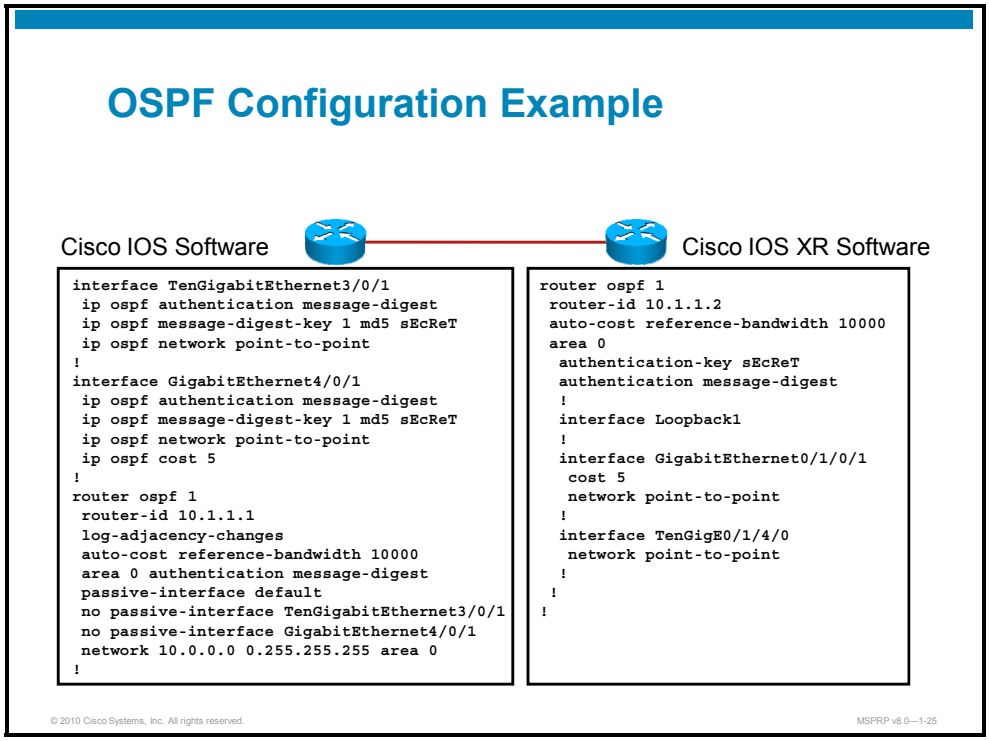

The sample configurations in the figure show a simple implementation of OSPF in Cisco IOS Software and Cisco IOS XR Software. The configurations include the following functionality:

- Manual selection of interfaces on which to run OSPF
- MD5-based authentication of OSPF
- Reference bandwidth to support at least 10-Gb/s links
- Manual configuration of a router ID: OSPF processes use router IDs as unique identifiers that are used when building shortest path trees using the LSP path calculation (Dijkstra algorithm). By default, routers will use the highest loopback IP address or the highest IP address if there is no loopback interface. Routers can also be statically configured with a router ID to ensure consistency (as shown in the example).
- Usage of point-to-point mode on point-to-point Ethernet links

Subsequent lessons and modules will provide additional information to build more advanced configurations to improve the performance of OSPF.

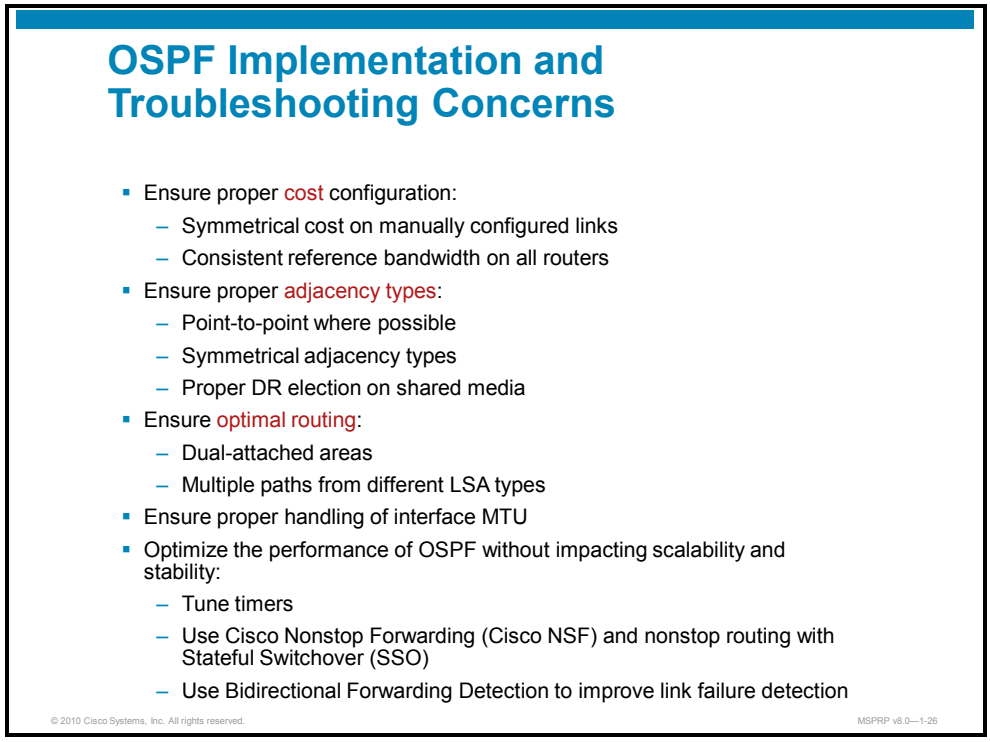

When designing, implementing, monitoring, and troubleshooting OSPF, you should be aware of the following:

- Set link costs properly to ensure that links between pairs of routers are configured with the same cost (symmetric cost configuration).
- For automated cost calculation, configure the same reference bandwidth on all routers in an OSPF domain. It is also important to ensure that all interfaces have a proper default bandwidth or are configured with appropriate bandwidth.
- Optimize OSPF by using the appropriate adjacency mode. For example, Ethernet links will default to broadcast mode—convert all point-to-point Ethernet connections to use OSPF in point-to-point mode. Also make sure that adjacent routers are configured with the same mode.
- In a hierarchical OSPF design with areas that attach to Area 0 using two or more ABRs, make sure that you properly manage summary LSAs to ensure optimal forwarding between areas.
- To prevent suboptimal forwarding, make sure all routes that have multiple paths are of the same type (that is, the same type of LSA).
- Make sure that adjacent routers are configured with the same maximum transmission unit (MTU) in order for OSPF to operate properly. Alternatively, you may disable MTU checking. Also be aware of the difference in MTU configuration between Cisco IOS Software and Cisco IOS XR Software (in Cisco IOS XR Software, MTU configuration includes the Layer 2 overhead).
- Optimize the performance of OSPF without impacting scalability and stability using various software and hardware mechanisms:
  - Tune OSPF timers.
  - Use Cisco Nonstop Forwarding (Cisco NSF) and nonstop routing with Stateful Switchover (SSO).
  - Use Bidirectional Forwarding Detection to improve link failure detection.

Overview of IS-IS

This topic describes the characteristics of IS-IS in service provider environments.

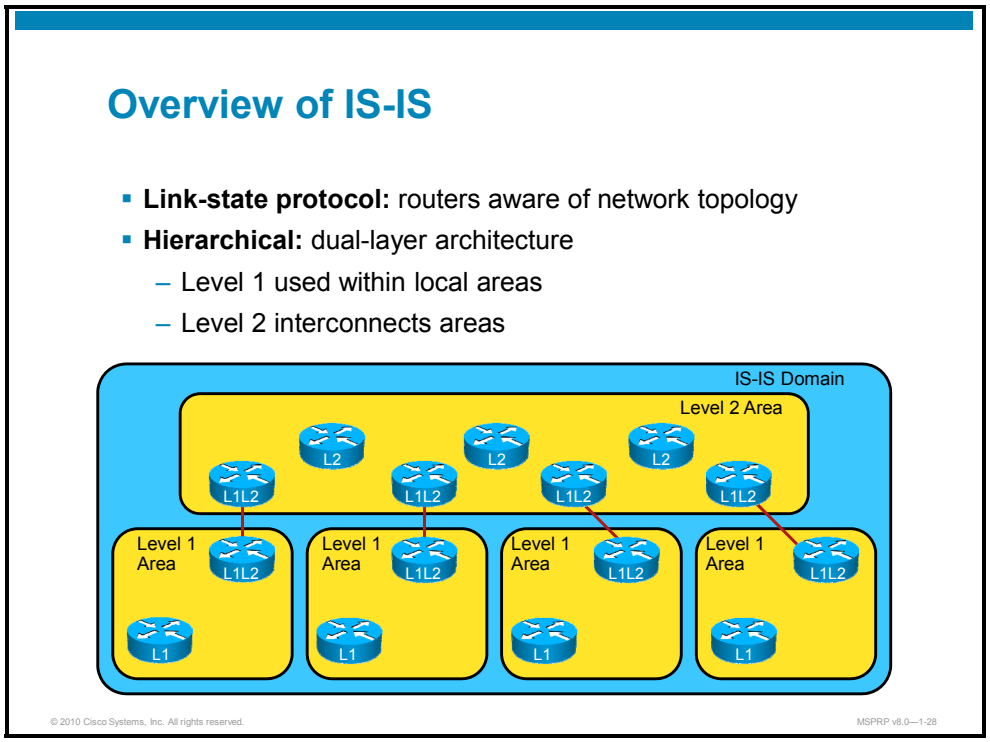

The figure illustrates a hierarchical IS-IS design with the following characteristics:

- Level 2 area with Level-2-only routers in the core
- Level 2 area edge with Level 1-Level 2 routers to connect to other areas using Level 1
- Level 1 areas with Level-1-Level-2 routers connecting areas to Level 2
- Level 1 areas with routers unique to Level 1

This is a very generic representation of what can be implemented using IS-IS. Like OSPF, IS-IS can also be implemented in a simpler fashion with fewer areas or simply one area and one level. Unlike OSPF, all areas do not have to connect to a common backbone area.

IS-IS Characteristics

- Each router has topology information for its area.
- IS-IS is part of OSI and was originally used with CLNS only.
- IS-IS still uses CLNS to maintain adjacencies and build an SPF tree.
- Integrated IS-IS can also carry IP routing information in its updates, such as its type, length, and values (TLVs).
- A wide-style metric should be used for a large, high-speed service provider network (24-bit link metric, 32-bit path metric).
- Link cost defaults to 10.
- Each router is identified using a unique NSAP address.

Like OSPF, IS-IS is a link-state protocol using the Dijkstra algorithm, in which each router has topology information for its area. IS-IS is part of the Open Systems Interconnection (OSI) standard protocol suite and was originally used with Connectionless Network Service (CLNS).

Each router is identified using a unique network service access point (NSAP) address, which is part of the CLNS protocol. IS-IS still uses CLNS to maintain adjacencies and build SPF trees. However, the integrated version of IS-IS can be used for other protocols, such as IP, and can also have extensions for MPLS TE.

A wide-style metric should be used for a large, high-speed service provider network (24-bit link metric, 32-bit path metric). Link cost defaults to 10, but can be modified to reflect the desired cost. The narrow-style metric can only accommodate 64 metric values. This number of values is typically insufficient in modern networks and may not even be compatible with IS-IS extensions such as those for MPLS TE.

Router and Link Types

- Router types:
  - Level 1 routers only peer with other Level 1 routers
  - Level 2 routers only peer with other Level 2 routers
  - Level 1 and level 2 routers can peer with any router
- Link types:
  - Level 1: only for Level 1 adjacencies within the same area
  - Level 2: only for Level 2 adjacencies
  - Level 1 and level 2: for level 1 adjacencies within the same area and level 2 adjacencies

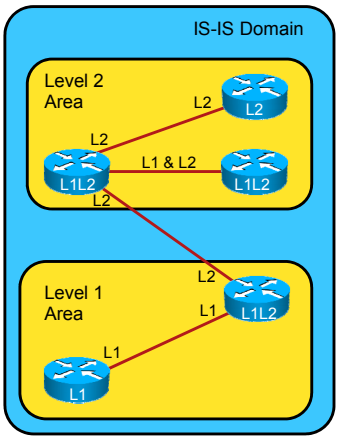

The interaction between routers in an IS-IS domain depends on the following characteristics:

- Router type:
  - Level 1 router can only maintain Level 1 adjacencies
  - Level 2 router can only maintain Level 2 adjacencies
  - Level 1 and Level 2 routers can maintain Level 1 and Level 2 adjacencies
- Link type:
  - Level 1 links only support Level 1 adjacencies
  - Level 2 links only support Level 2 adjacencies
  - Level 1 and Level 2 links support Level 1 and Level 2 adjacencies (concurrently)
- Area:
  - Level 1 and Level 2 adjacencies can be formed between routers in the same area
  - Only Level 2 adjacencies can be formed between routers in different areas

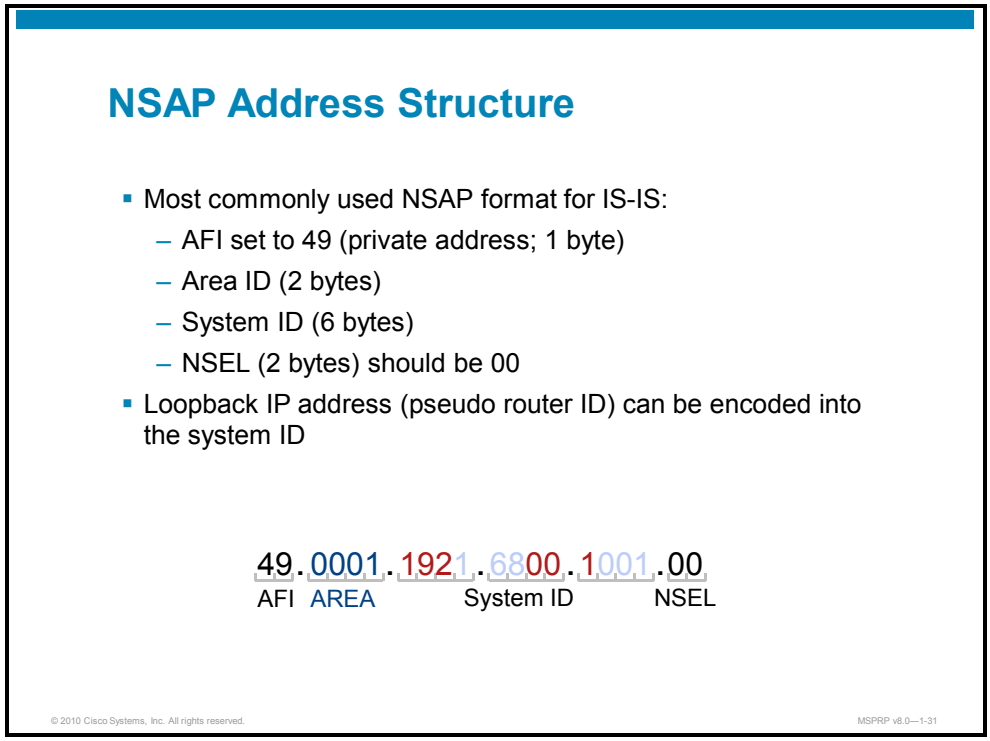

The figure illustrates the most common form of the NSAP address used with IS-IS:

- The authority and format identifier (AFI) is typically set to 49, representing private address space. The field requires 2 hexadecimal digits (1 byte).
- If private addresses are used, there is no interdomain ID.
- Area ID is encoded using 4 hexadecimal digits (2 bytes).
- System ID is encoded using 12 hexadecimal digits (6 bytes). This portion can be used to encode an IP-based router ID (for example, a loopback IP address) to help when troubleshooting, so that NSAP addresses can easily be mapped to IP addresses and vice versa.
- The NSAP selector (NSEL) is encoded using two hexadecimal digits (one byte) and should always be set to 00.

Note In reality, Cisco IOS Software regards AFI and Area as a single entry that can be from 1 to 13 bytes long and represents an Area ID.

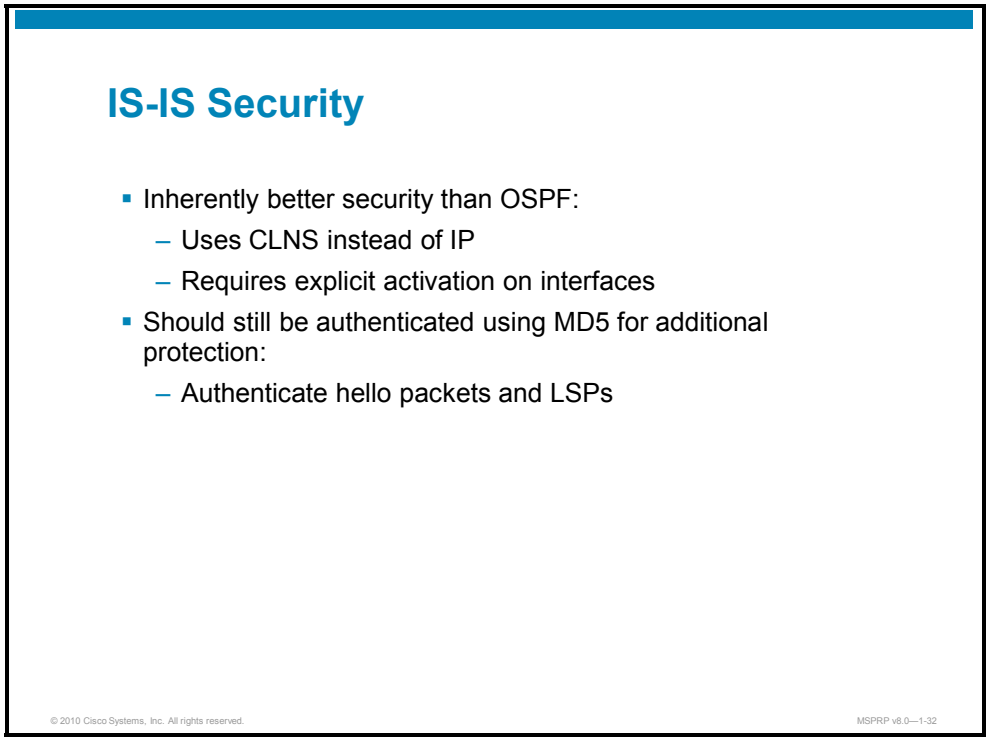

Unlike OSPF, IS-IS will, by design, prevent accidental misconfigurations that could result in reduced security posture:

- Interfaces must be specifically enabled for IS-IS.
- Because CLNS is used, it is slightly less likely that it can be accidentally or even intentionally misused.

Nevertheless, you should still add another layer of security by using MD5-based authentication of both hello packets and link-state protocol data units (PDUs), or of link-state packets (LSPs).

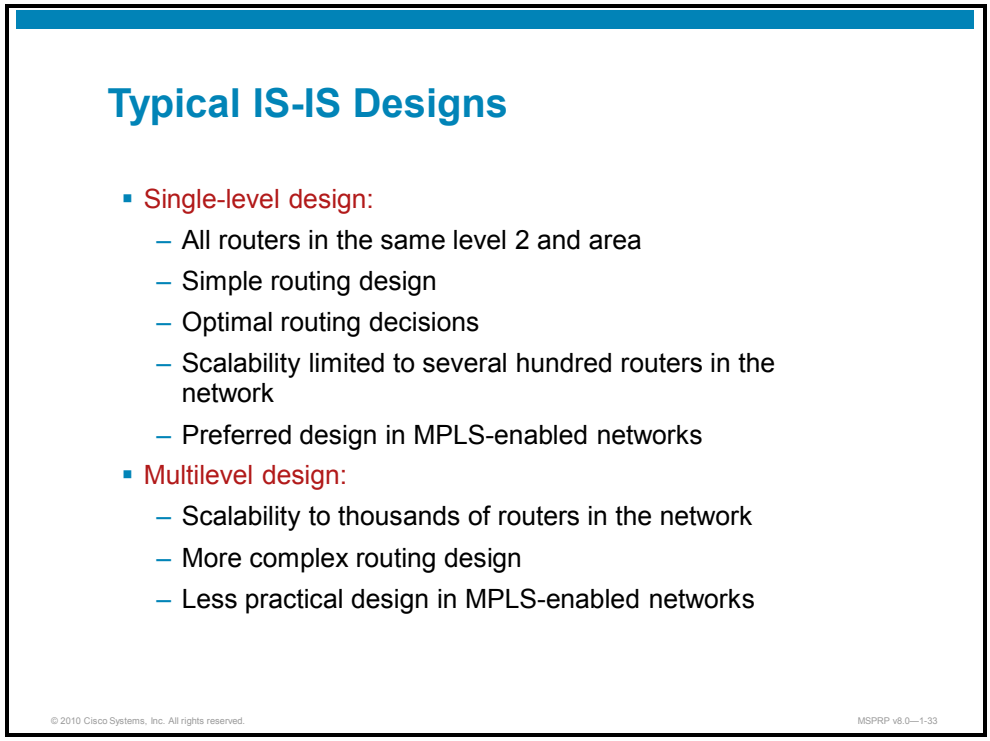

Like OSPF, IS-IS can be implemented in two major ways:

- For the same reasons as with OSPF, it is often desirable to implement IS-IS using one area and level.
- Also as with OSPF, it may be necessary to split IS-IS into multiple levels in order to improve scalability.

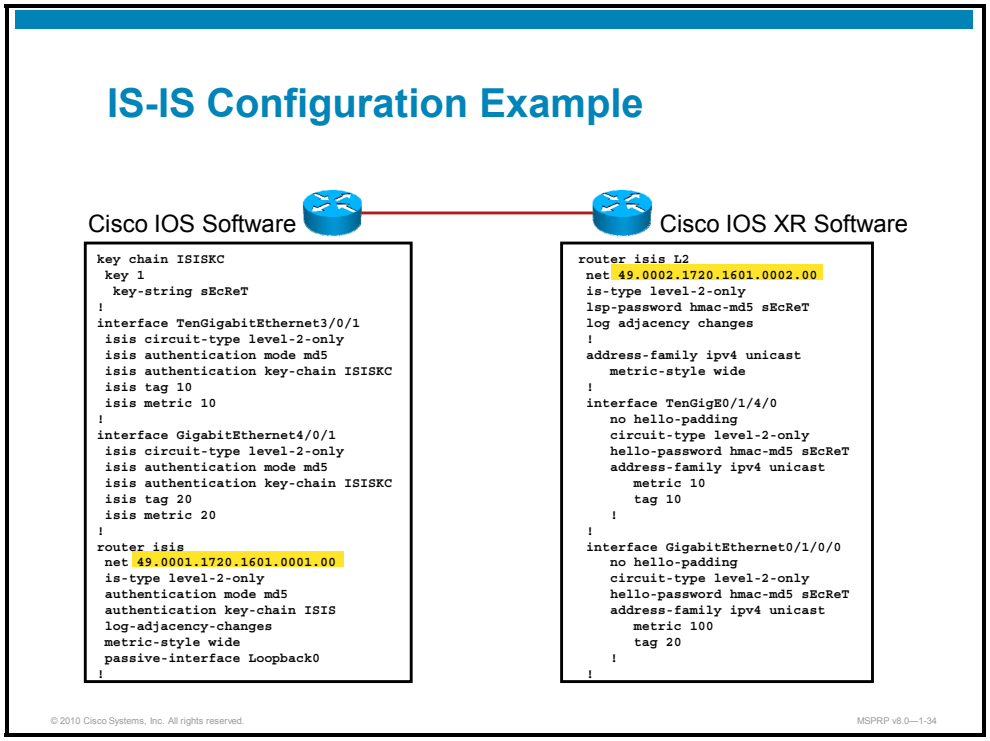

The sample configurations in the figure illustrate a simple implementation of IS-IS in Cisco IOS Software and Cisco IOS XR Software. The configurations include the following functionality:

- Manual selection of interfaces on which to run IS-IS, for better security (to prevent accidental adjacencies on interfaces where there are not supposed to be any IS-IS peers)
- MD5-based authentication of IS-IS to further secure IS-IS by only establishing IS-IS adjacencies with peers that have been configured with the same password
- Optional tagging of routes to support more complex routing policies; IS-IS routes can carry a tag that can be set or matched by route maps (Cisco IOS Software) or routing policies (Cisco IOS XR Software)
- Encoding of an IP into the NSAP address to simplify troubleshooting
- Running of a single level (Level 2) to ensure complete visibility and optimize path calculation (note that running a single level may reduce scalability)
- Inclusion of loopback addresses using the **passive-interface** command

Subsequent lessons and modules will provide additional information to build more advanced configurations to improve the performance of IS-IS.

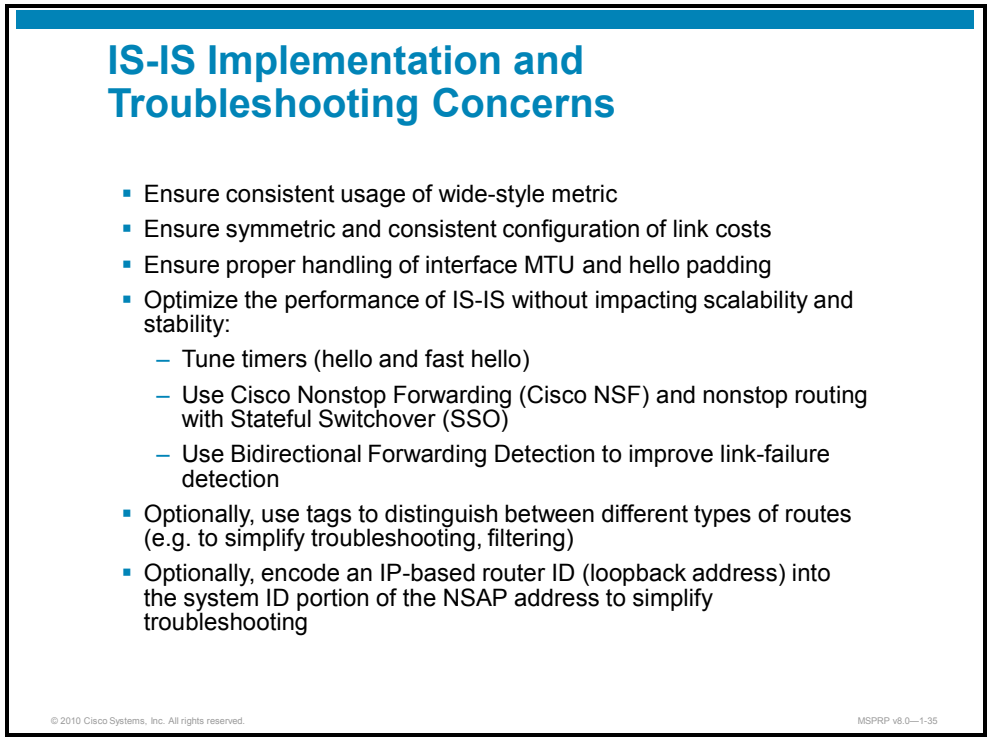

When designing, implementing, monitoring, and troubleshooting IS-IS, you should be aware of the following:

- All routers in an IS-IS domain should be configured to support the wide-style metric.
- Link costs must be properly set to ensure that links between a pair of routers are configured with the same cost (symmetric cost configuration).
- Optimize the performance of IS-IS without impacting scalability and stability using various software and hardware mechanisms:
  - Tune IS-IS timers.
  - Use Cisco NSF and nonstop routing with SSO.
  - Use Bidirectional Forwarding Detection to improve link failure detection.

Overview of BGP

This topic describes the characteristics of BGP in service provider environments.

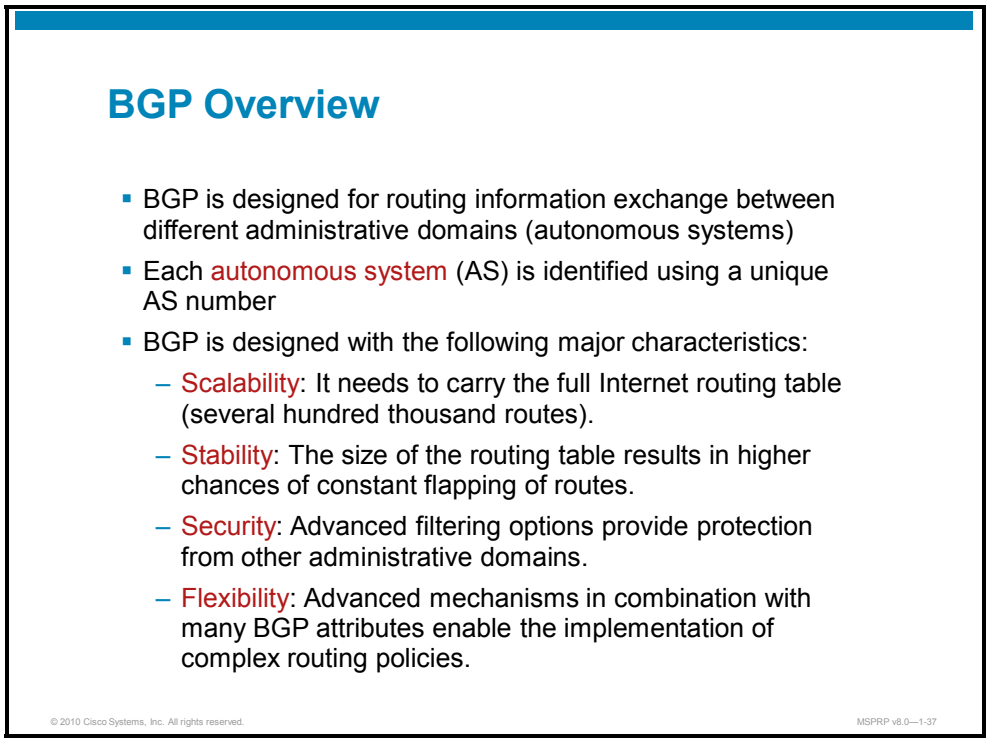

BGP is a distance-vector routing protocol that is designed to meet the following criteria:

- **Scalability:** BGP is intended for distributing Internet routing information that constantly grows.
- **Stability:** It is important for BGP to be able to manage constant flapping of routes in an ever growing Internet, where the likelihood of flapping is also increasing.
- **Security:** Because it is used between administrative domains and in a public environment, BGP must also include powerful security mechanisms. These mechanisms protect routers from intrusions from other administrative domains or the Internet in general.
- **Flexibility:** Complex topologies and diverse requirements demand that BGP support advanced mechanisms to implement complex routing policies.

BGP Architecture

There are two types of BGP sessions:

- External BGP (EBGP) sessions exchange routing information.
- Internal BGP (IBGP) sessions exchange routing information between routers within the same AS.

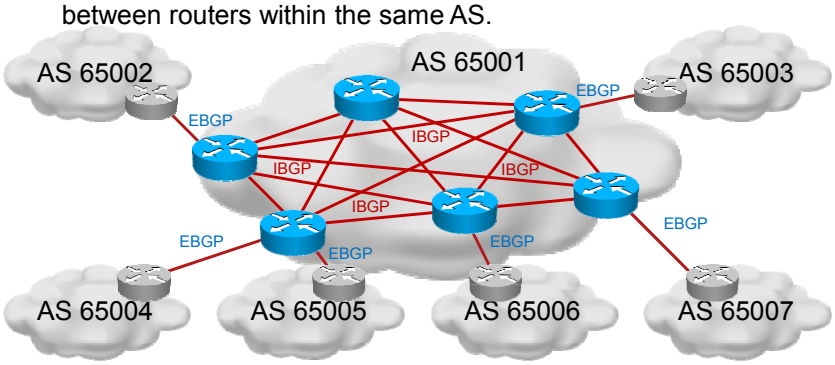

The figure illustrates a general architecture of BGP with the following characteristics:

- Each administrative domain is identified using a unique AS number.
- BGP sessions within an AS are called IBGP sessions and differ from EBGP sessions that are used between different AS.

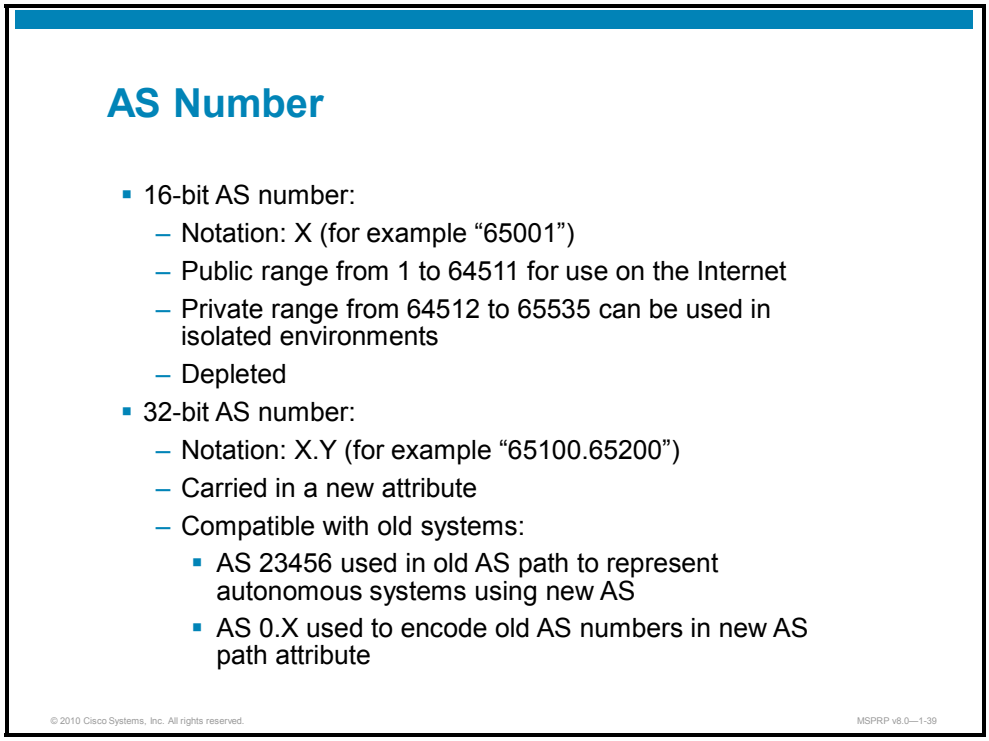

AS numbers come in two forms:

- 16-bit AS numbers have been depleted due to a relatively small number space and the increased demand for companies to be multihomed (increased availability).
- 32-bit AS numbers were introduced to provide a larger number space while maintaining backward compatibility with 16-bit systems to ease the migration. A 32-bit AS can use two notations: A single 32-bit number (X) or two 16-bit numbers joined using a dot (for example, X.Y).

Refer to the following online white paper for a detailed explanation of 32-bit AS numbers and how they interact with older systems:

http://www.cisco.com/en/US/prod/collateral/iosswrel/ps6537/ps6554/ps6599/white_paper_C11_516823.html

BGP Sessions

- BGP uses TCP on port 179 to establish adjacencies
- OPEN messages are used at session setup to negotiate fundamental session parameters and capabilities:
  - AS numbers must match the configuration and determine session type (EBGP versus IBGP)
  - EBGP peers must be reachable through a directly connected link (by default)
  - IBGP sessions are typically established between loopbacks (IGP ensures reachability of loopback addresses)
  - IP addresses must match the configuration
  - Hold time (default is 180 seconds)

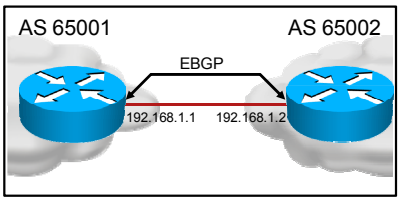

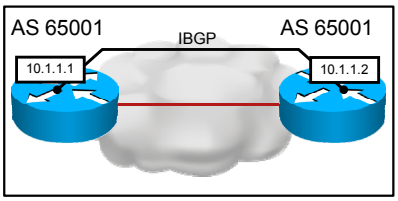

BGP sessions are established over TCP using port number 179. Many session parameters and capabilities are negotiated during session setup by exchanging OPEN messages.

Based on the exchanged AS numbers, both routers will determine if the exchanged numbers match their configurations and select which type the session is (IBGP or EBGP). For EBGP sessions, routers will check to see if the neighbor address is in the routing table as a directly connected address (default requirement). IBGP sessions, on the other hand, can be several hops away, and loopback addresses are typically used to implement IBGP sessions for consistency and stability. The source IP address of a neighbor must also match the configured IP address. Holdtime values are also exchanged and both routers choose the lower value (the keepalive is one third of the holdtime value).

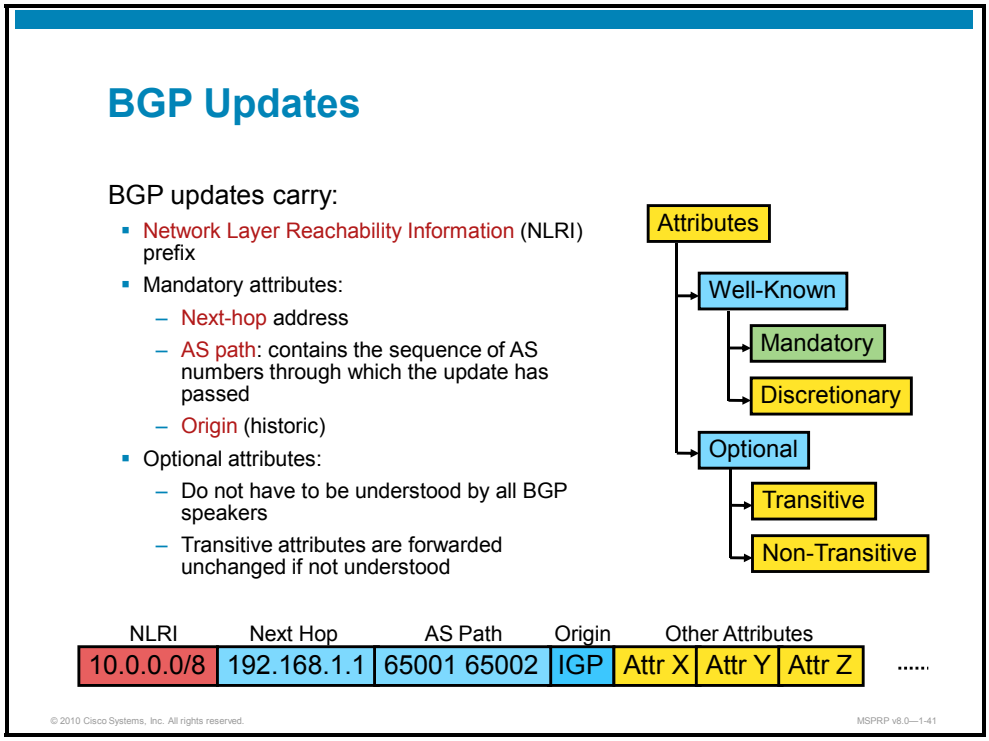

Each BGP update can carry multiple prefixes in the Network Layer Reachability Information (NLRI) portion of the update. Each set of updates must contain at least the three mandatory attributes:

- **Origin code:** This is more or less a cosmetic attribute, which indicates how the BGP routes were originated. For example, the code will be "igp" if a route was natively originated by BGP using network commands, or "unknown" if a route was redistributed into BGP.
- **Next-hop address:** This is important because a BGP entry, when inserted into a forwarding table, will take the prefix and point it to the next-hop address.
- **AS path:** This attribute is primarily used to prevent routing loops. Each border router will check the AS path attribute for the occurrence of its own AS number in the AS path. If a router AS number appears in the AS path, it is dropped. The AS path is also used when selecting the best path for a prefix with two or more paths. The path with the shortest AS path attribute is regarded to be the best.

BGP will typically also carry a number of other attributes. The figure illustrates the types of attributes that BGP supports:

- Well-known attributes are defined by standards and must be supported by all implementations of BGP. Well-known attributes are divided into two categories:
  - Mandatory attributes
  - Discretionary attributes, which may appear in updates, depending on the implementation of BGP
- Optional attributes can be defined by standards or be proprietary. Interoperability between routers should be available even if one of the routers does not support an optional attribute that the other router wants to use.
  - Transitive attributes instruct routers that do not understand them to pass them on to other BGP speakers without any modification.
  - Nontransitive attributes are deleted by routers that do not understand them.

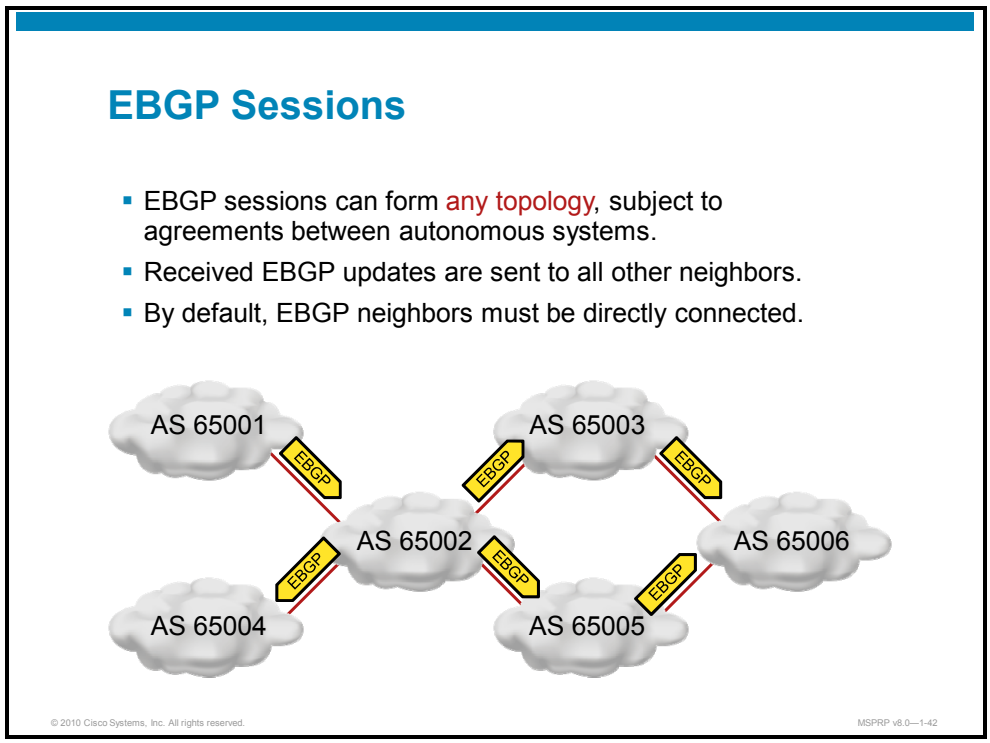

The figure illustrates an arbitrary topology of EBGP sessions between autonomous systems. EBGP has very simple forwarding rules:

- EBGP updates can be sent to all other neighbors (EBGP and IBGP).
- IBGP updates can be sent to EBGP peers.

AS paths are used to prevent updates from looping.

IBGP Sessions

- By default, IBGP sessions require a full mesh between all routers within an autonomous system:
  - A simple split-horizon mechanism prevents internal updates from being sent to other internal neighbors
  - Does not scale in large autonomous systems
- IBGP neighbors can be multiple hops away

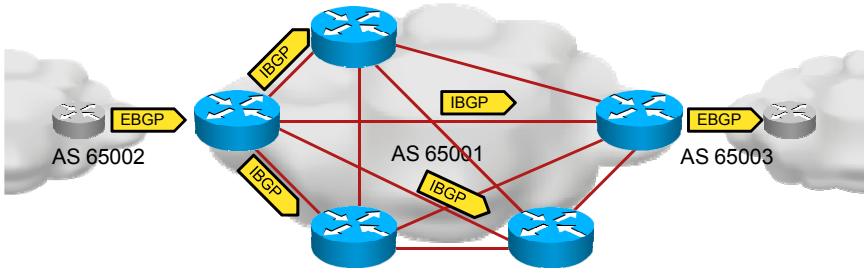

IBGP sessions have different forwarding rules, which include a type of split horizon mechanism:

- IBGP updates can only be sent to EBGP peers.
- EBGP updates can be sent to all neighbors (IBGP and EBGP).

In order to ensure that all routers receive all updates, you must configure a full mesh of IBGP sessions between BGP routers in an AS.

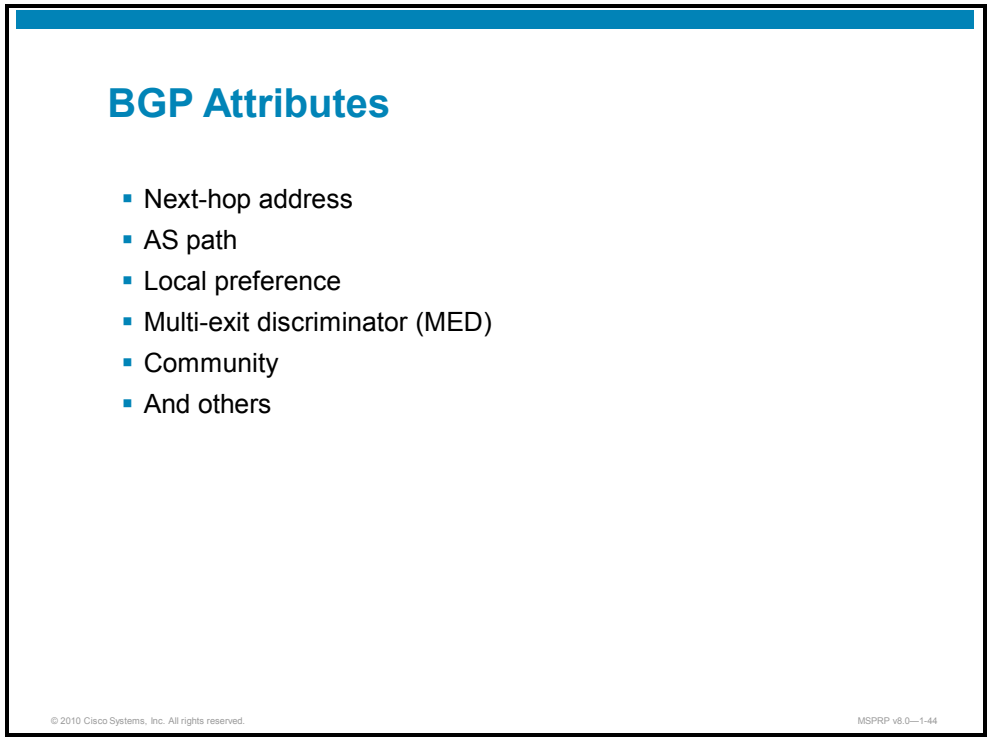

The figure shows some of the most commonly used attributes when implementing routing policies in BGP:

- Next-hop address
- AS path
- Local preference
- Multi-exit discriminator (MED)
- Community

Other common attributes include:

- Extended community
- Originator ID
- Cluster list
- Atomic aggregate
- Aggregator ID
- 4-byte AS path

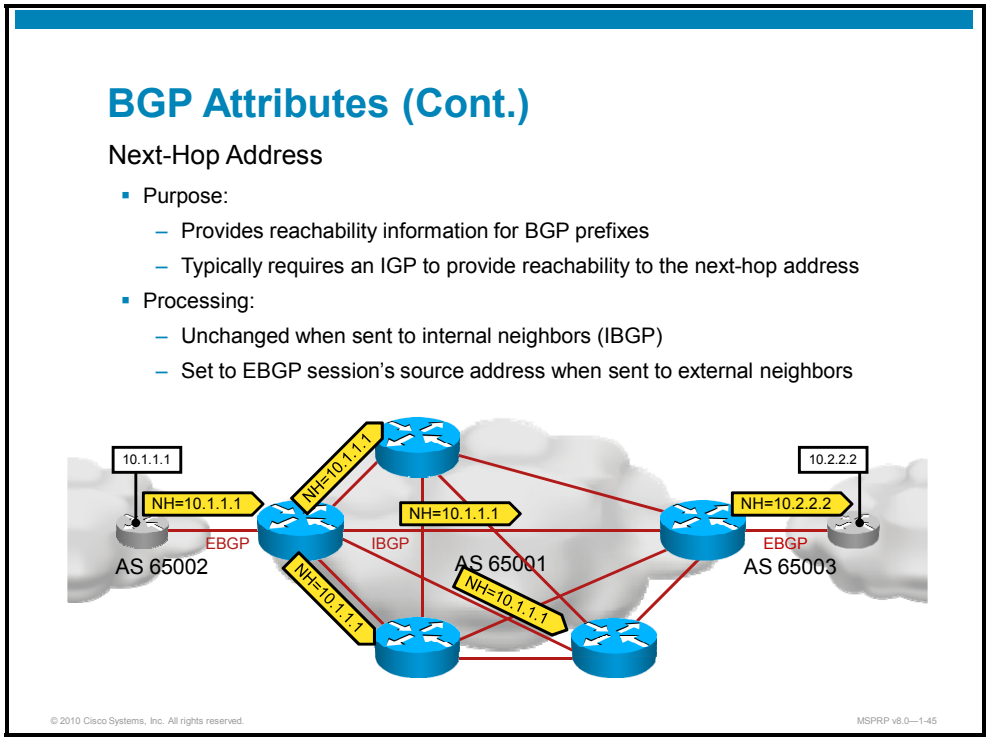

The figure illustrates the propagation of a routing update and how it affects the next-hop attribute:

- Routers sending updates to internal neighbors will not modify the next-hop attribute unless they are configured with the next-hop-self option.
- Routers sending updates to external neighbors will set the next-hop attribute to the source address of the EBGP session.

An IGP is required in an AS to make sure that next-hop addresses are reachable on all routers in the AS.

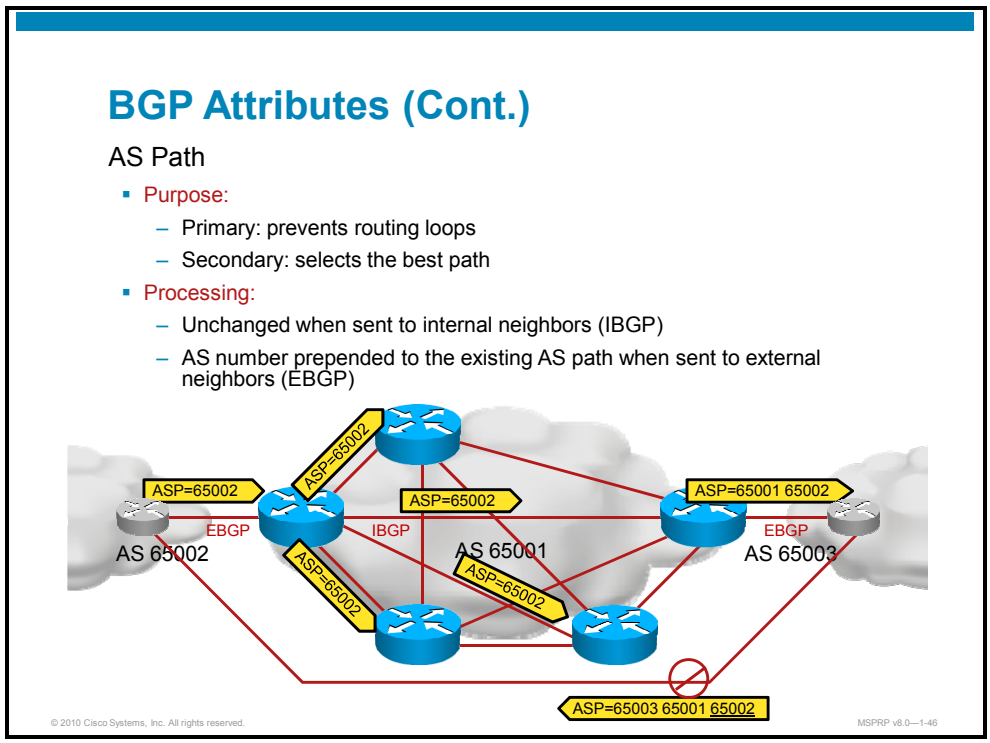

The figure illustrates the processing of updates based on the AS path attribute:

- Routers sending updates to internal neighbors will not modify the AS path attribute.
- Routers sending updates to external neighbors will prepend their own AS number to the existing AS path.

A router that receives an update and finds its own AS number anywhere in the AS path will regard the update as a routing loop and therefore drop the update.

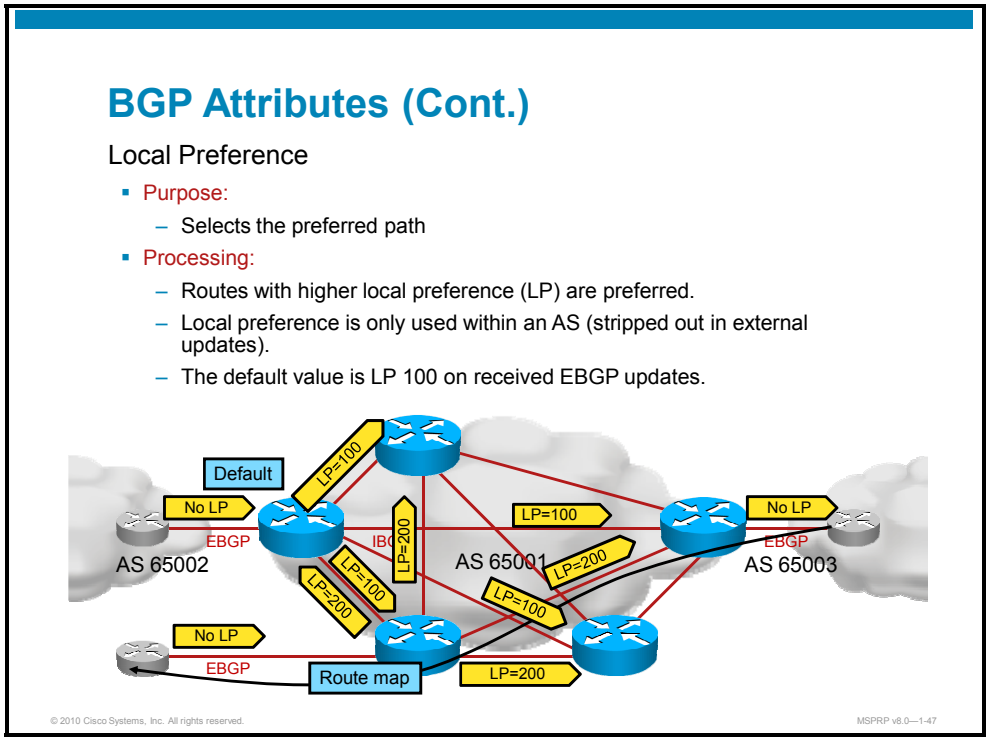

The figure illustrates the processing of updates in relation to the local-preference attribute.

- A router receiving an update that does not contain a local preference will add the attribute and set it to the default value ("100," unless the default has been modified).
- A route map or a route policy can be used to arbitrarily set local preference values to individual routes.
- When comparing routes with multiple paths, a router will prefer routes with higher local preference values.
- When forwarding updates to external neighbors, routers will remove the local preference attribute.

The figure illustrates how the central AS will prefer the path through the lower link to the left AS due to the higher local preference of that path.

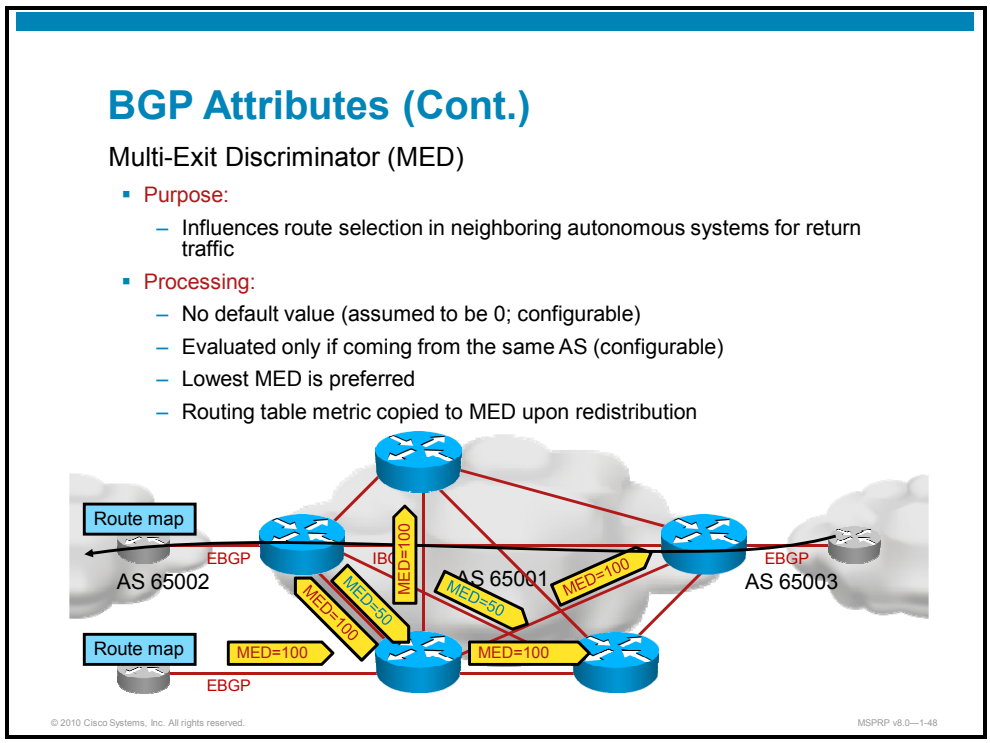

The figure illustrates how an AS can influence the route selection in a neighboring AS by sending EBGP updates with the MED attached. The receiving AS will prefer the path with the lower MED value.

The processing of a MED, by default, is only performed for routes coming from the same AS. The following configuration options modify the default behavior in MED processing:

- **bgp always-compare-med** causes MED values to be compared even if routes come from different autonomous systems.
- **bgp bestpath med missing-as-worst** makes routers RFC-compliant by assuming an infinite value of MED if the MED attribute is not present in an update. By default, Cisco routers will assume a value of "0."
- **bgp deterministic-med** modifies the algorithm in the comparison of more than two paths to make it deterministic and not dependent on the order in which they are compared.

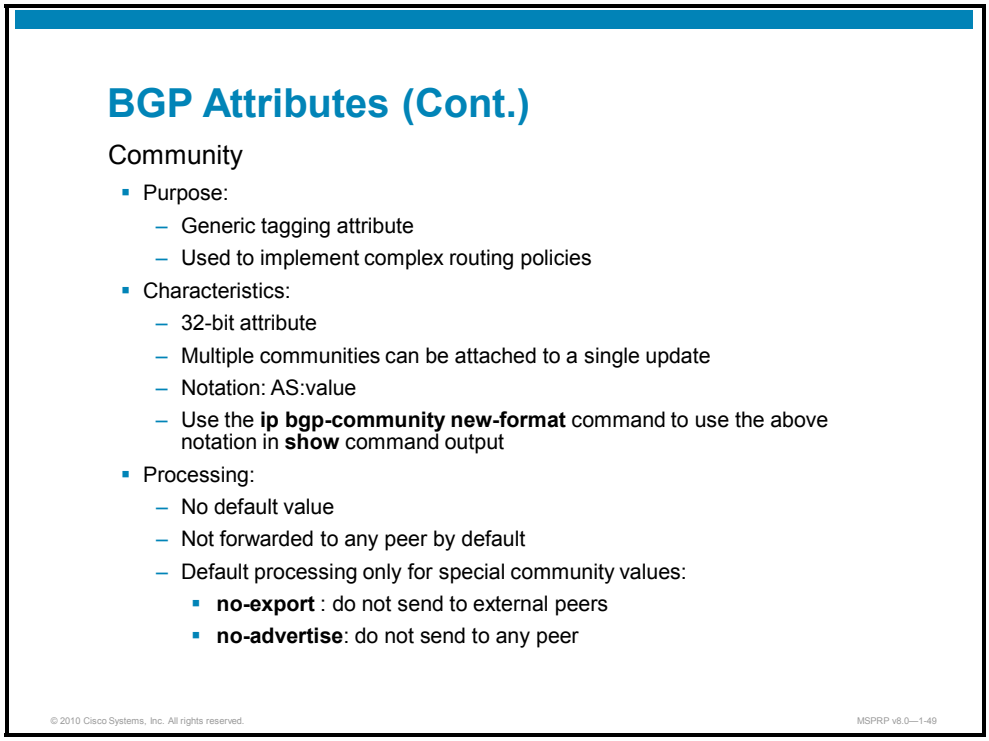

The BGP community attribute by itself has no special meaning unless some standard community values are used. In general, it is a tagging (coloring) tool that allows designers to implement AS-wide routing and filtering policies. These policies depend on the source (for example, ingress edge or customer router) to properly tag a route upon which another router will perform an action.

The standard community values will result in a default action:

- **no-export** allows routes to be propagated throughout an AS, but routers will not forward it to EBGP neighbors.
- **no-advertise** will keep routes only within a router and will not send them to any neighbors.

Each route can be tagged with multiple BGP community values. To ensure uniqueness, it is recommended to put the AS number into the first half of the community attribute and an arbitrary value into the second half.

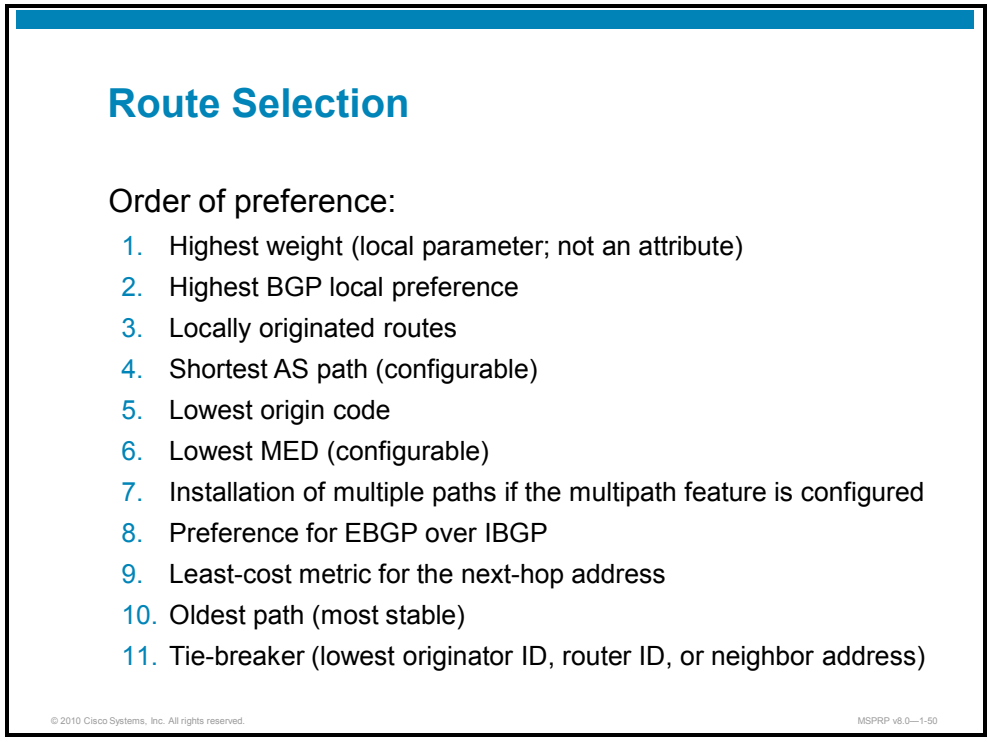

The figure shows the order in which routes for the same prefix are compared.

Refer to the following online document for a more complete description of the route selection algorithm, including the various configuration options that affect it:

http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a0080094431.shtml

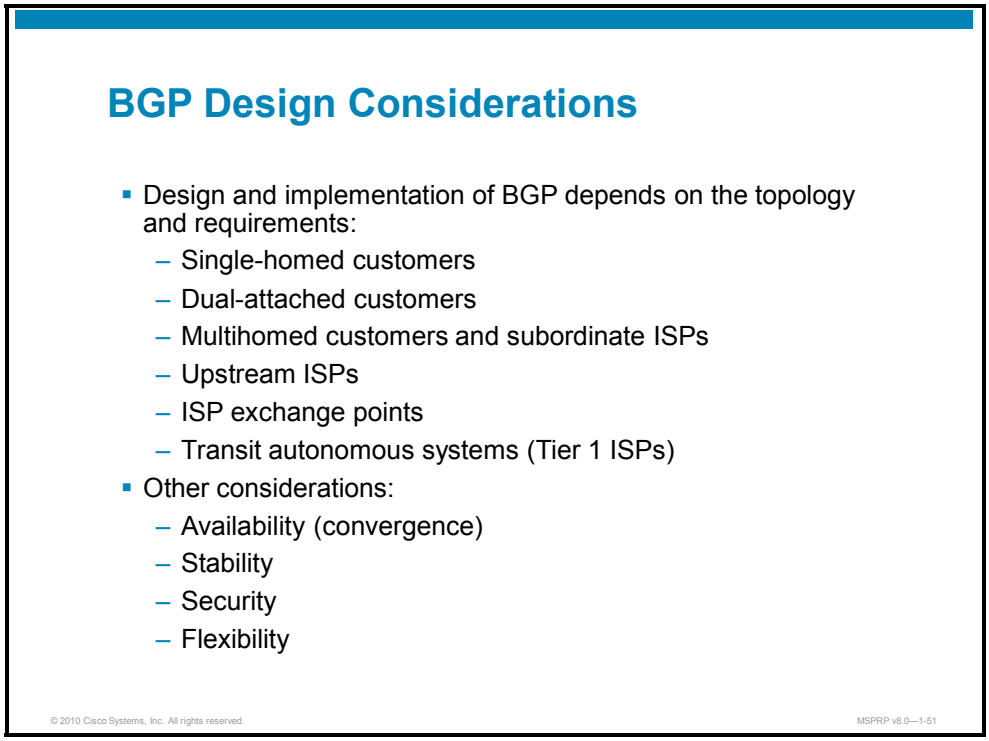

When designing and implementing BGP, you should identify the route or roles that an AS and individual routers play. The first list in the figure shows some design choices for customers and providers (topology).

Additionally, other characteristics need to be identified to properly design and implement high availability, stability, adequate security, and flexibility.

Single-Homed Customers

- Typically do not require BGP:
  - Static route for customer's ISP-assigned address space on edge router
  - Static default route on customer router
- BGP can be used to detect link failures and trigger dial backup:
  - ISP originates only the default route
  - Customer originates its address space
  - Private AS numbers can be assigned to customers by the ISP

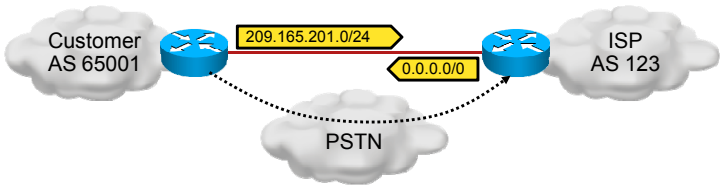

Most single-homed customers (customers that are connected to one ISP using one link) only require static routing. The reason is that there is no alternative path if the primary path fails (such as a router, link, or ISP failure).

Some single-homed customers may deploy a dial-backup solution in which it is beneficial for them to receive notification if their primary path has failed. ISPs can use BGP to send them a default route. If a failure occurs, the BGP session will go down and the customer can initiate a dial backup connection.

Dual-Attached Customers

- Mitigate link and device failures
- Two design options:
  - Primary and backup
  - Load balancing

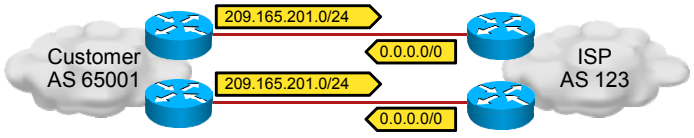

A dual-attached customer (a customer that is connected to the same ISP over two or more links) should require that BGP exchange routing information, enable primary and backup routing or load balancing, and have the ability to detect failed links.

This solution can mitigate router and link failures but it cannot mitigate ISP failures.

Multihomed Customers

- Mitigate link, device, and path failures
- Should connect to independent ISPs
- Two design options:
  - Primary and backup
  - Load balancing

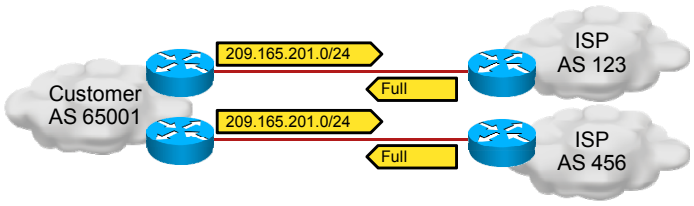

Multihomed customers (that is, customers that are connected to two or more independent ISPs) have the most resilient setup. The setup can mitigate any single failure of routers, links, or ISPs. These customers will often require complete Internet routing information from all ISPs to give them the most flexibility when implementing load balancing.

Multihomed customers have the following requirements:

- Public AS number
- Provider-independent address space

Upstream ISP

- Mitigates link, device and path failures
- Should connect to independent upstream ISPs
- Two design options:
  - Primary and backup
  - Load balancing
- ISP receives full Internet routing table
- ISP forwards:
  - Summaries for owned address space
  - Prefixes from BGP customers using the independent address space

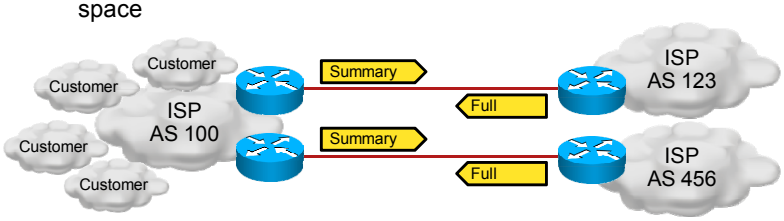

For ISPs, the network is their business. They always have multiple links to other ISPs:

- Upstream ISPs, which will provide complete Internet routing information
- Peering ISPs over a local exchange point, which provide routing information for their customers

ISPs forward summaries for their IP supernets and the individual routes of their BGP-based customers.

Transit ISP

- A transit ISP mitigates link, device, and path failures.
- The routing policy depends on agreements with other ISPs.
- A Tier 1 ISP forwards the full Internet routing table.

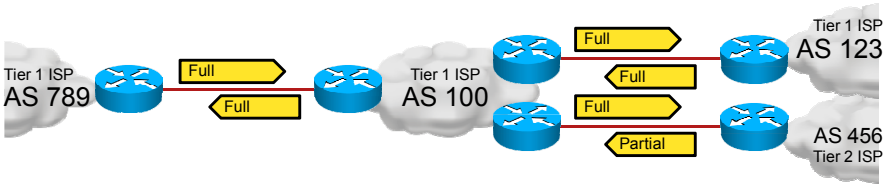

There are many types of transit ISPs, which can be summarized into two major groups:

- Tier 1 ISPs, which peer with other Tier 1 ISPs to form the backbone of the Internet
- Tier 2 and Tier 3 ISPs, which depend on Tier 1 ISPs to reach the rest of the Internet

Relations between these large ISPs heavily depend on the agreements between them and the rates they charge each other. The routing policy must be implemented in accordance with the agreements.

In most cases, Tier 1 ISPs will exchange complete Internet routing tables. In contrast, Tier 2 and Tier 3 ISPs will only forward the routing information for their subordinate ISPs and end customers.

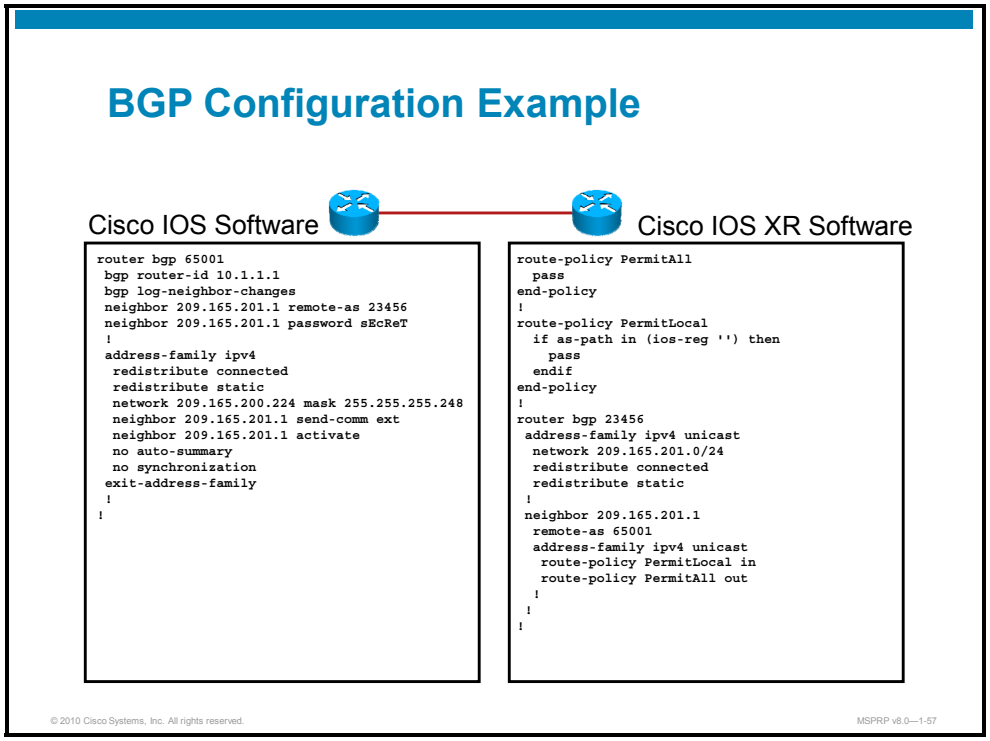

The figure shows a simple implementation of BGP in Cisco IOS Software and Cisco IOS XR Software. The configurations include the following functionality:

- A single EBGP session
- MD5 authentication of the BGP session
- Static configuration of a router ID
- Redistribution of static and connected routes into BGP (not typically recommended; IGP should be used for local routes)
- Enabled forwarding of BGP communities

Additionally, a router using Cisco IOS XR Software requires a routing policy to send and receive updates on an external session.

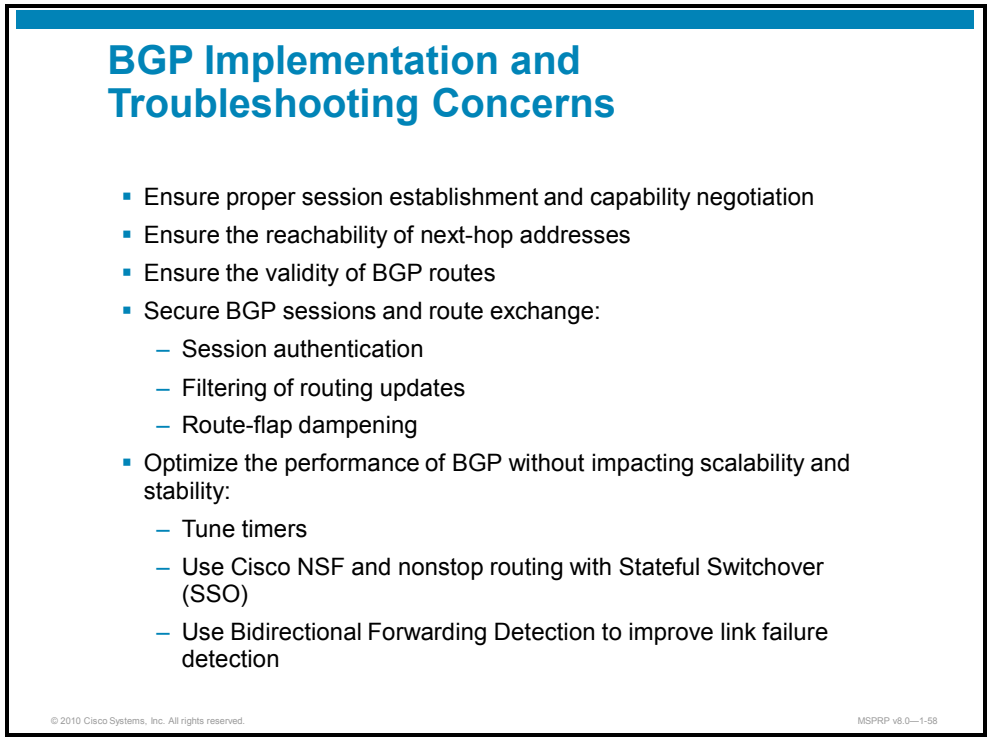

When designing and implementing BGP, it is important to consider the many characteristics and features of BGP:

- Ensure proper session establishment and capability negotiation.
- Ensure the reachability of next-hop addresses.
- Ensure the validity of BGP routes.
- Secure BGP sessions and route exchange:
  - Session authentication using MD5
  - Filtering of routing updates using prefix lists, route maps, or routing policies
  - Route-flap dampening to improve the stability of BGP
- Optimize the performance of BGP without impacting scalability and stability:
  - Tune BGP timers (for example, the scan timer, advertisement interval, keepalives, and fast external failover)
  - Use Cisco NSF and nonstop routing with SSO
  - Use Bidirectional Forwarding Detection to improve link failure detection

Summary

Summary

- ISPs provide IP connectivity within the Internet:
  - To end customers
  - To subordinate ISPs
  - To upstream ISPs
  - To other ISPs through exchange points or direct peerings
- BGP is used to exchange Internet routing information.
- IGPs are used to provide IP connectivity within autonomous systems.
- OSPF and IS-IS are the most commonly used IGPs.
- The following characteristics are desirable:
  - Scalability
  - Performance
  - High availability
  - Security

Lesson Self-Check

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

Q1) In a service provider environment, which two are the main purposes of BGP and IGP? (Choose two.) (Source: Overview of Routing Protocols)

A) BGP is used to provide connectivity within an AS.
B) BGP is used to exchange Internet routing information with other ISPs and those customers that require it.
C) IGP is used to provide reachability for BGP neighbors and BGP next-hop addresses.
D) IGP is used to exchange external routing information.

Q2) In OSPF, which area type typically carries all the routing information? (Source: Overview of OSPF)

A) stubby area
B) totally stubby area
C) regular nonbackbone area
D) backbone Area 0

Q3) Which LSA type carries routing information from one area into another area? (Source: Overview of OSPF)

A) type 3 summary LSA
B) type 5 external LSA
C) type 1 router LSA
D) type 2 network LSA

Q4) Which is the default value of the BGP hold time? (Source: Overview of BGP)

A) 60 seconds
B) 180 seconds
C) 40 seconds
D) 120 seconds

Q5) Which BGP attribute is used to prevent routing loops? (Source: Overview of BGP)

A) origin code
B) local preference
C) AS path
D) next-hop address

Lesson Self-Check Answer Key

Q1) B, C

Q2) D

Q3) A

Q4) B

Q5) C

Lesson 2

Using Routing Protocol Tools

Overview

The lesson focuses on the mechanisms that are available in combination with routing protocols to filter routing information or to implement desired routing policies using Cisco IOS and Cisco IOS XR routers. The lesson describes prefix lists, autonomous system (AS) path access lists, route maps, and the Routing Policy Language (RPL).

Objectives

Upon completing this lesson, you will be able to identify the main characteristics of routing protocols that are used in service provider environments. This ability includes being able to meet these objectives:

- Describe the characteristics and requirements for routing policies in service provider environments
- Describe the characteristics and usage scenarios for prefix lists
- Describe the characteristics and usage scenarios for AS path-based filtering in service provider environments
- Describe the characteristics and usage scenarios for route maps in service provider environments
- Describe the characteristics of RPL

Routing Protocol Tools Overview

Objective: describe the characteristics and requirements for routing policies in service provider environments.

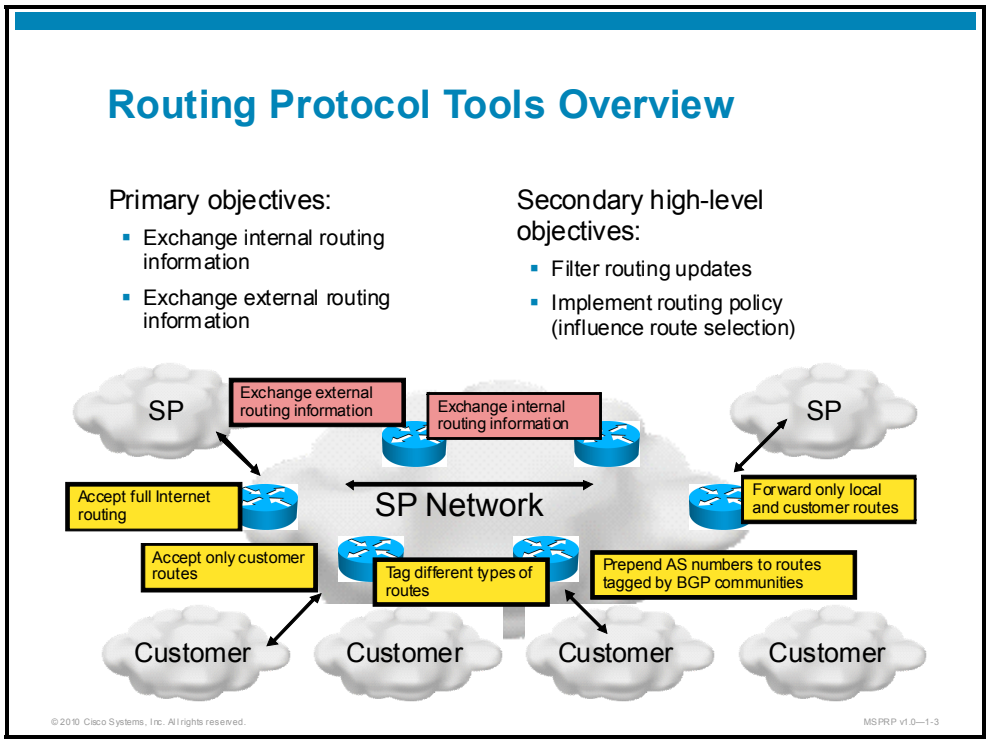

The figure illustrates various actions that are performed on routing updates in a typical service provider environment. The actions can be divided into two main categories:

- Exchanging routing information (the primary objective of routing protocols)
- Implementing a routing policy and filtering of routing information

To exchange routing information, a typical service provider uses two routing protocols:

- An interior gateway protocol (IGP) such as Open Shortest Path First (OSPF) or Intermediate System-to-Intermediate System (IS-IS) to exchange local routing information.
- Border Gateway Protocol (BGP) to exchange external routing information (that is, customer routing information and complete Internet routing information from other service providers).

BGP is always combined with advanced filtering and policy mechanisms for security and performance reasons. This lesson will discuss various mechanisms that can be used for filtering and routing policy implementation.

Typical Filtering Objectives

Filter:

- Incoming updates
- Outgoing updates
- Redistributed routes from other routing protocols

Filter based on:

- Prefix and prefix length (subnet mask)
- Update parameters (specific to a routing protocol)

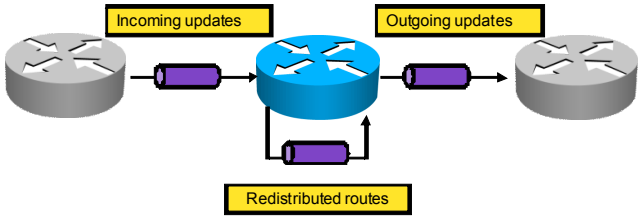

Routing information can be filtered in the following locations:

- On incoming updates as they are received from a neighboring router
- On outgoing updates before they are sent to a neighboring router
- On redistribution from another routing protocol including static and connected route redistribution

Filtering itself can be based on the prefix or based on another parameter available in the routing protocol or source routing protocol (when doing redistribution).

Example: Typical OSPF Filtering Objectives

Filter OSPF based on:

- Prefix and prefix length (subnet mask)
- LSA type (internal, external, NSSA-external)
- Route source

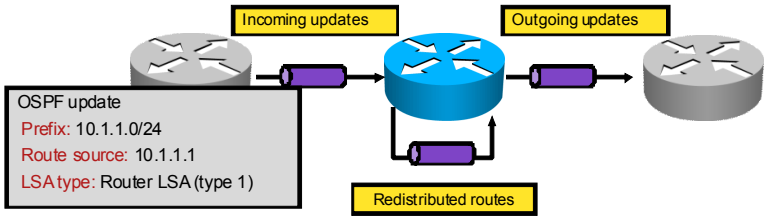

The figure illustrates an OSPF update that carries information that can be used for filtering purposes:

- Prefix and prefix length
- Route source (that is, advertising router's IP address)
- OSPF link-state advertisement (LSA) type

Redistributed routes can be filtered on any router that effectively becomes an Autonomous System Boundary Router (ASBR). In contrast, regular filtering of OSPF updates can only be performed on Area Border Routers (ABRs) for routes forwarded from one area into another.

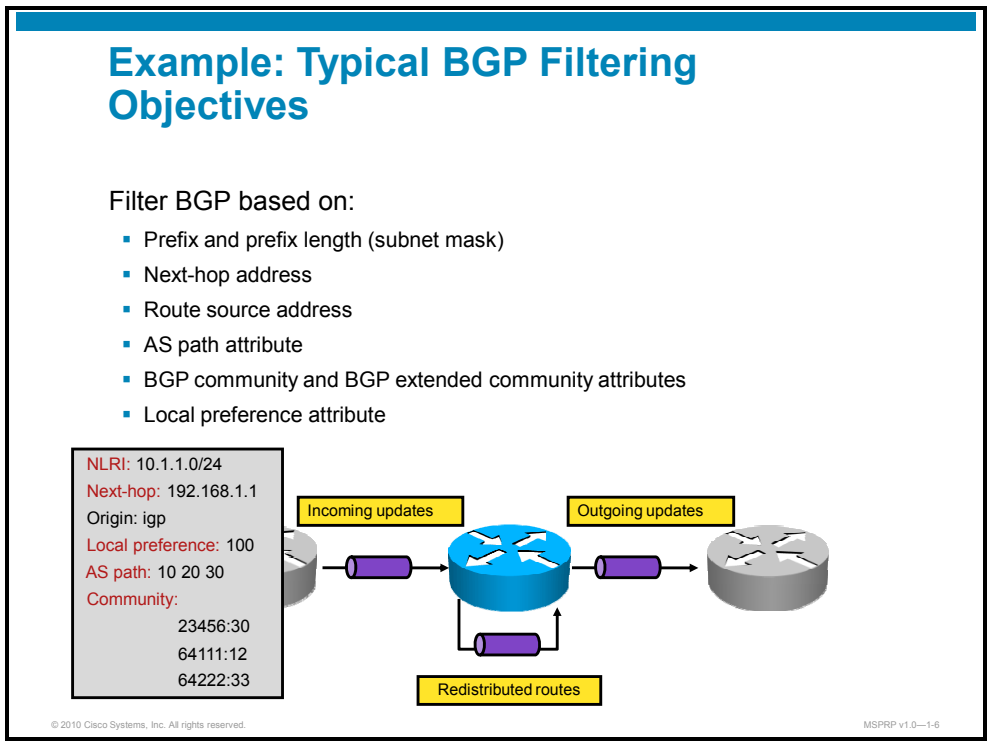

The figure illustrates a BGP update, which has a much richer metric (that is, collection of BGP attributes) that can also be used for filtering purposes.

BGP updates can be filtered based on:

- Prefix and prefix length (subnet mask) found in the BGP Network Layer Reachability Information (NLRI)
- Next-hop address found in the BGP next-hop attribute
- Route source address (that is, the neighbor's IP address)
- AS path attribute
- BGP community and BGP extended community attributes
- Local preference attribute

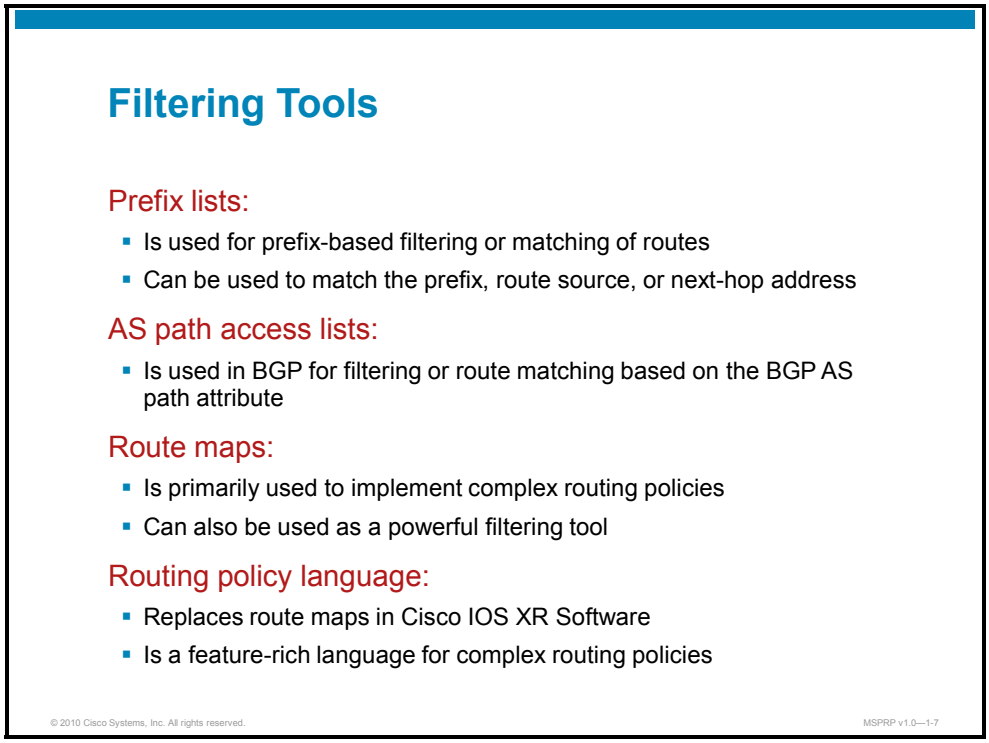

The following tools are most commonly used to implement filtering and routing policies in Cisco IOS and Cisco IOS XR Software:

- **Prefix lists** can be used to implement filtering or matching of routing updates based on IP addresses or IP network information such as prefixes, next-hop addresses, or neighbors addresses. Prefix lists are available in Cisco IOS Software. Prefix lists are also available in Cisco IOS XR Software, with slight differences.
- **AS path access lists** can be used with BGP to implement filtering or matching of routing updates based on the contents of the AS path attribute. A regular expression is used to process the AS path as a string of characters. AS path access lists are only available in Cisco IOS Software. Cisco IOS XR Software matches AS path attributes directly in routing policies.
- **Route maps** are primarily used to implement routing policies that can also modify routing protocol parameters as well as perform filtering. Route maps are only available in Cisco IOS Software.
- **Routing policies** are more powerful and flexible versions of route maps that are available in Cisco IOS XR Software.

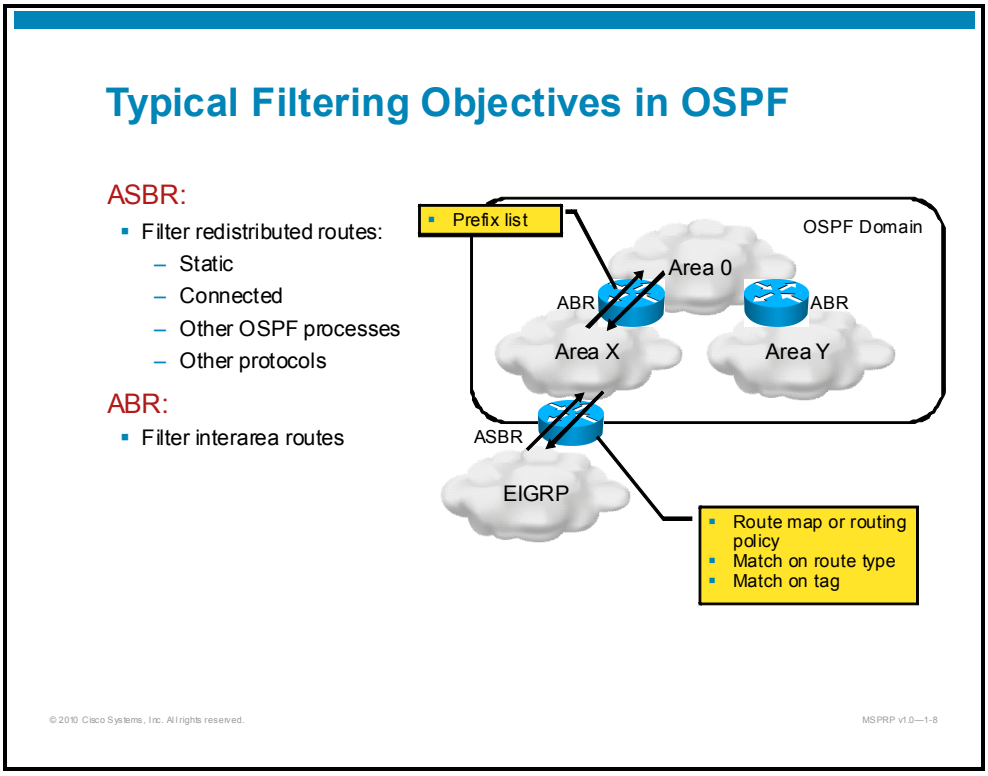

The figure illustrates a sample OSPF domain using multiple OSPF areas and a connection to an external Enhanced Interior Gateway Routing Protocol (EIGRP) AS.

ASBRs can filter redistributed routes using route maps or routing policies from any of these:

- Connected routes
- Static routes
- Other OSPF processes
- IS-IS
- EIGRP
- Routing Information Protocol (RIP)

BGP (not recommended)ABRs exchange routing information between OSPF areas within the same OSPF domain according to OSPF rules. Prefix lists can be used to control the exchange of routing information between areas.

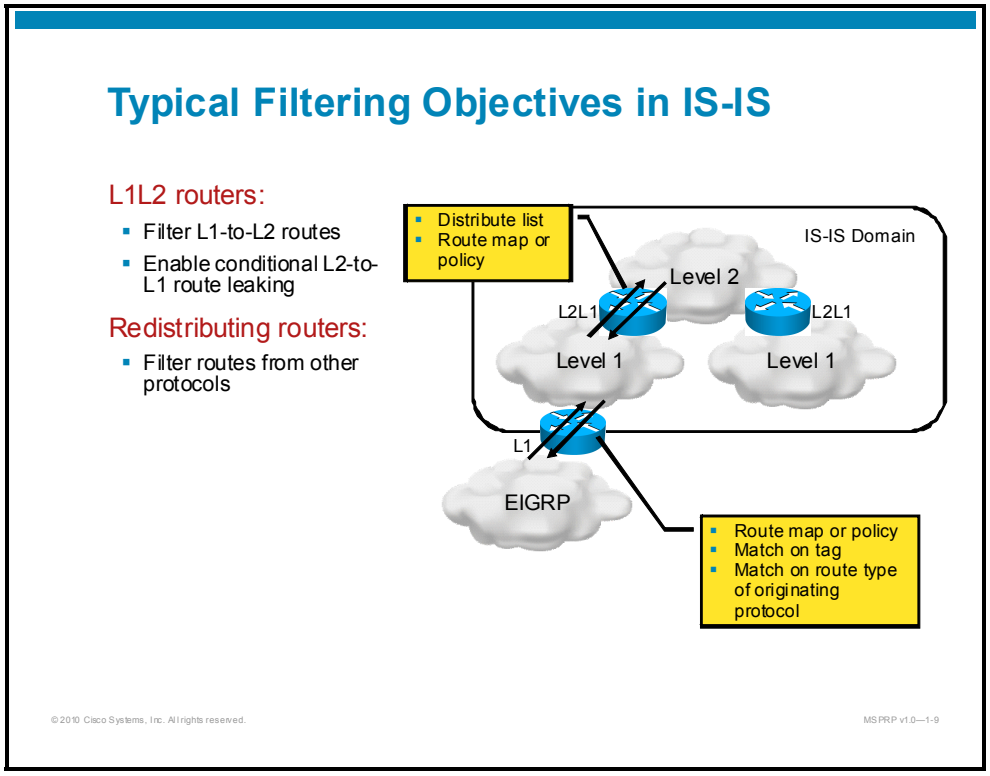

The figure illustrates a sample IS-IS domain using multiple IS-IS levels and a connection to an external EIGRP AS.

L2L1 routers (like ABRs) perform routing exchange for both IS-IS levels. Prefix lists, route maps, or routing policies can be used to filter exchange of routing information between IS-IS levels. Route leaking can also be used to control the distribution of Level 2 routes into Level 1.

Any IS-IS router can perform redistribution from other routing protocols using a route map or routing policy to control the redistribution of routes.

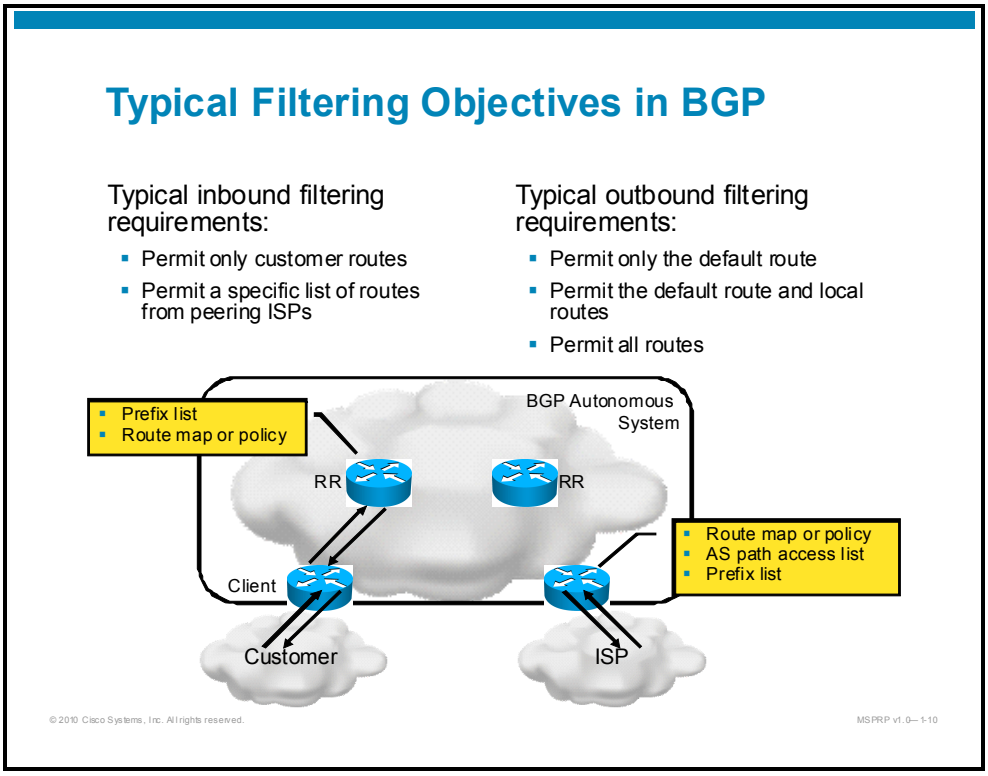

The figure illustrates a sample BGP AS using BGP route reflectors (to reduce the full-mesh Internal Border Gateway Protocol [IBGP] requirements) and edge BGP routers to implement routing for external destinations.

Inbound filtering can depend on the type of neighboring AS:

- Permit only customer routes for end customers
- Permit a specific list of routes from subordinate ISPs or ISPs peering at an exchange point
- Permit the complete Internet routing information from upstream ISPs

Outbound filtering depends on the type of neighboring AS or on the customer requirements:

- Permit only the default route (for example, for single-homed customers that do not require more specific information; most single-homed customers do not even require a routing protocol)
- Permit the default route and local routes (for example, for multihomed customers using a specific ISP as a backup provider but still accessing local destinations directly)
- Permit all routes (for example, for multihomed customers that require complete Internet routing information)

Typical Routing Objectives

- Implementation of complex routing policies using BGP
- Complex routing policies mostly implemented using BGP
  - Outgoing traffic
  - Incoming traffic
- Routing decision influenced
  - Locally
  - Remotely (such as by a customer or downstream ISP)

Routing policies are most commonly implemented for external routing information using BGP. A routing policy can address the outgoing path or the return path. Additionally, you can use BGP to implement a policy locally within your AS or have a neighboring AS influence the route selection in your AS. For example, you can use AS path prepending or signal a policy using BGP communities, which are translated to local preference in your AS.

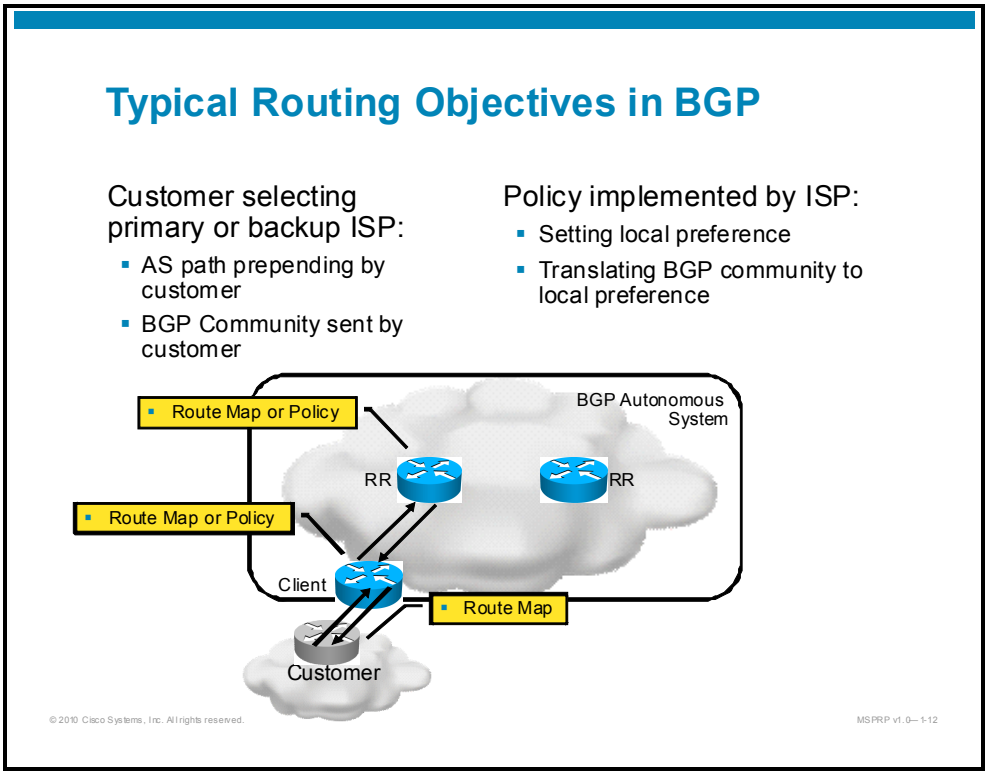

A routing policy is always implemented using route maps in Cisco IOS Software or routing policies in Cisco IOS XR Software. You should implement policies that are consistent across the entire AS (that is, implement policies on edge routers).

The figure lists some commonly implemented policies in service provider environments:

- Customers often use AS path prepending to artificially lengthen the AS path attribute, thus making it less desirable (that is, to signal that this ISP is the backup ISP).
- Customers can alternatively signal their ISP preference by using some BGP communities an ISP offers. The ISP will then translate the BGP communities that it receives from the customers to another BGP attribute. For example, the ISP might use AS path prepending or local preference to influence the outbound traffic to the customers.
- ISPs can use the BGP local preference attribute to influence route selection internally within the ISP AS (for example, select preferred upstream ISPs).

Prefix Lists

Describe the characteristics and usage scenarios for prefix lists.

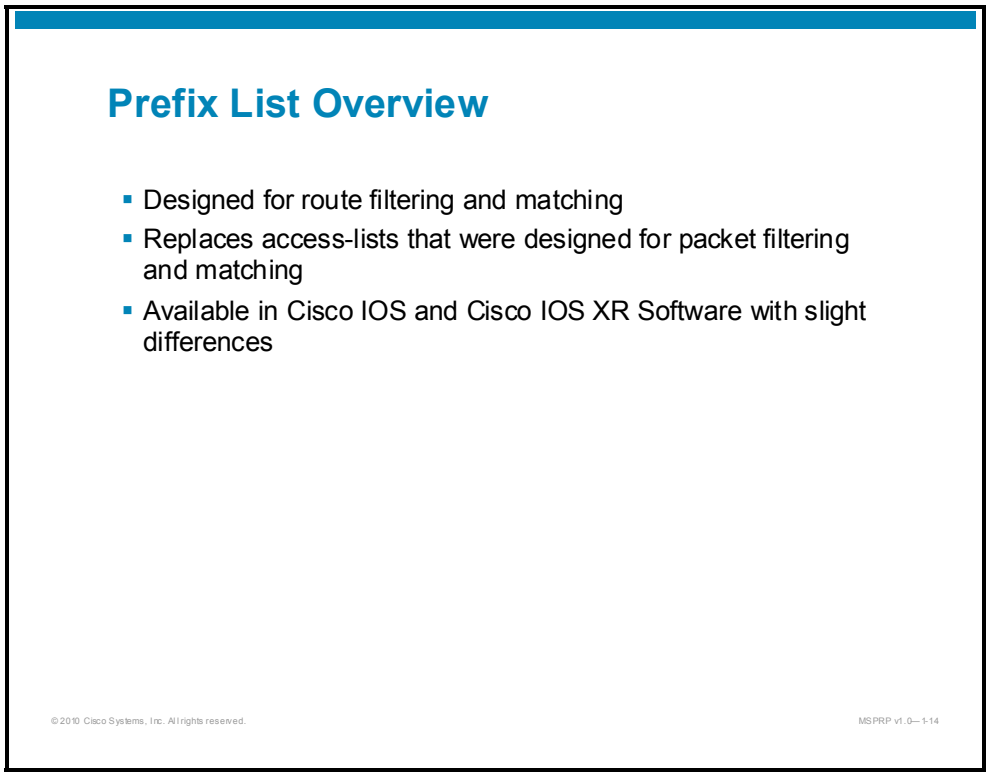

Prefix lists are designed to simplify the filtering of routing updates. They are available in Cisco IOS, Cisco IOS XE (for the ASR router family) and Cisco IOS XR Software (with some slight differences).

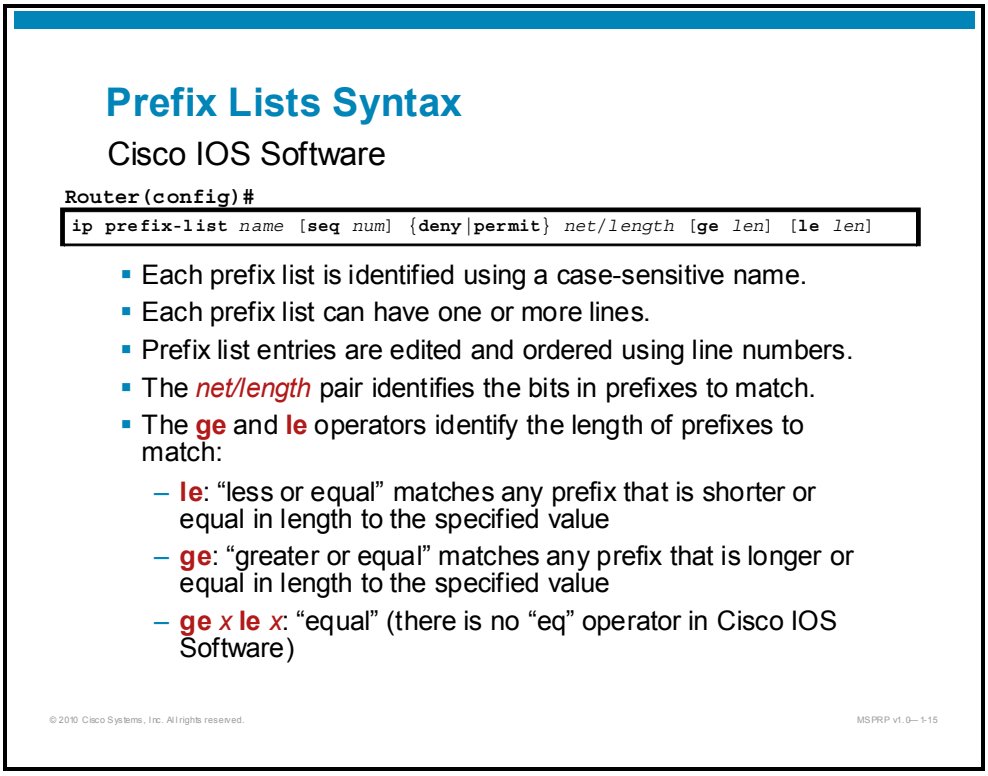

Each prefix list is identified using a case-sensitive name (like all other named objects in Cisco IOS and Cisco IOS XR Software). A prefix list can have multiple lines that are ordered using line numbers.

The network and length pair identifies the bits in prefixes to match. The **ge** and **le** operators identify the length of prefixes to match. A combination of both operators can be used to match a range of prefix lengths or a specific length—**ge** *x* **le** *x* ~ "equal" (there is no "eq" operator in Cisco IOS Software).

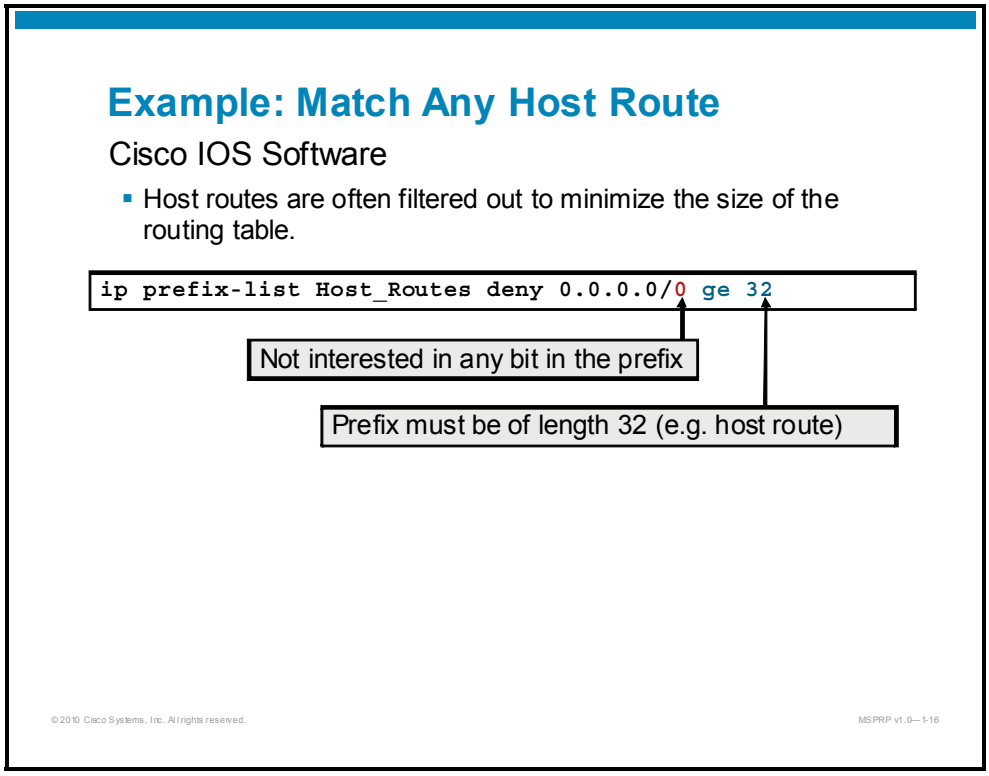

The sample prefix list shows how to match any host route:

- The **0** in the prefix length indicates that you are not interested in any bit in the prefix itself.
- The **ge 32** indicates that the length of the prefix (that is, the subnet mask) must be 32 (that is, 255.255.255.255), thus matching host routes.

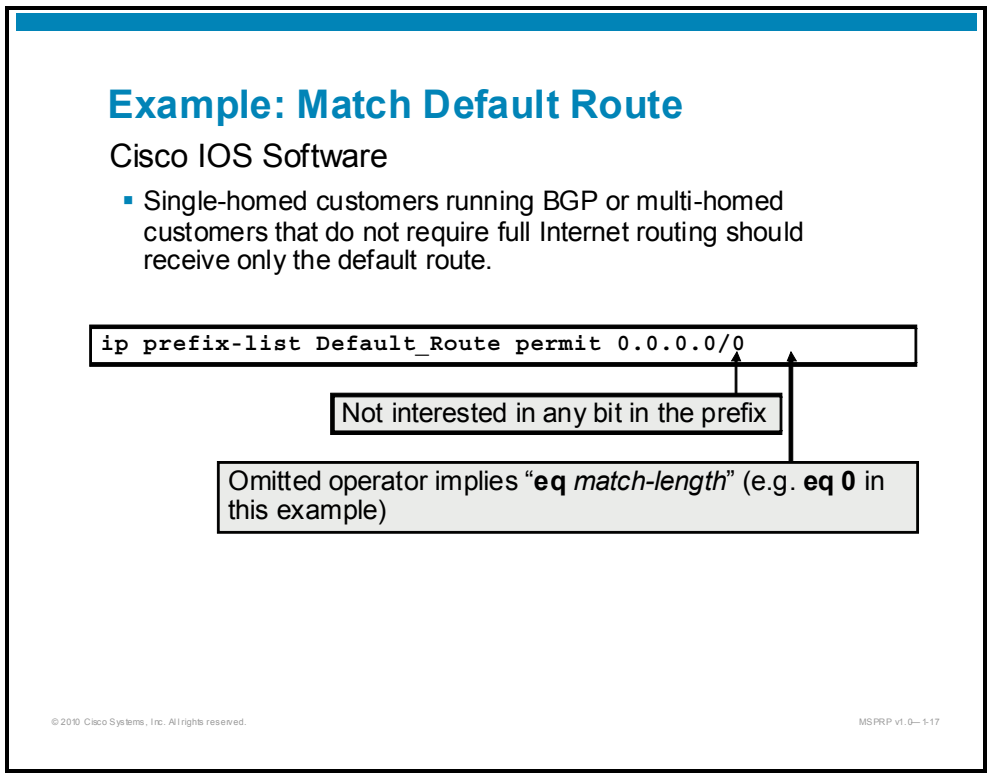

The sample prefix list shows how to match a default route:

- The **0** in the prefix length indicates that you are not interested in any bit in the prefix itself.
- The omitted operator indicates that the prefix length should be the same as the number of bits in the prefix you are trying to match (that is, **0**, which is the length of the subnet mask in a default route).

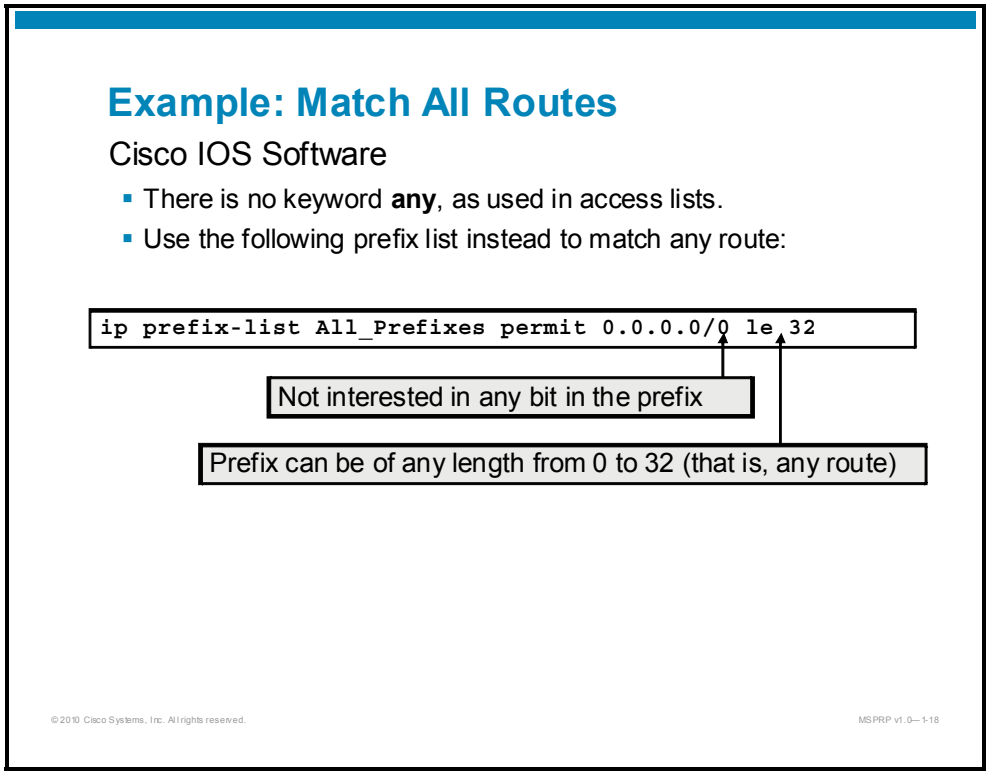

The sample prefix list shows how to match any route:

- The **0** in the prefix length indicates that you are not interested in any bit in the prefix itself.
- The **le 32** operator indicates that the prefix can be of any length from 0 to 32, thus matching any subnet mask.

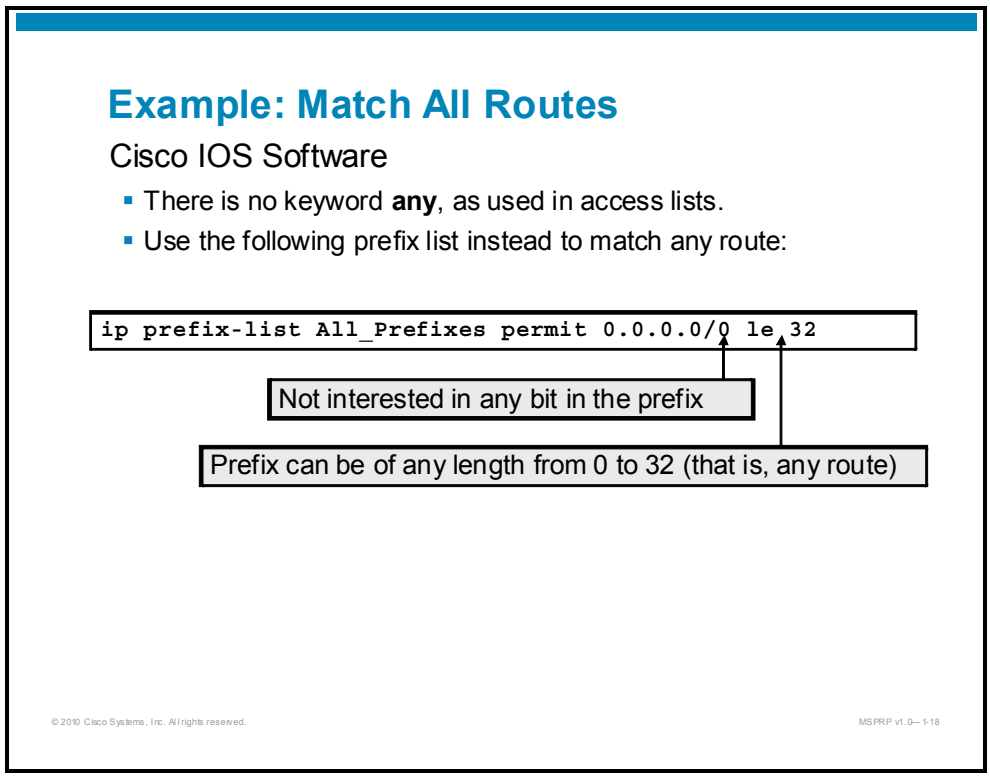

The sample prefix list shows how to match any prefix:

- The **0** in the prefix length indicates that you are not interested in any bit in the prefix itself.
- The **le 32** operator indicates that the prefix length can be anything from 0 to 32, thus matching any subnet mask from 0.0.0.0 to 255.255.255.255.

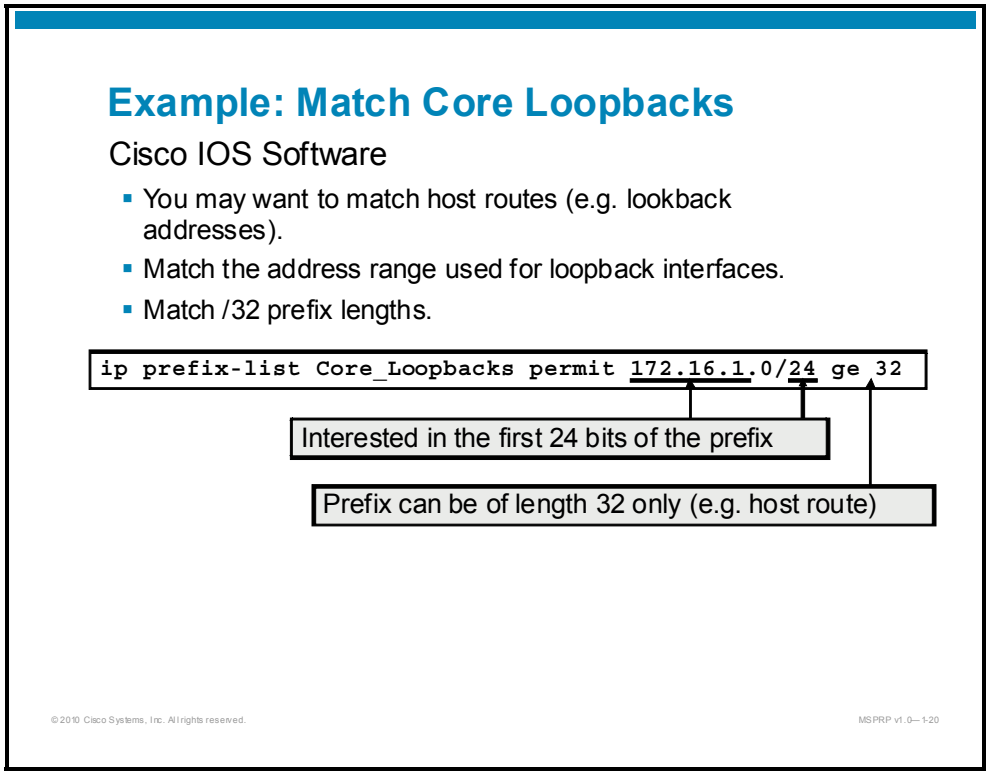

The sample prefix lists matches all host routes in a given range of prefixes (for example, 172.16.1.1/32, 172.16.1.2/32, and so on). This type of prefix list is useful for matching loopback addresses.

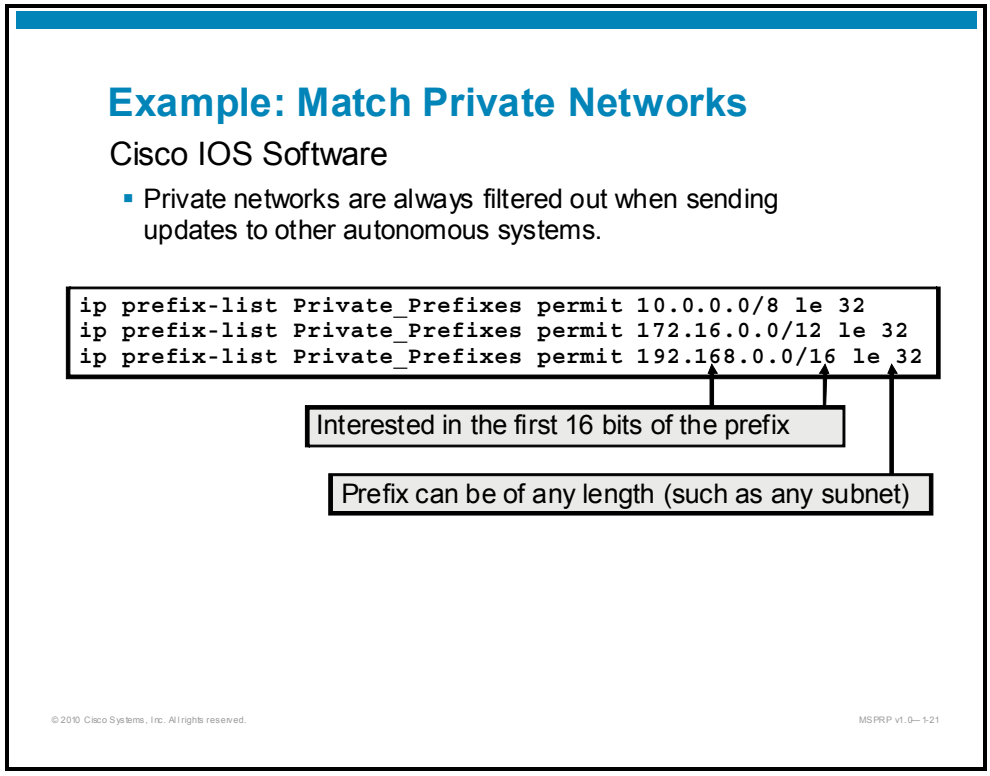

The sample prefix list matches any network or subnet in the RFC 1918 range of IP addresses (that is, the private address space). These private networks are typically filtered out on routing exchange between AS.

The **le 32** operator is used whenever you are not interested in the size of the prefix (that is, to match any subnet).

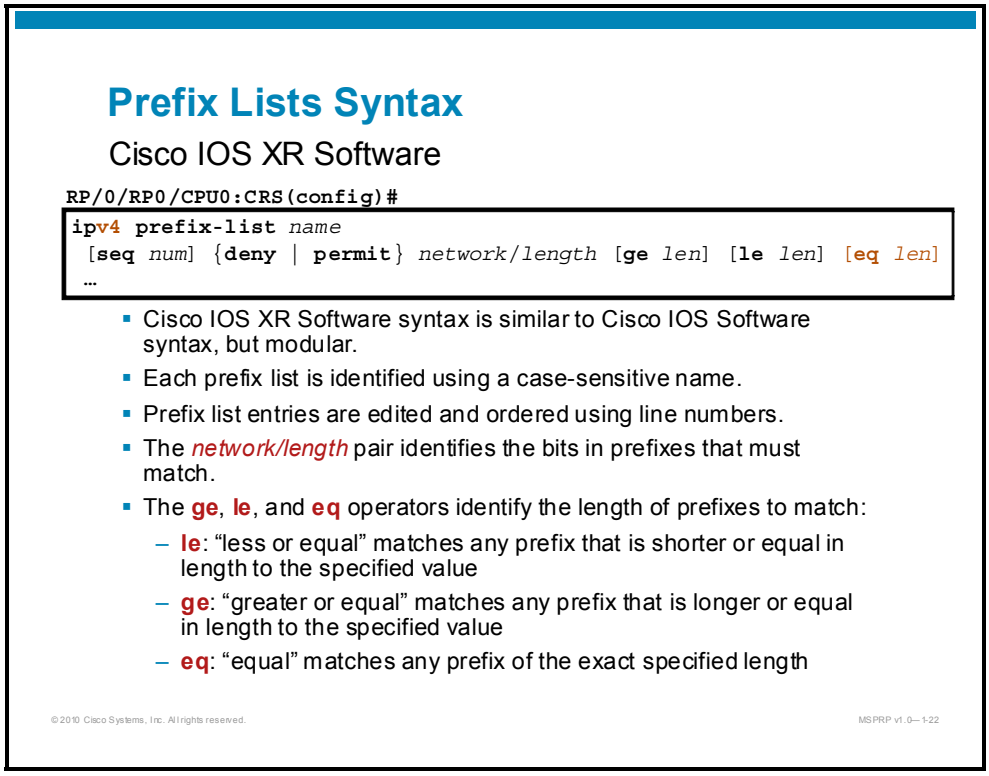

Prefix list syntax in Cisco IOS XR Software is different from Cisco IOS Software only in that it also implements the **eq** operator to match an exact prefix length.

Example: Prefix Lists

Cisco IOS XR Software

```
ipv4 prefix-list Private_Prefixes
 deny 10.0.0.0/8 le 32
 deny 172.16.0.0/12 le 32
 deny 192.168.0.0/16 le 32
 permit 0.0.0.0/0 le 32
!
ipv4 prefix-list Core_Loopbacks
 permit 172.16.1.0/24 eq 32
!
ipv4 prefix-list Host_Routes
 permit 0.0.0.0/0 eq 32
!
ipv4 prefix-list Default_Route
 permit 0.0.0.0/0
!
ipv4 prefix-list All_Prefixes
 permit 0.0.0.0/0 le 32
!
ipv4 prefix-list Small_Prefixes
 permit 0.0.0.0/0 le 24
!
```


The Cisco IOS XR Software example lists all of the previous examples for Cisco IOS Software. The matching of core loopbacks was modified to use the **eq** operator, although it would also work with the **ge** operator.

The sample Private_Prefixes prefix list shows how to filter out all RFC 1918 prefixes. These types of filters are commonly used on incoming and outgoing updates on External Border Gateway Protocol (EBGP) sessions.

The sample Core_Loopbacks prefix list illustrates how to match host routes that can be used to match loopback addresses from a given address range.

The sample Host_Routes prefix list illustrates how to match any host route.

The sample Default_Route prefix list illustrates how to only match the default route.

The sample All_Prefixes prefix list shows how to match any (all) prefixes. The prefix list line is equivalent to the “any” keyword that is used in access lists to match any network.

The sample Small_Prefixes illustrates how to filter out all prefixes that are longer than 24 bits (filter to small prefixes).

AS Path-Based Filtering

Describe the characteristics and usage scenarios for AS path-based filtering in service provider environments.

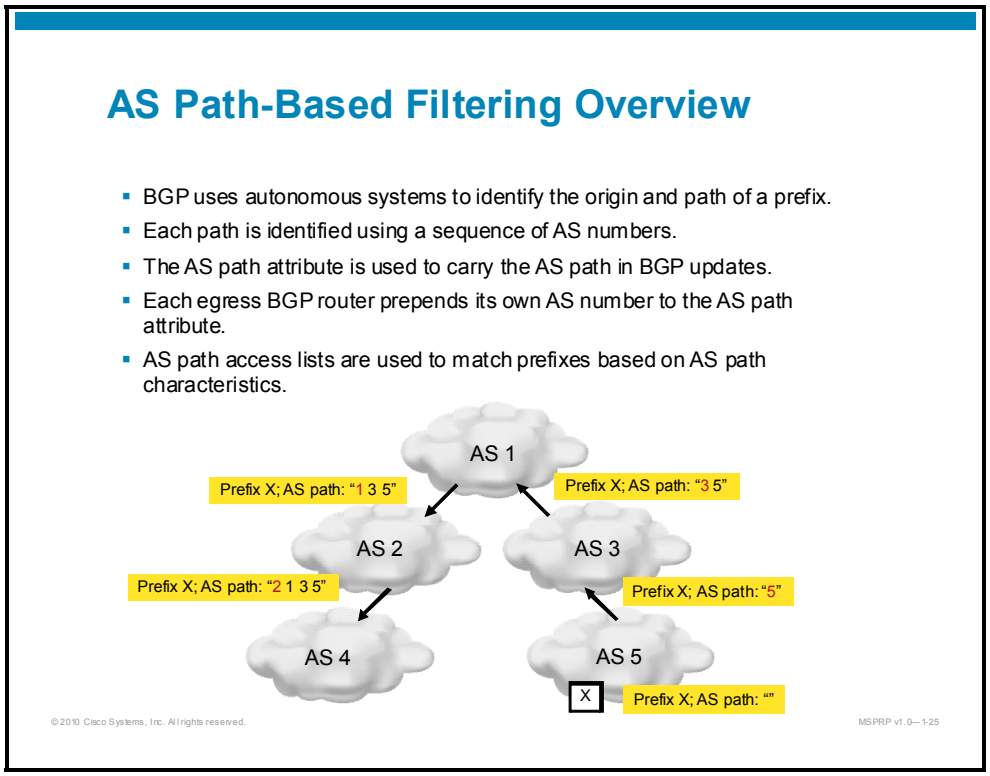

The figure illustrates the automatic prepending that is done by all egress routers when sending updates to neighboring autonomous systems. It shows that the first number in the AS path is always the number of the neighboring AS from which the update was received. The last number in the AS path is the number of the originating AS.

An AS path access list can be used to identify various updates based on the characteristics of their AS path attribute. Regular expressions are used to process AS path attributes.

AS Path Access List Syntax

Cisco IOS Software

```
Router(config)#
ip as-path access-list acl-number {permit | deny} regexp
```

- Each AS path access list is identified using a unique number.
- Regular expressions are used to match prefixes based on the contents of the AS path attribute.
- The AS path is processed as a string of characters.

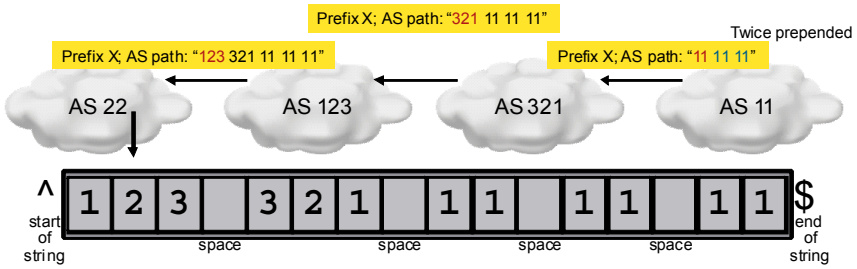

Each AS path access list is identified using a unique number in the range from 1 to 500. Regular expressions are used to match prefixes based on the contents of the AS path attribute that is converted to a string of characters.

The figure illustrates an AS path attribute as seen in AS 22. The AS path is converted to a string of characters that starts with the character "1" and also ends with the character "1" in this AS path example. Regular expressions must be written to take into account that you typically want to identify AS numbers and not individual characters.

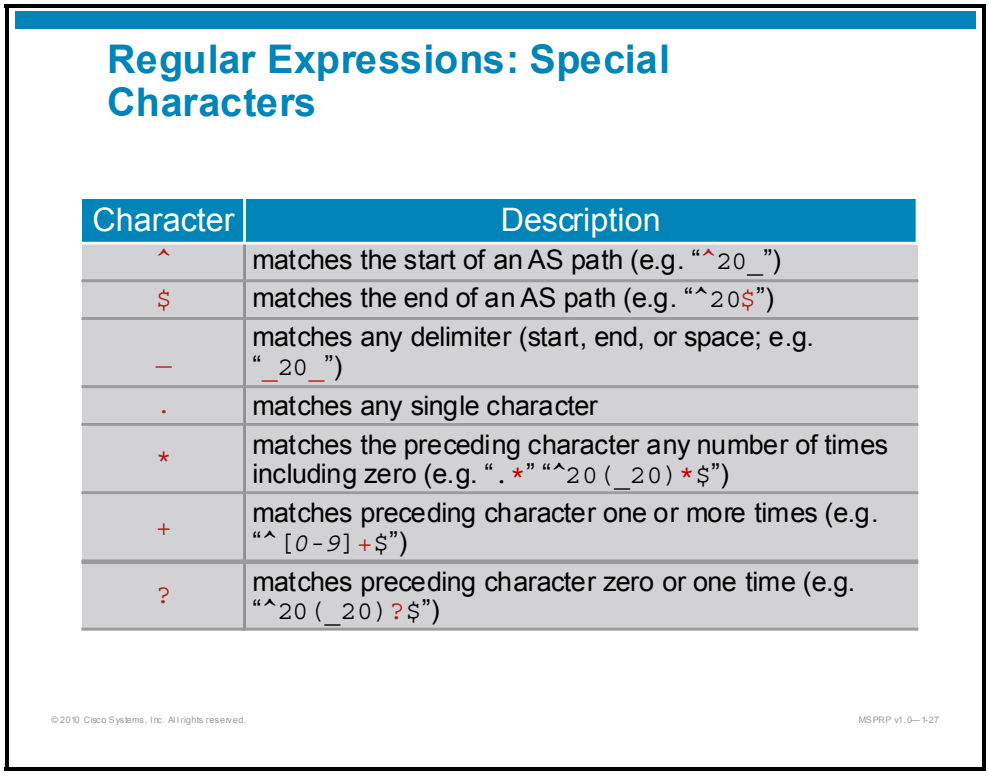

Regular Expression Special Characters

Character	Description
^ (caret sign)	Matches the start of the AS path string
$ (dollar sign)	Matches the end of the AS path string
_ (underscore)	Matches any delimiter, space and including a space and the start or end of the AS path string
. (dot)	Matches any single character
* (asterisk)	Matches any single preceding character zero or more times
+ (plus)	Matches any single preceding character one or more times
? (question mark)	Matches any single preceding character once or not at all

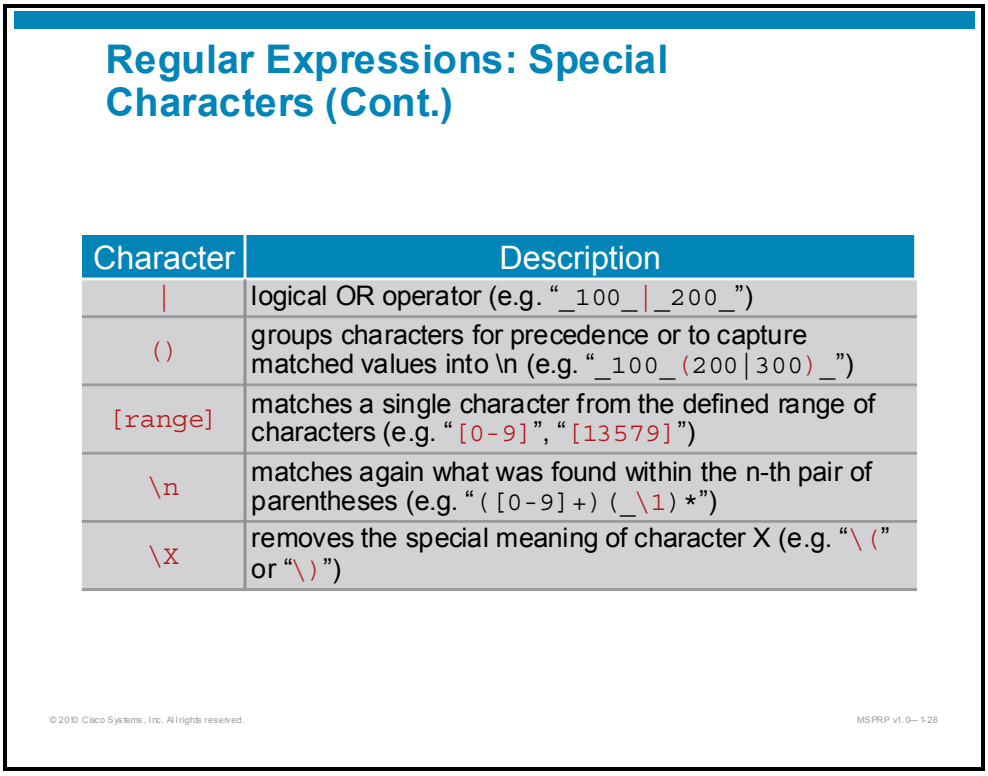

Regular Expression Special Characters (Cont.)

Character	Description
| (vertical bar)	Used to represent a logical OR operator. This character has the lowest precedence. It means that the regular expression must either match what is to the left of the sign, to the right of the sign, or both.
() (parentheses)	Used to group characters in a regular expression: ■ in order to affect the precedence of operators (for example, "_20_30_|_40_50_ vs. _20_(30|40)_50_") ■ in order to affect the grouping of characters for quantifiers "*" and "+" and "?," which normally only apply special meaning to a single preceding character (for example, "(_20)*") ■ to store the matched character in a temporary variable that can later be referenced using the \n expression
[] (square brackets)	Used to match a single character from a defined range of characters.
\ (backslash sign)	Used followed by a number *n* to match once more what was matched within the *n*-th parentheses in the same expression. The backslash character can also be used followed by a character to remove the special meaning from the character.

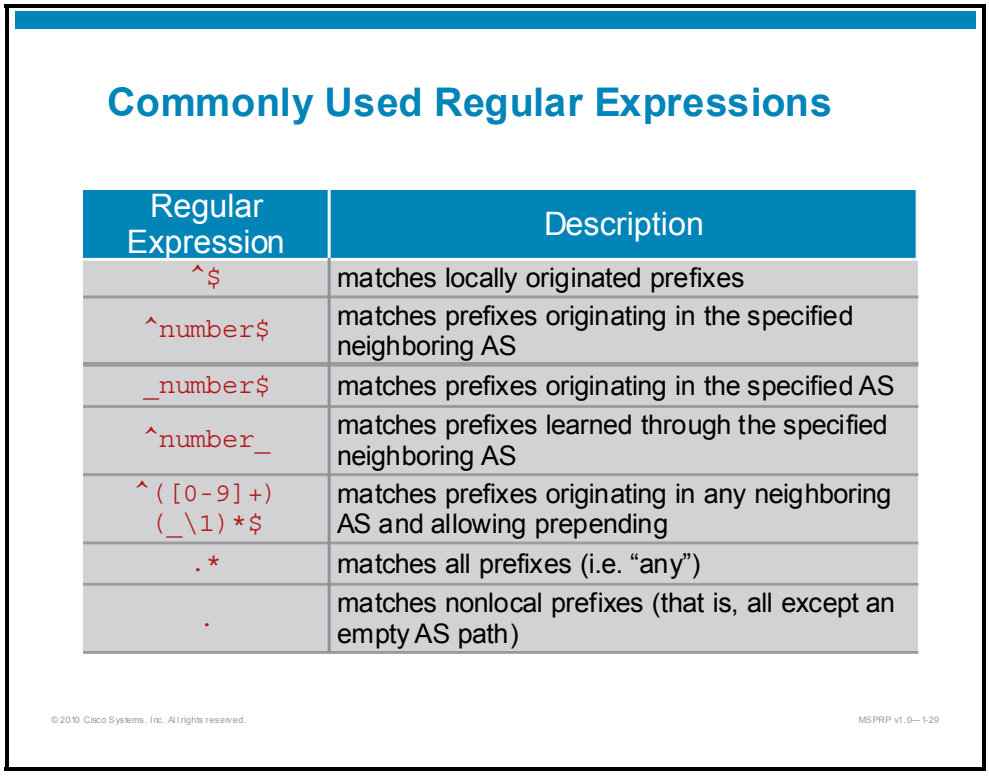

Commonly Used Regular Expression

Regular Expression	Description
`^$`	Matches all local routes (local routes have an empty AS path attribute).
`^10$`	Matches any route originating in a neighboring AS 10.
`_20$`	Matches any route originating in AS 20.
`^10_`	Matches any route that is received from a neighboring AS 10.
`([0-9]+)(_\1)*`	Matches any AS number, which can optionally repeat any number of times (that is, prepending). The **\1** references whatever is matched in the first pair of parentheses.
`.*`	Matches any character any number of times. This regular expression is used to match any prefix.
`.`	Matches any single character.
`_20_`	Matches any route originating or passing through AS 20.
`[0-9]+`	Matches any AS number from 0 to 65335 (maximum range in BGP).
`[13579]$`	Matches routes originating in odd-numbered AS.
`[02468]$`	Matches routes originating in even-numbered AS.

Example: Permit All Routes

Cisco IOS Software

- There is no keyword “any,” as used in access lists.
- Use this regular expression instead to match any route.
- Example:
  - The filter matches any prefix from any neighbor.

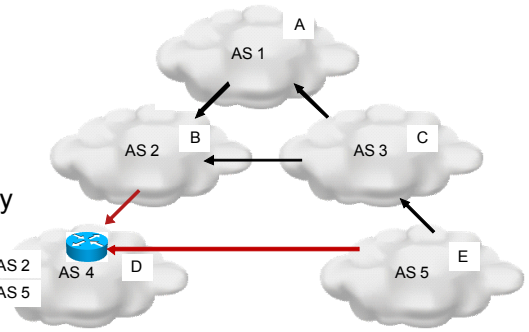


```
ip as-path access-list permit .*
```

Matches any character (.) any number of times (*)

The sample regular expression **ip as-path access-list permit .*** matches any character any number of times. This AS path access list entry is used to permit any route (that is, the equivalent of the “any” keyword in access lists).

The figure illustrates five autonomous systems, each represented by one prefix that it advertises. AS 1, for example, advertises the prefix “A,” which can be learned by AS 4 from AS 2, AS 5, or both.

If you apply the example filter in AS 4 to incoming updates from AS 2 or AS 5, you will accept all routes.

Example: Permit Local Routes

Cisco IOS Software

- Locally originated routes have an empty AS path attribute.
- Customers typically use this outbound filter to only announce their prefixes to their ISPs.
- Example:
  - AS 4 only sends its own prefixes.

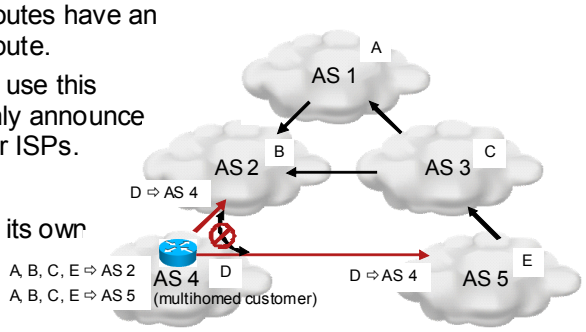


```
ip as-path access-list permit ^$
```

Matches an empty AS path attribute (i.e. no character from start to end)

The sample regular expression **ip as-path access-list permit ^$** matches any route that has an empty AS path attribute (that is, no character from start to end). Only locally originated routes have an empty AS path attribute, hence this regular expression is used when matching local routes.

Multihomed customers use this type of filter to only send their address space to their service providers to prevent them from becoming transit autonomous systems. In the figure, AS 4 only advertises its own prefix (“D”) to its providers. Other prefixes that are received from one provider are not forwarded to the other provider.

Example: Permit Routes From a Neighbor

Cisco IOS Software

- The first number in the AS path is the last prepended number.
- A directly connected neighboring AS is always found as the first number in the AS path.
- This is typically used for routing policies.
- Example:
  - AS 4 matches any prefix from neighboring AS 5.

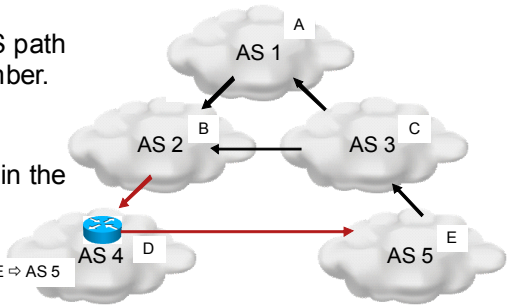


```
ip as-path access-list permit ^5_
```

Matches routes coming from a specific neighbor

The sample regular expression **ip as-path access-list permit ^5_** matches any route that is received from neighboring AS 5—the first number in the AS path. All prefixes that are received from AS 5 are accepted. If the same filter is applied to incoming updates from AS 2, the prefixes will be denied.

This type of filter is typically used when creating routing policies (for example, assigning different local preference values).

Example: Permit Routes Originating in a Specific AS

Cisco IOS Software

- The last number in the AS path is the first prepended number.
- The originating AS is always found as the last number in the AS path.
- Typically used for routing policies.
- Example:
  - AS 4 matches prefixes originating in AS 1 from any neighboring AS.

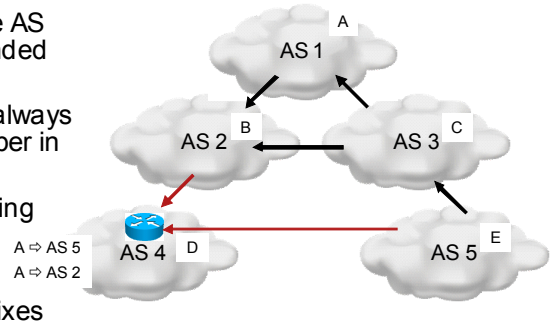


```
ip as-path access-list permit _1$
```

Matches routes coming from a specific neighbor

The sample regular expression **ip as-path access-list permit _1$** matches any route that is originated in AS 1—the last number in the AS path. If this filter is applied to incoming updates from AS 2 or AS 5, it will permit the prefix “A” originating in AS 1. This type of filter is commonly used to implement routing policies in which you can assign preference for certain prefixes coming from a preferred ISP.

Example: Permit Neighbors' Local Routes

Cisco IOS Software

- A single AS number in an AS path denotes prefixes originating in a neighboring autonomous system.

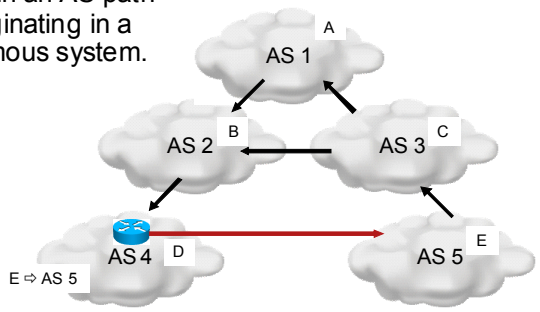


```
ip as-path access-list permit ^5$
```

Matches a single AS number in the AS path (i.e. a prefix originating in a neighboring AS)

The sample regular expression **ip as-path access-list permit ^5$** matches any route that is originated in neighboring AS 5—the first number in the AS path. In the example, AS 4 only accepts the prefix "E" from AS 5, because other prefixes originate in other autonomous systems.

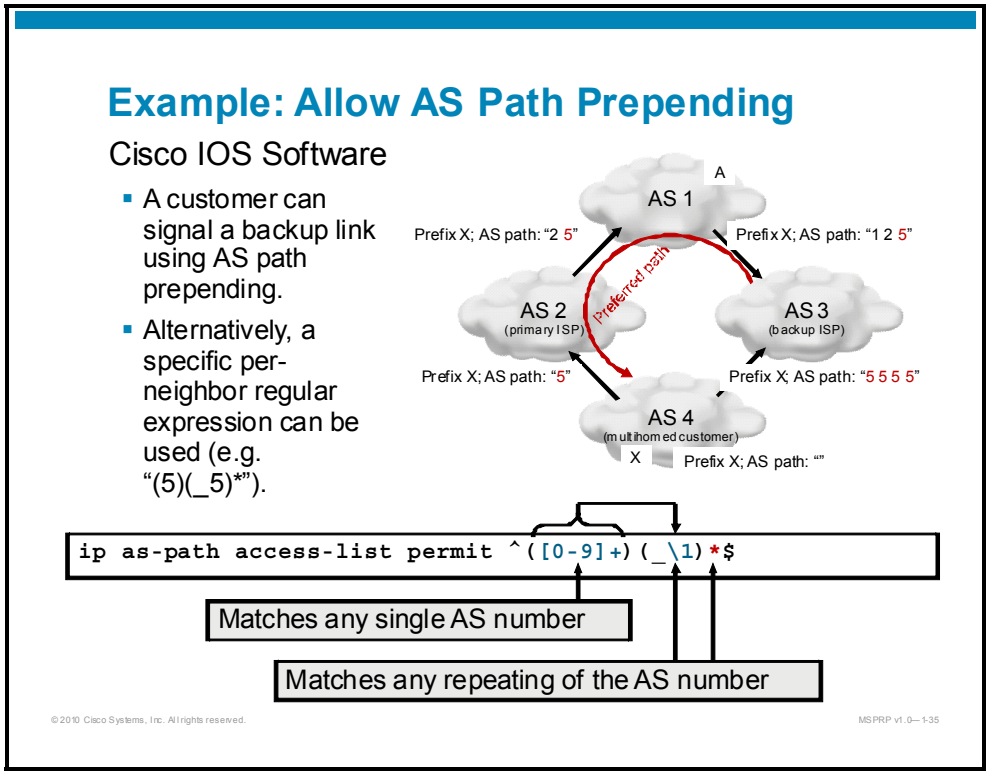

The figure illustrates the generic filter **ip as-path access-list permit ^([0-9]+)(_\1)*$**. This filter can be used on any neighboring AS where you wish to accept the neighbor's local prefixes while still allowing the neighbor to perform AS path prepending.

Enclosed within the first parentheses is a range of digits that can appear multiple times. In BGP, this filter effectively matches any number from 0 to 65535. Enclosed within the second pair of parentheses you reference whatever was matched in the first parentheses and allow that number to repeat zero or more times.

AS 3 can reach prefix X via AS 1 or via AS 5 directly. However, since AS 4 is using AS path prepending for its updates to AS 3, AS 3 will prefer the seemingly shorter AS path, which is through AS 1 to reach prefix X.

Route Maps

Describe the characteristics and usage scenarios for route maps in service provider environments.

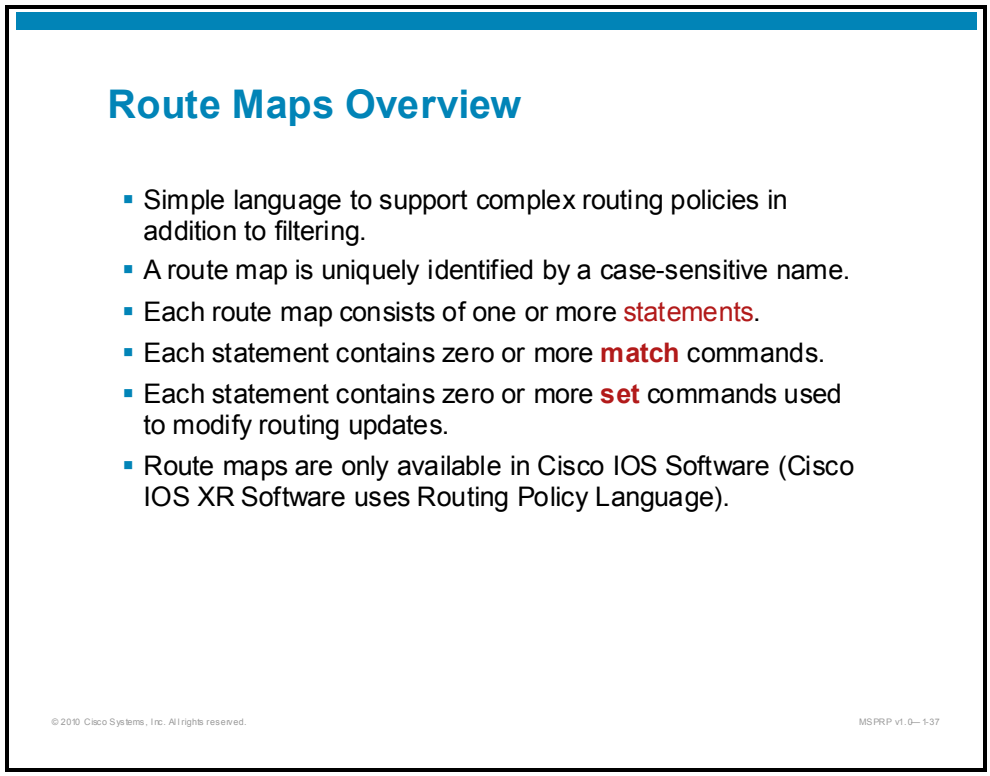

Route maps are a simple language to support complex routing policies in addition to filtering. Each route maps is uniquely identified by a case-sensitive name and consists of one or more statements. Each statement contains zero or more **match** and **set** commands. The **match** command is used to identify which routes should be processed in a given statement. The **set** command specifies which parameters should be modified or added in a routing update.

Route maps are not available in Cisco IOS XR Software. Instead Cisco IOS XR Software uses RPL.

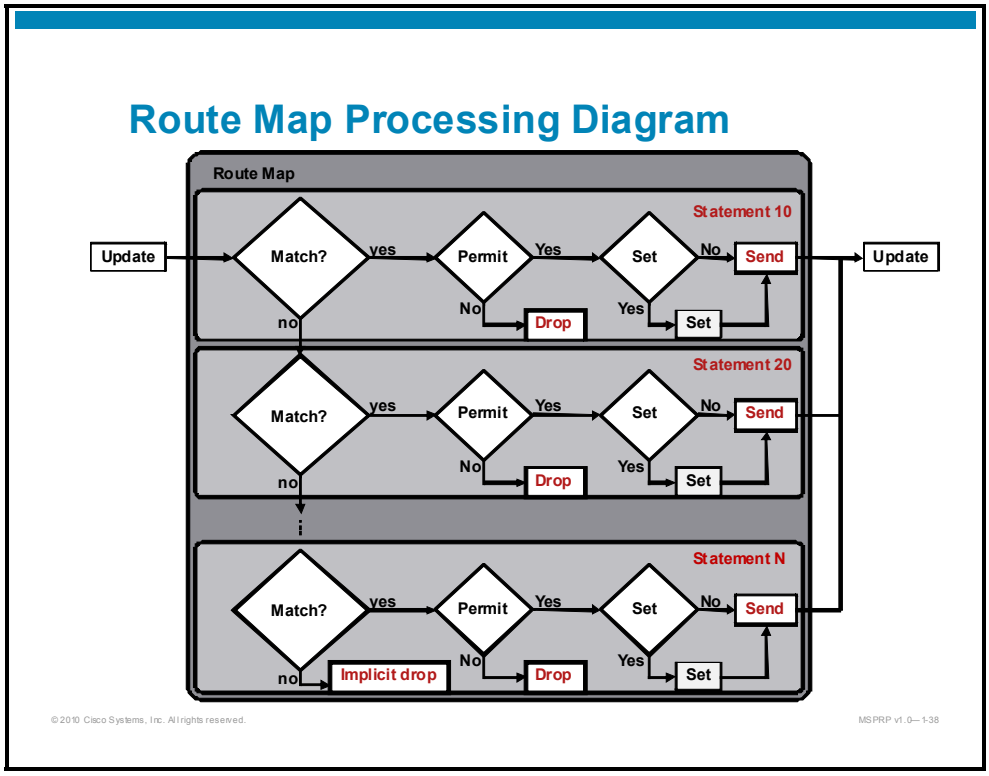

The diagram illustrates the components and processing of a route map:

- A route is processed by route-map statements in the order that is defined by sequence numbers.
- If a route matches the match conditions, it is processed by that statement.
- If the statement uses the “deny” options, the route is immediately dropped.
- If a route has one or more **set** commands, the **set** commands are processed, resulting in modified or added parameters and attributes.

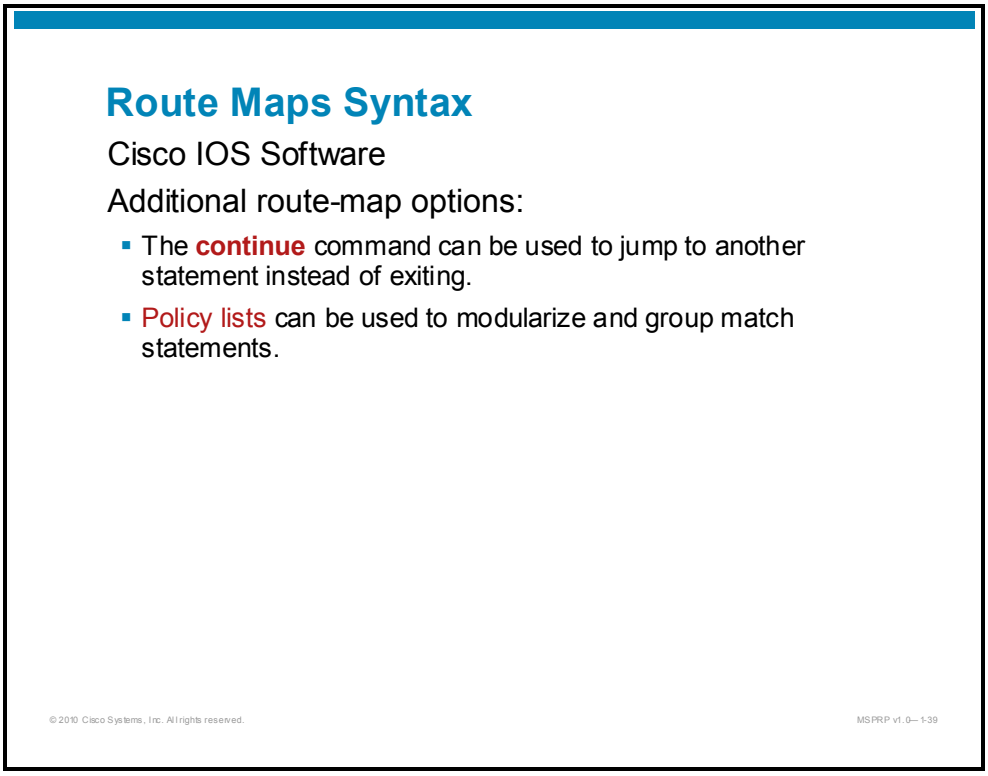

In addition to the components shown in the previous diagram, route maps also have the **continue** command. This command allows processing to continue in another statement (that is, jump to the next or specified **route-map** command).

Complex match options can be grouped in policy lists and then reused in various route maps for more modularity and reusability.

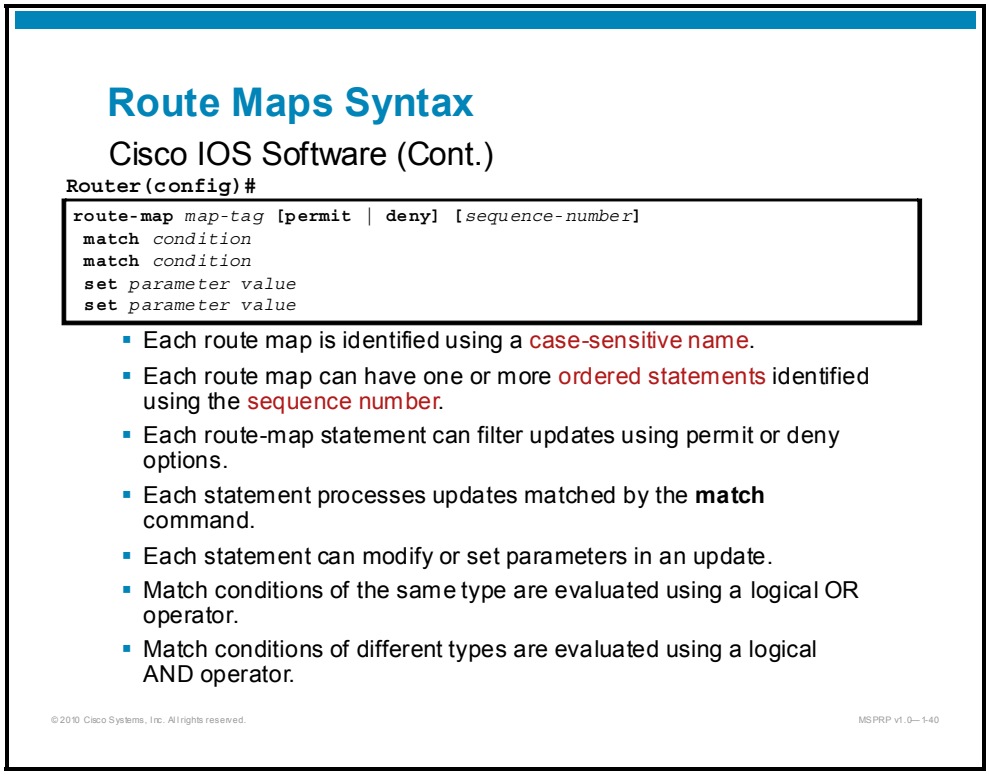

A route-map statement processes routes that match a match condition. For example, any prefix that is permitted by a prefix list will be matched; any prefix denied by a prefix list will not be processed by the statement and will instead be evaluated by the next route-map statement. If a route matches, the route-map statement can then permit or deny it.

If there are multiple match conditions, they are evaluated using the following rules:

- Match conditions of the same type are evaluated using the logical OR operator (that is, the prefix must match at least one condition).
- Match conditions of different types are evaluated using the logical AND operator (that is, the prefix must match all conditions).

Routes that are matched and permitted by a statement can be modified using **set** commands.

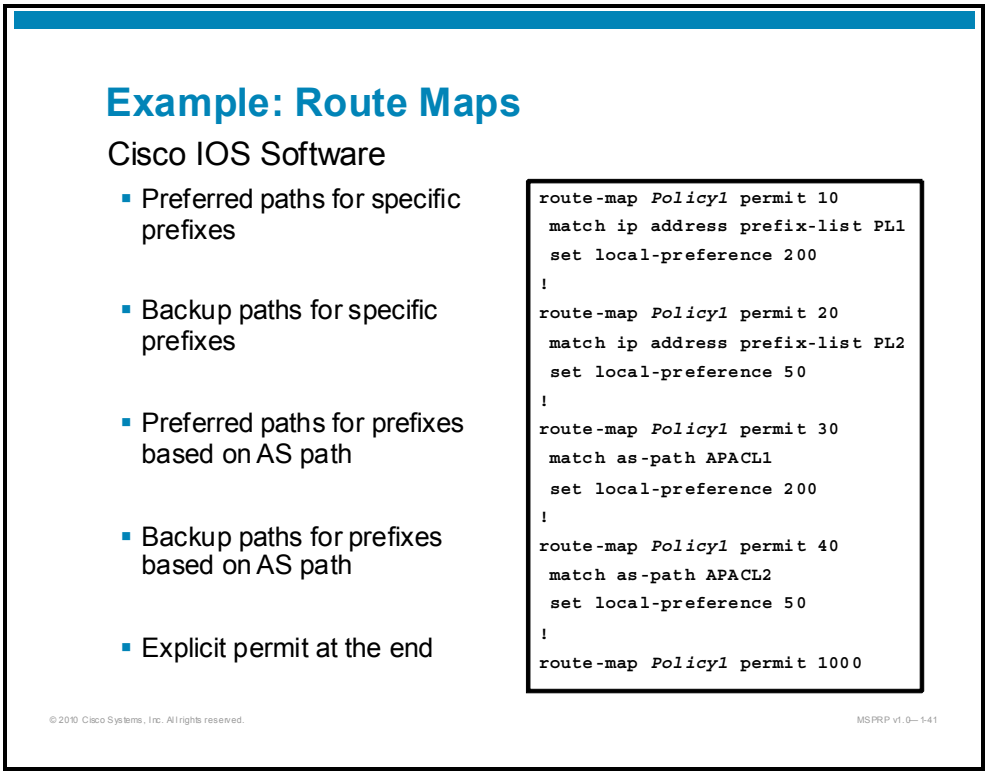

This sample route-map configuration consists of five route-map statements. The first two process routes that are matched by a prefix list based on the prefix and set appropriate BGP local preference attributes. The next two statements match routes using AS path access lists and also set appropriate BGP local preference values. All routes that do not match are passed unchanged by explicitly permitting them at the end.

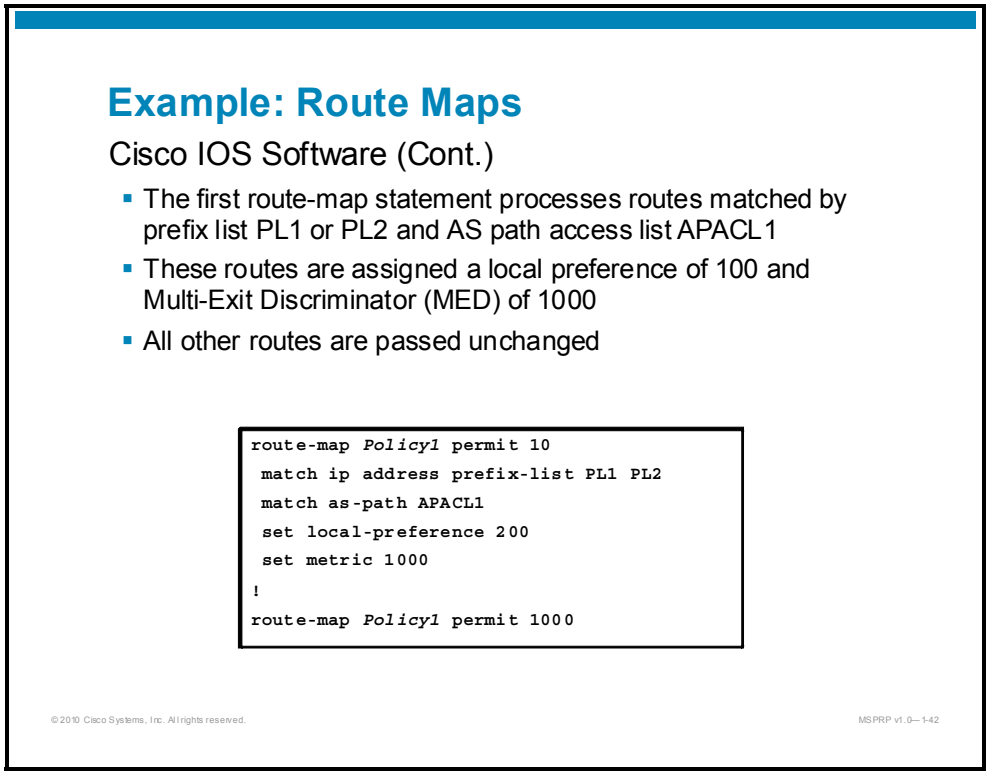

This sample configuration illustrates the logical processing of multiple match conditions. The first condition uses two prefix lists, of which a route must match at least one (logical OR). In order for this statement to process a route, the route must also match the second match command, which uses an AS path access list (logical AND).

A single match statement may contain multiple conditions of the same type (prefix lists PL1 and PL2, in the previous example). At least one condition in the match statement must be true for that match statement to be considered a match (logical OR).

A route-map statement may also contain multiple match statements of different types (prefix lists and AS path access lists in the previous example). All match statements must be true for the route-map statement to be considered a match (logical AND).

The previous example can be illustrated as "(PL1 **OR** PL2) **AND** APACL1."

Routing Policy Language

Describe the characteristics of RPL.

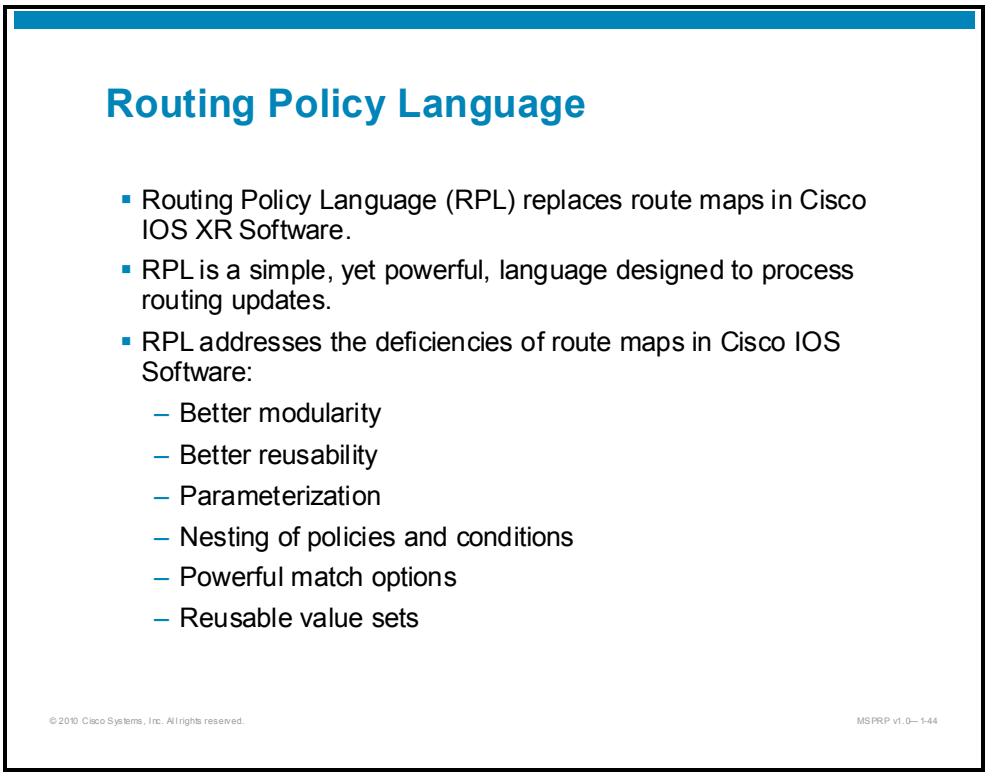

RPL is a newer mechanism that was introduced into Cisco IOS XR Software as a replacement for and improvement upon the route maps that are used in Cisco IOS Software.

RPL offers a more powerful set of tools to process routes:

- Modularity by allowing policies to reference other objects such as prefix list, value sets, and other policies (that is, nesting of policies)
- Parameterization for optimization and better reusability of policies

RPL Overview Modularity

- Each routing policy is identified by a case-sensitive name.
- The entire policy is defined between the **route-policy** and **end-policy** commands.
- The main RPL functions are as follows:
  - Filtering of updates (**pass** and **drop** commands)
  - Modification of attributes (**set** commands)

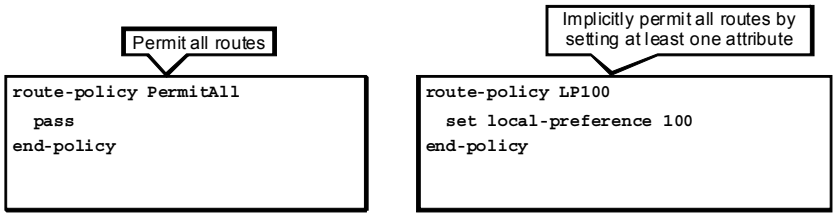

Like route maps or many other objects, routing policies are identified using case-sensitive names. Each routing policy is a single object (no sequence numbers, multiple lines, or statements).

Like route maps, routing policies are also filtering tools that allow you to permit or deny routing updates. The explicit commands to permit or deny are **pass** and **drop**, respectively.

Like route maps, routing policies can also modify or add parameters or attributes using the **set** command. A single **set** command also implicitly includes the **pass** command.

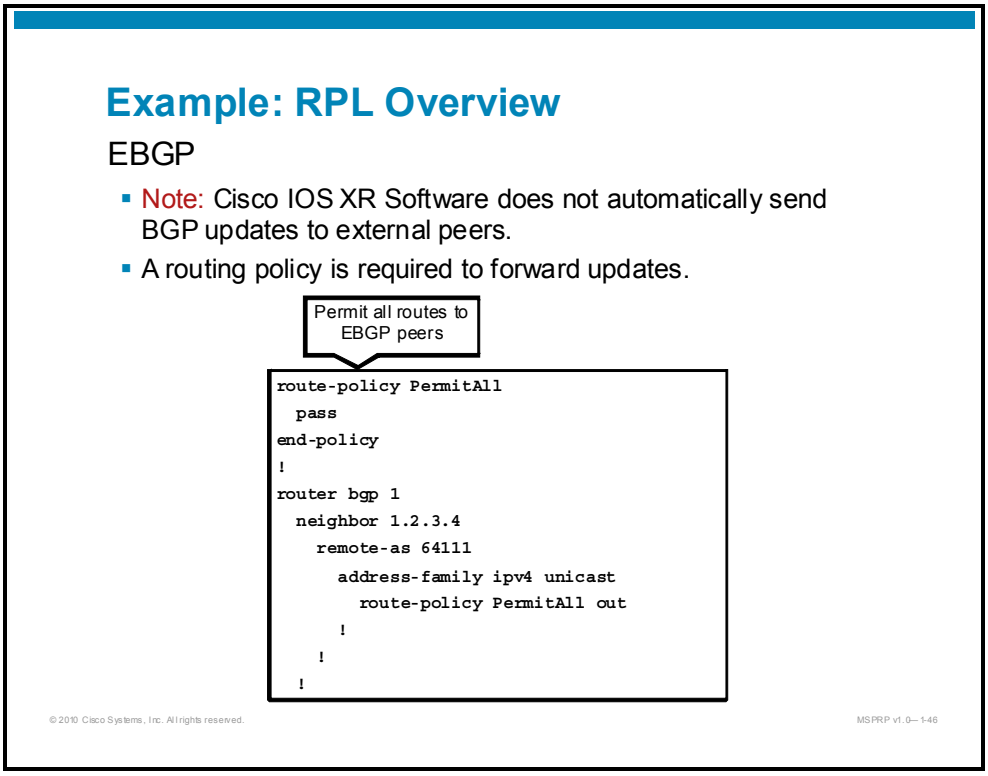

The sample configuration shows how to enable the forwarding of routing updates to an external BGP neighbor. In Cisco IOS XR Software, updates are not forwarded to an external neighbor unless an outbound policy is attached to the neighbor.

The sample configuration uses a simple policy to permit all routes.

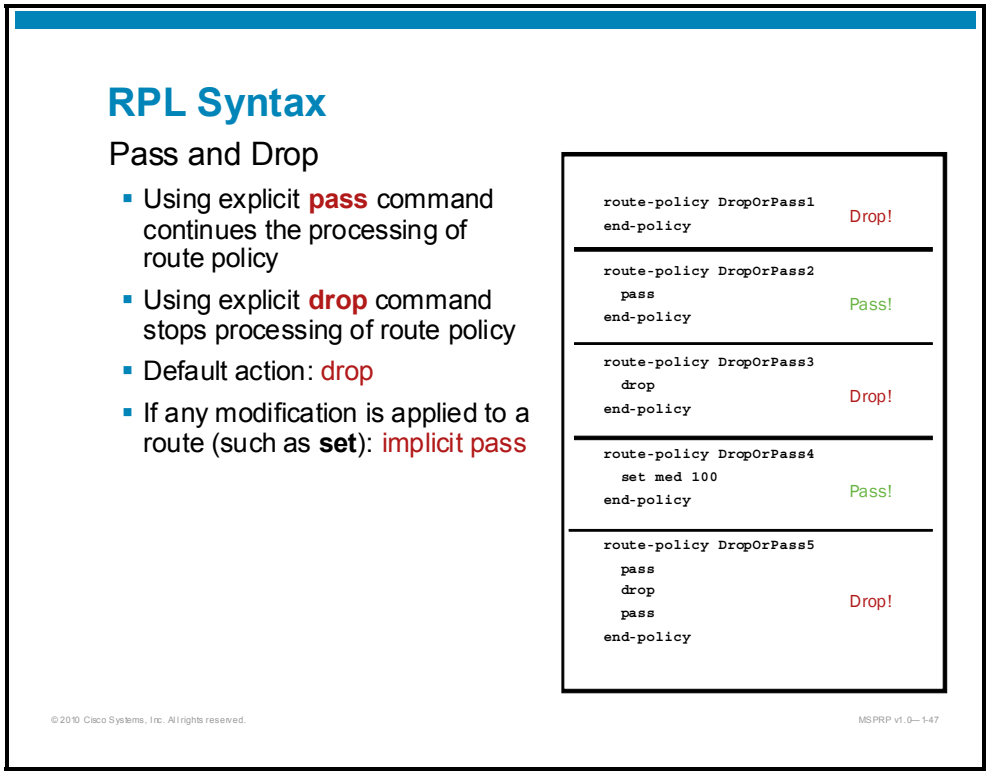

The sample configurations illustrate the routing policy rules:

- An empty policy implicitly denies all routes.
- An explicit **pass** without any conditions will forward all routes without any modifications.
- An explicit **drop** command will do the same thing as an implicit **drop** (deny all routes).
- A **set** command will modify the attribute accordingly and forward all routes.
- An explicit **drop** command will stop the processing of a policy and deny a route.

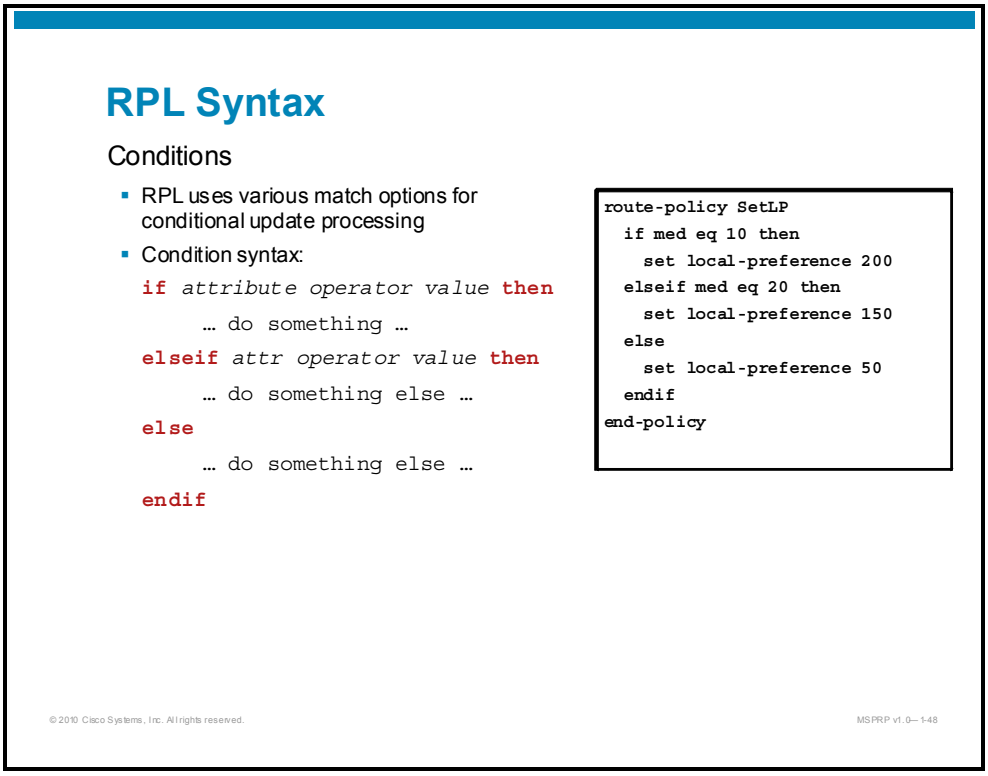

RPL uses conditional statement syntax that is found in many programming languages:

if *condition* **then** *operation1* **else** *operation2* **endif**

if *condition1* **then** *operation1* **elseif** *condition2* **then** *operation2* **else** *operation3* **endif**

The sample configuration illustrates how the multi-exit discriminator (MED) attribute can be used to influence a routing policy by setting a more powerful local preference attribute.

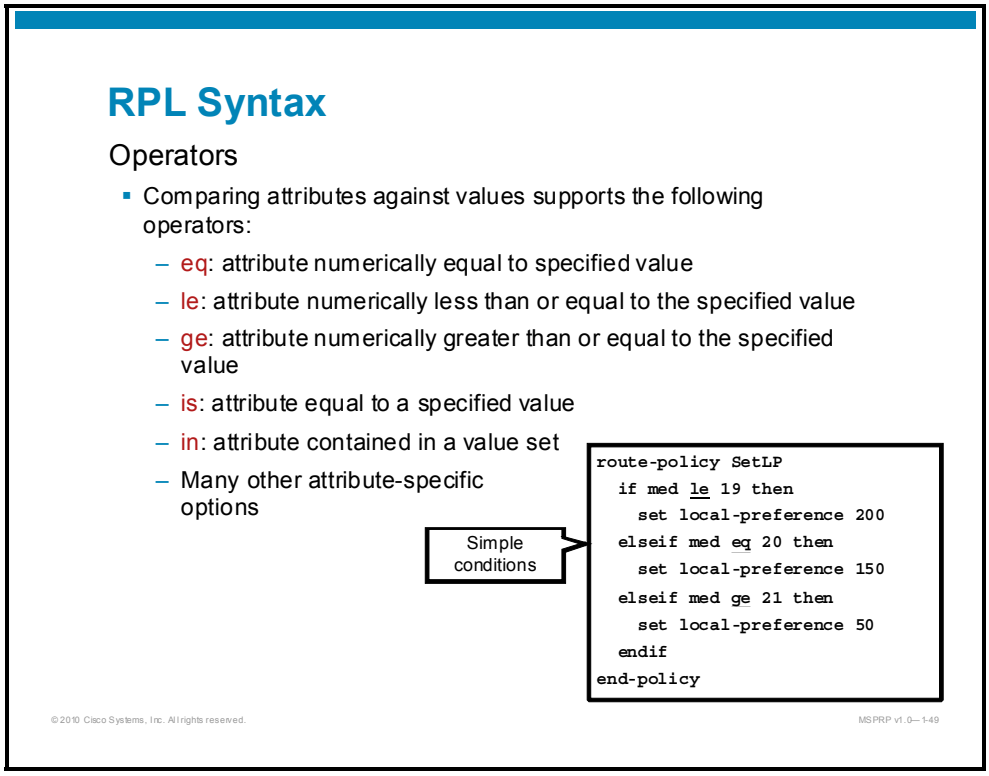

RPL conditions can use various operators:

- **eq**: attribute numerically equal to specified value
- **le**: attribute numerically less than or equal to the specified value
- **ge**: attribute numerically greater than or equal to the specified value
- **is**: attribute equal to a specified value (used for non-numerical values)
- **in**: attribute contained in a value set
- Many attribute-specific conditions (for example, AS path matching)

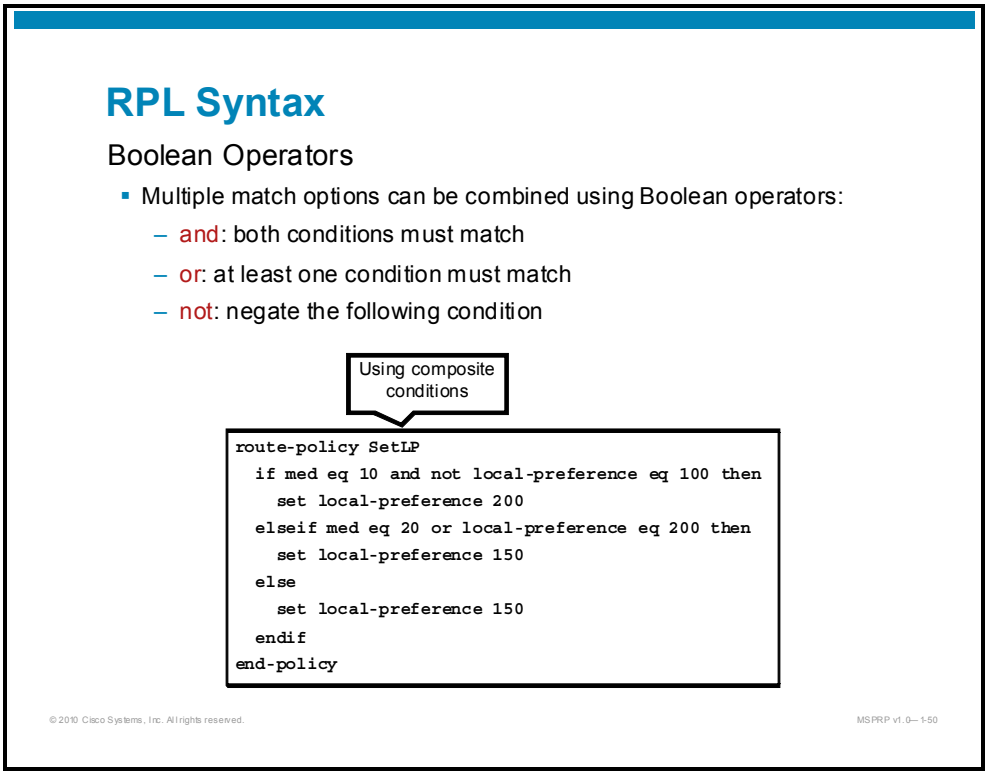

Boolean operators can be used to create complex compound conditions:

- Use the “and” operator if two or more conditions must match.
- Use the “or” operator if at least one of two or more conditions must match.
- Use the “not” operator to negate a condition.

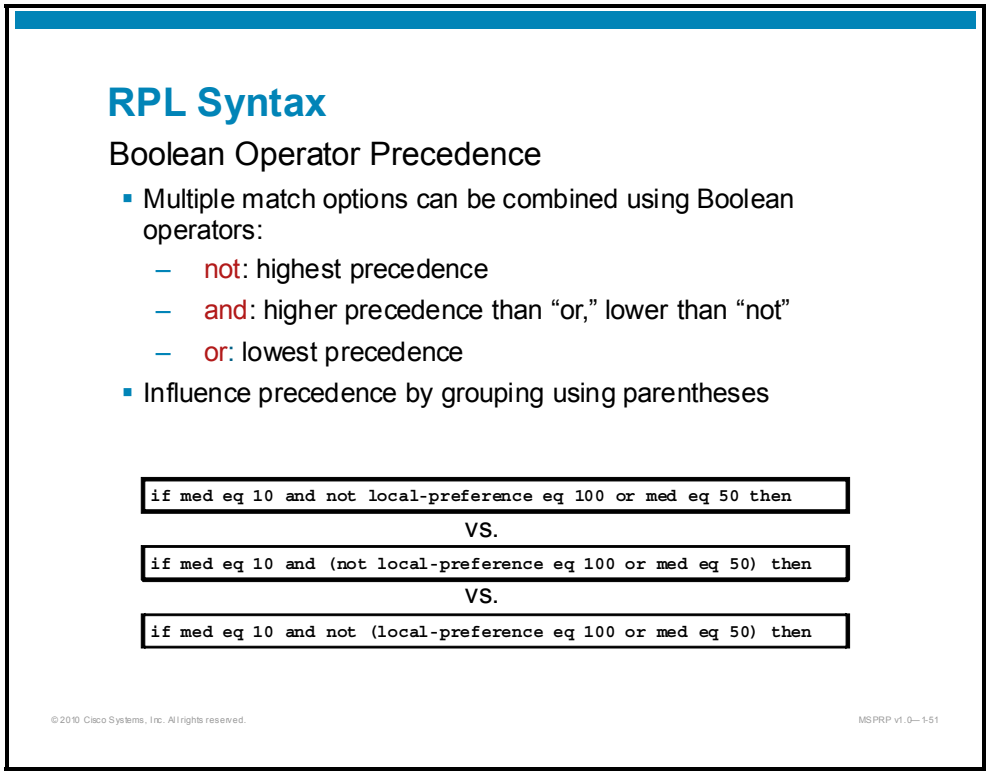

Use parentheses to influence the precedence of operators and achieve the desired result. The operators have the following precedence:

1. “not” is always evaluated first
2. “and” is evaluated second
3. “or” is evaluated last

The first example does not use parentheses. It can be written with parentheses to ensure the proper understanding of the condition:

if ((med eq 10) and (not (local-preference eq 100))) or (med eq 50) then

The second and third examples will result in different conditions.

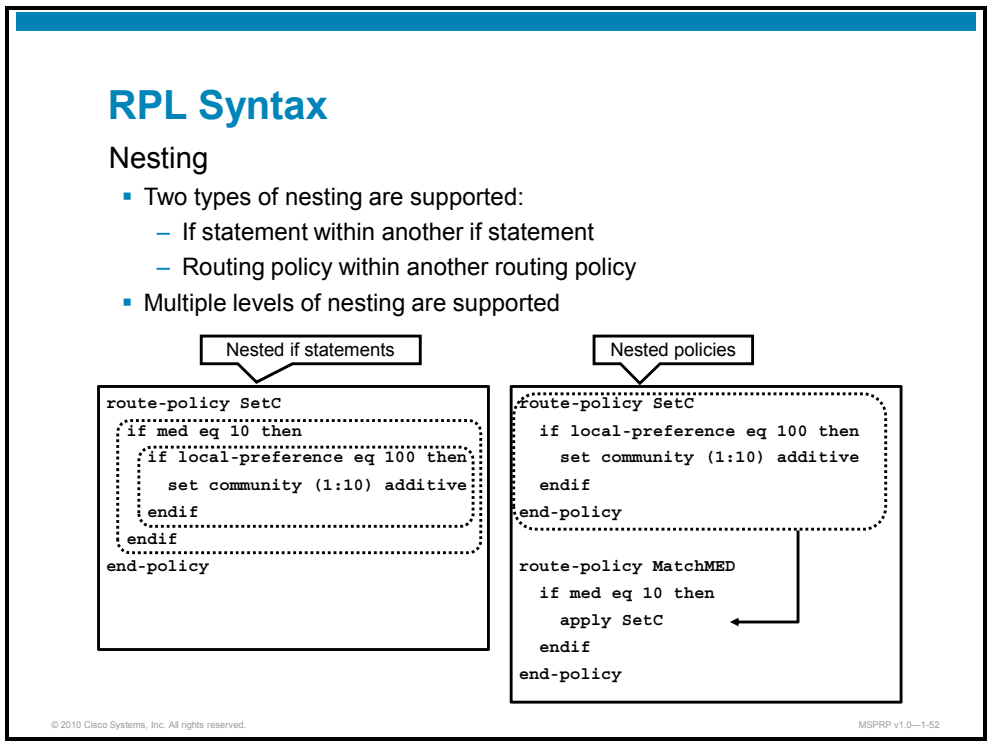

Large and complex routing policies should preferably be optimized by using modularization as much as possible. The left example shows a nested **if** statement and the right example shows a nested route-policy. The **SetC** policy in the right example can be reused in multiple policies to conditionally assign a BGP community. The **apply** command is used within a route-policy to call another route-policy.

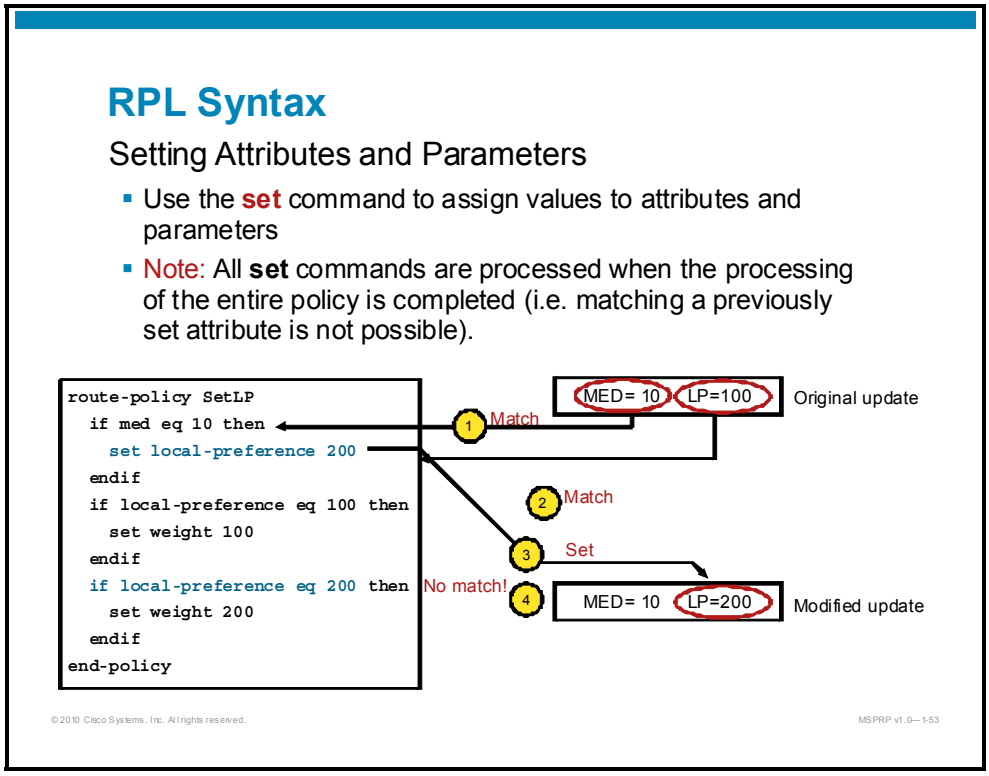

The figure illustrates a policy with multiple conditions and modifications based on the same parameter (BGP local preference). It is important to remember that modifications to an attribute are only executed when processing of the policy is completed, so conditions cannot be used based on previously modified values.

RPL Syntax

Setting Attributes and Parameters (Cont.)

- Note: Last **set** wins when multiple **set** commands are evaluated for a unique parameter.

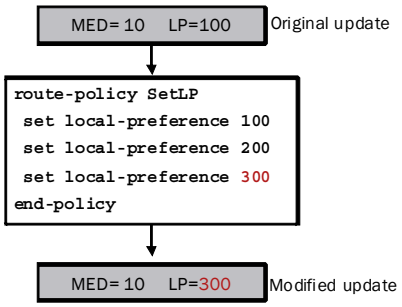

If multiple **set** commands are processed for the same attribute, the last one will be used when the processing of the policy is completed.

RPL Syntax

Setting Attributes and Parameters (Cont.)

- Note: All **set** commands are evaluated in the same order for non-unique attributes and operations.

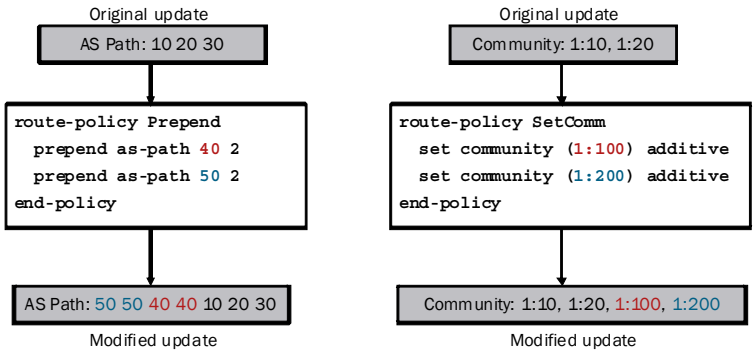

If a **set** command is processing an attribute that is a set of values, it may happen that all **set** commands will take effect.

The left example shows how the first **prepend** command modifies the AS path attribute by prepending 40 twice. The second **prepend** command then additionally prepends 50 twice.

The right example shows how two **set** commands add two values to a set of BGP Community attributes.

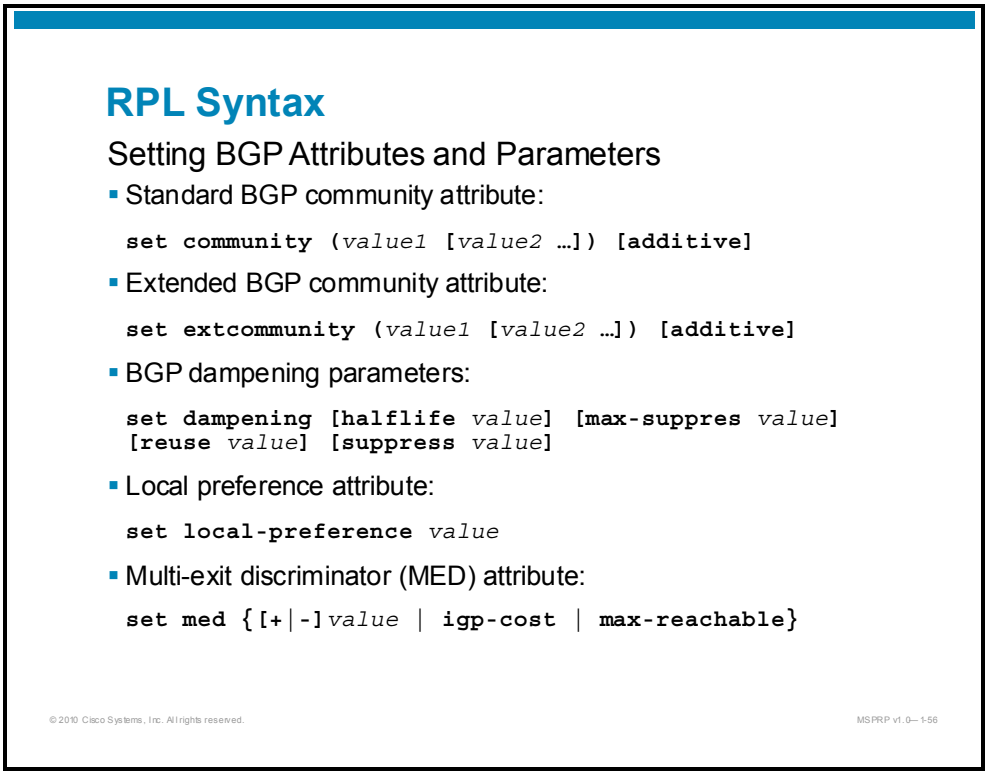

Setting standard and extended BGP community attributes:

- One or more values can be assigned to the BGP community attribute.
- If the **additive** keyword is used, the new communities will be added to the existing BGP communities.
- Omitting the **additive** keyword will cause existing BGP communities to be overwritten.

BGP Route Flap Dampening can be enabled and tuned using the **set dampening** command, where the dampening parameters can be specified to determine the aggressiveness of the dampening mechanism.

BGP local preference can be set on incoming updates or upon redistribution. The default BGP local preference is 100 and is set on all updates coming from external peers or being redistributed.

The BGP MED attribute can be set to a specific value or modified using the “+” or “-“ options. The MED can also be set to the underlying IGP cost or a maximum value.

Example: Setting BGP Attributes and Parameters

Route Flap Dampening

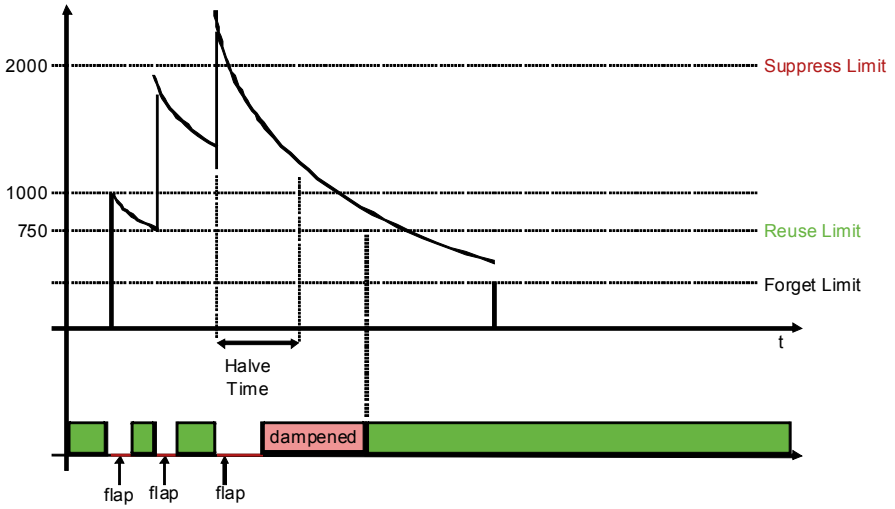

BGP Route Flap Dampening is a feature that is designed to make BGP more stable and consequently scale better by "punishing" routes that flap (disappear and reappear) more often. The default behavior of dampening results in stopping propagation of routes that consecutively flap three or more times in a short period, for a certain period.

The default behavior can be summarized:

- Each flap is penalized by adding 1000 penalty points to the penalty.
- If the cumulative penalty exceeds the suppress limit (2000 points by default), the route is dampened. In other words, it is stored in the BGP table, but is not evaluated in the best-path selection. Consequently, it is not installed into the routing table or forwarded to any neighbor. Routers remember the penalty when the route is not reachable by storing it as a "history" entry.
- The penalty is gradually decreased. The penalty reduction is determined by the halve time, which is 15 minutes by default.
- When a penalty drops below the reuse limit (750 by default) or when the route has been dampened for more than the maximum suppress time (one hour by default), the route becomes valid again.
- When the penalty drops below one-half the reuse limit, all flap history and penalty are forgotten.

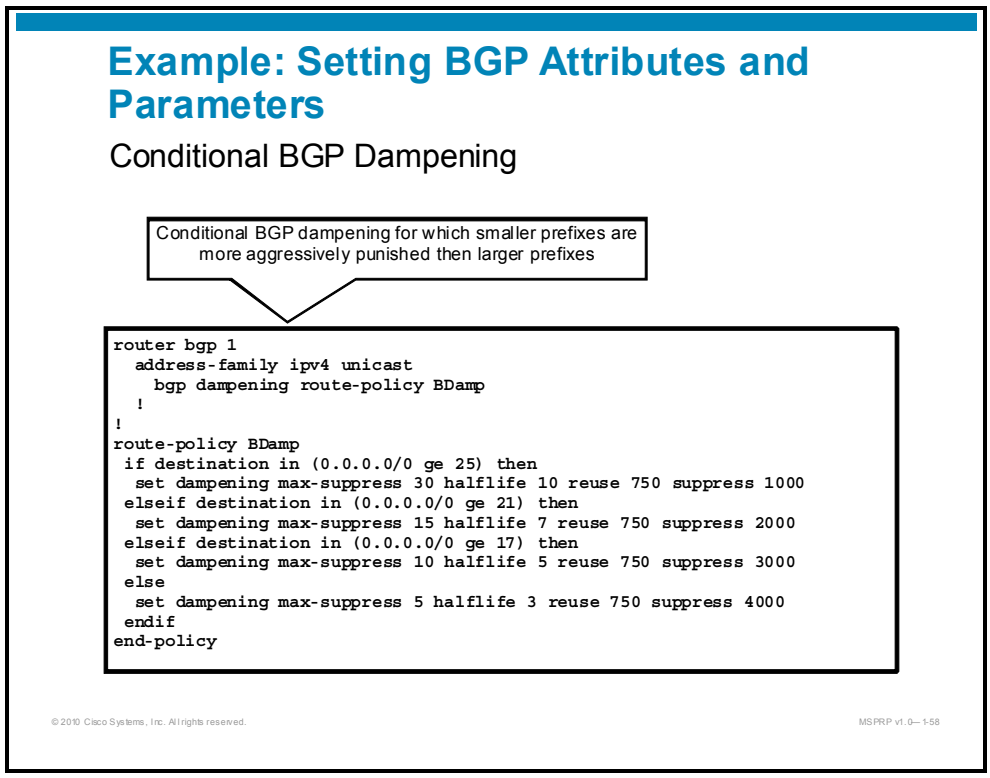

The sample configuration illustrates how graded BGP Route Flap Dampening is configured:

- Small prefixes (/25 to /32) are assumed to be more likely to flap and are hence more aggressively punished if they flap several times.
- Larger prefixes (/21 to /24) are assumed to be slightly more stable and are less aggressively punished (allow more flaps before suppression and become unsuppressed faster when they stabilize).
- Large prefixes (/17 to /20) are even less aggressively punished in case they flap.
- The largest prefixes (/0 to /16) are assumed to be the most stable (large summaries belonging to service providers). They are suppressed after more than four consecutive flaps and are unsuppressed within 10 minutes after stabilizing.

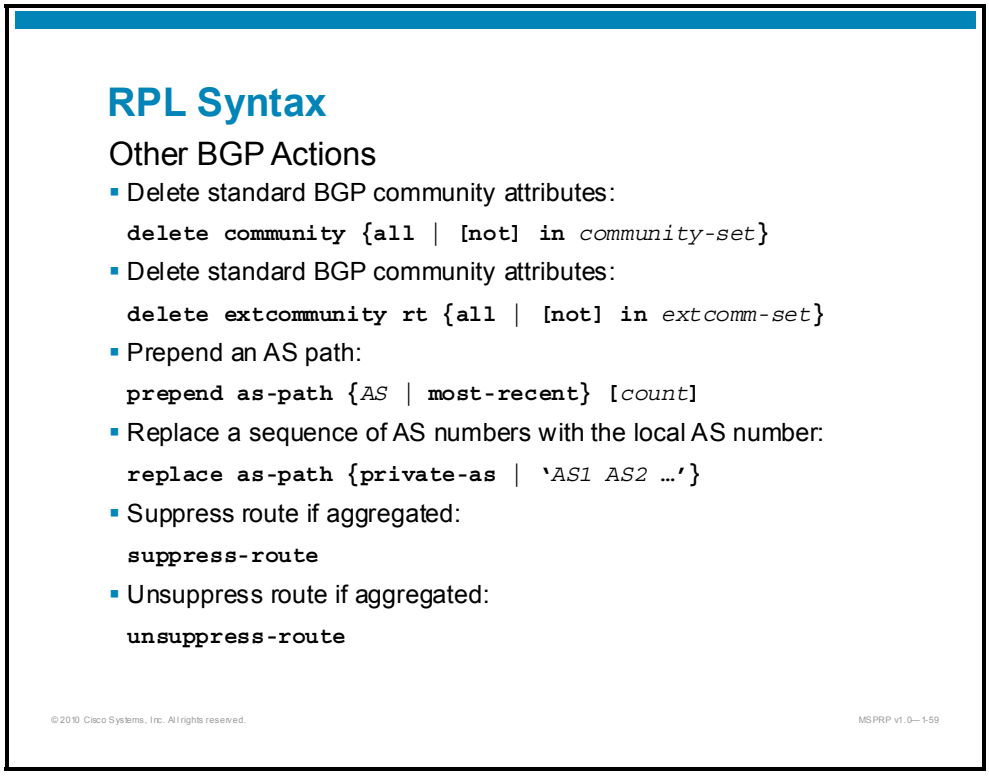

The **delete** command can be used in combination with standard and extended BGP Communities to delete some or all of the BGP Community attributes.

The **prepend as-path** command can be used to prepend an arbitrary number to the AS path a number of times.

The **replace as-path** command can be used to replace all occurrences of private AS numbers with the local AS number or to arbitrarily replace specified AS numbers with the local AS number.

Policies can be used in combination with summarization (aggregation) in order to set various parameters to the summary but also to specify which individual routes are suppressed or unsuppressed.

RPL Syntax

Setting OSPF and IS-IS Parameters

- OSPF metric type:

```
set metric-type {type-1 | type-2]
```

- OSPF metric:

```
set ospf-metric value
```

- IS-IS metric type:

```
set metric-type {external | internal}
```

- IS-IS metric type:

```
set isis-metric value
```

- IS-IS level for redistributed routes:

```
set level {level-1 | level-2 | level-1-2}
```


Routing policies can also be used in combination with OSPF and IS-IS to modify the routing information.

RPL Syntax

Parameterization

RPL supports two types of parameters:

- Global parameters:
  - Defined globally using the **policy-global** command
  - Available for use in all routing policies
- Parameters passed to a nested routing policy:
  - Defined when creating a routing policy
  - Available in match and set statements within a policy or when calling another nested routing policy.

In order to make policies modular and reusable, parameters can be used instead of fixed values when calling nested policies.

A policy can reference global parameters or parameters that are passed to it from a calling policy.

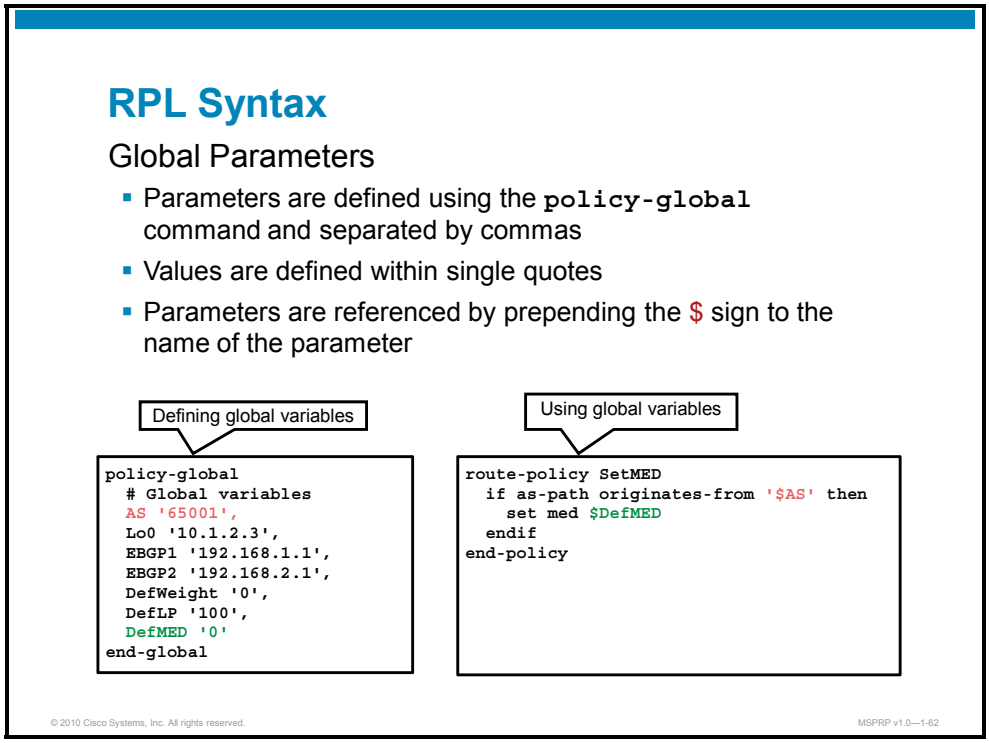

The left example illustrates the usage of the **policy-global** commands, for which all the global variables should be defined. These variables can then be referenced by any routing policy.

The right example illustrates a routing policy referencing two global variables.

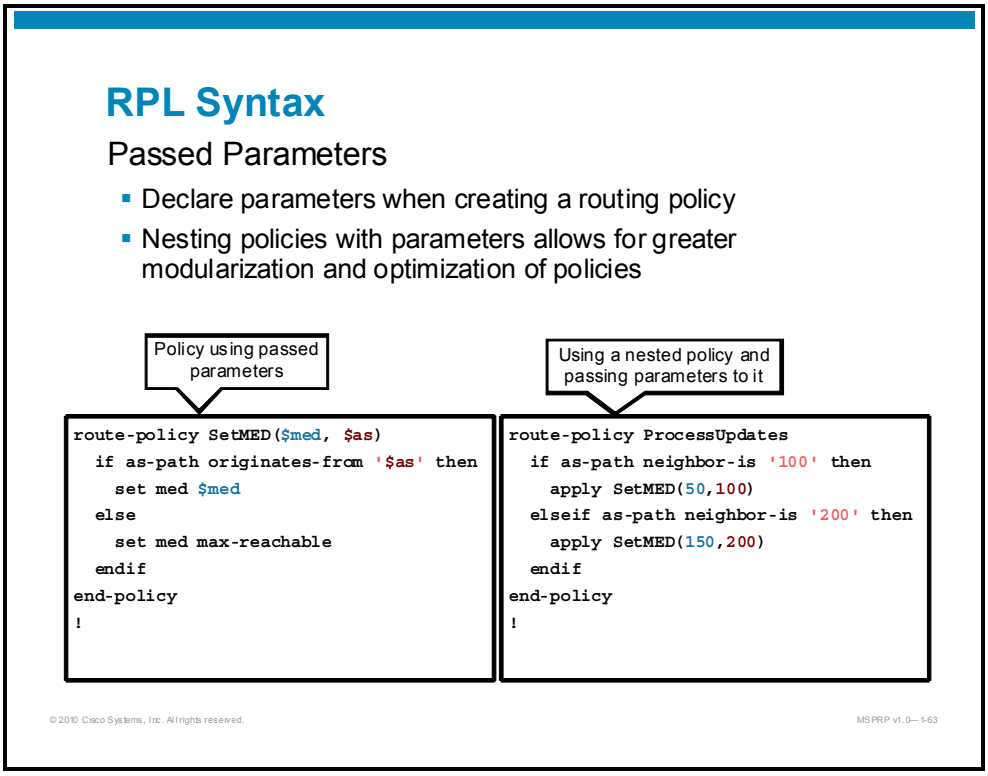

The sample configuration illustrates a modular approach to creating routing policies. The policy on the right calls the policy on the left. The left policy applies a MED value based on the AS from which the route originates. Note that matching based on the AS path is always done using regular expressions, which must be enclosed within single quotes.

The left routing policy is defined with two variables: “$med” and “$as.” When calling this policy from within another policy using the **apply** command, you should supply two parameters.

Applying Routing Policies

- Design a routing policy
- Configure the policy
- Test the policy by using **show** commands
- Apply the policy, if correct
- Routing policies can be used in many places (attach points):
  - Routing updates (e.g. BGP, OSPF, EIGRP, IS-IS, RIP)
  - Route origination (e.g. redistribution, network commands)
  - Route insertion into routing table
  - In **show** commands to filter output

When building a routing policy, you should clearly define the requirements for the route-policy. Often a route-policy will be derived from existing routers that use Cisco IOS Software. In this case the route-policy requires a route map to be “translated” to RPL. You should review the route-policy and optimize it to simplify maintenance of it.

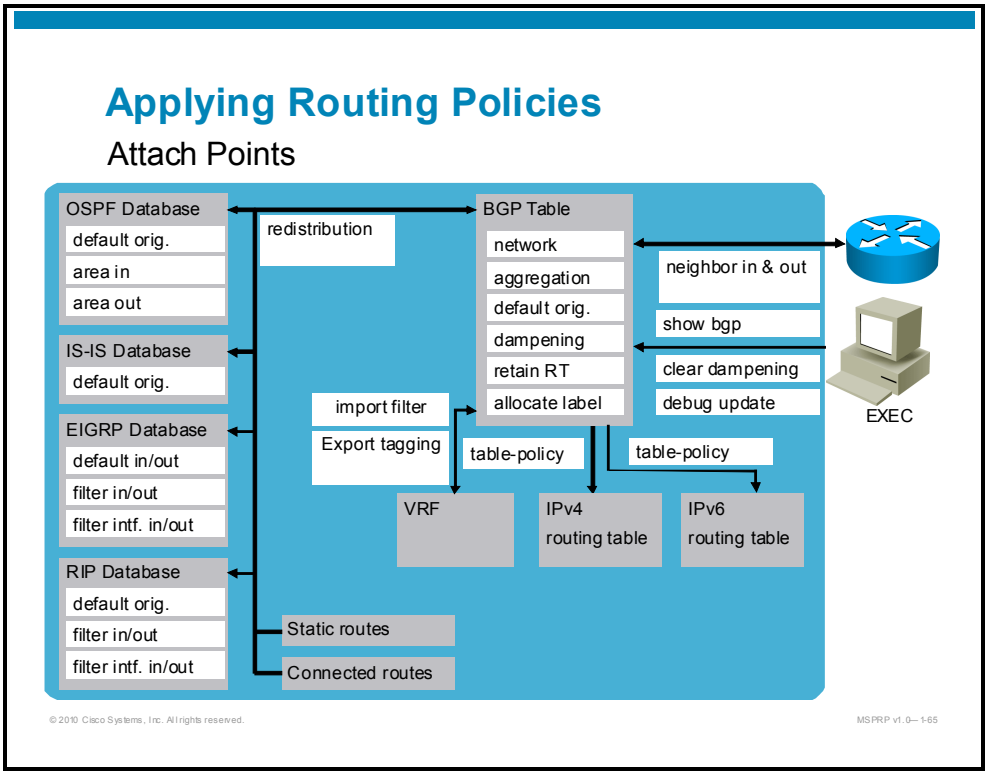

The figure illustrates the many attach points for routing policies:

- Redistribution between any pair of routing protocols
- Received or sent updates depending on the limitations of routing protocols (for example, ABRs in OSPF)
- Origination of routes in BGP by using network statements or summarization
- Injection of routes into the routing table from BGP
- **show** commands in BGP to filter the output or test the effect of the routing policy

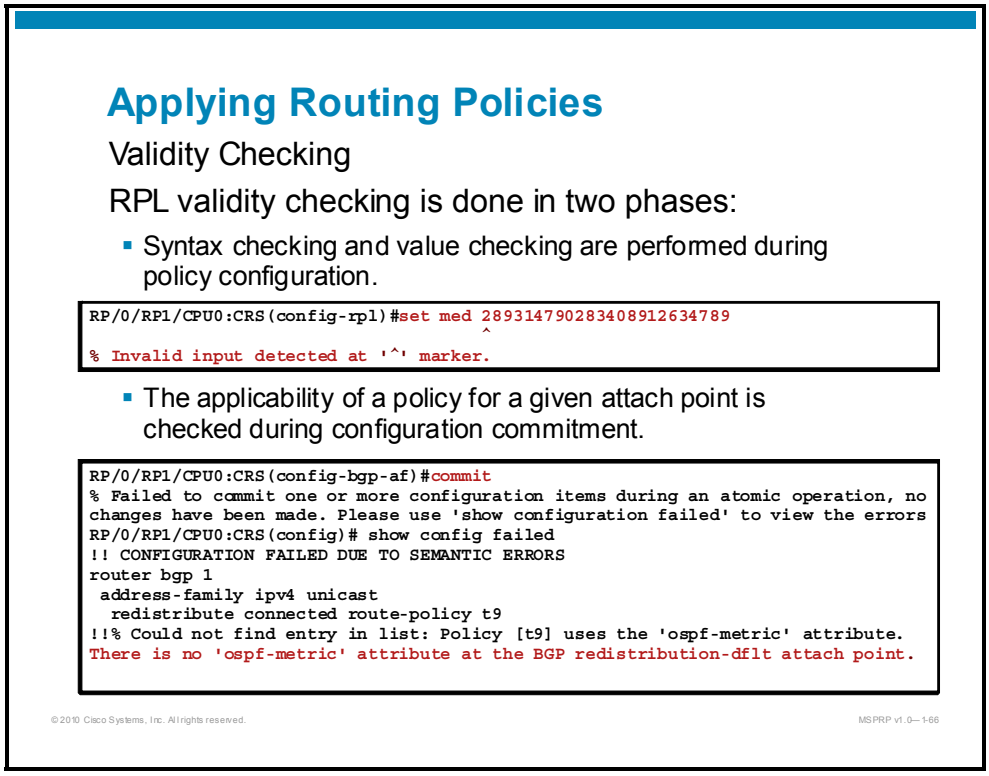

Cisco IOS XR Software performs validity checks in two phases:

- Basic syntax and value checking are performed when a command is entered. The first example illustrates the configuration of MED with a value that is out of range. The syntax checking will reject this command immediately.
- The applicability of a routing policy is verified for a given attach point when the configuration is committed. The second example illustrates that the configuration of a route policy was successful, and the policy was successfully applied to redistribution into BGP. However, when you try to commit the configuration, the router will reject the configuration because you tried to set OSPF parameters inside BGP.

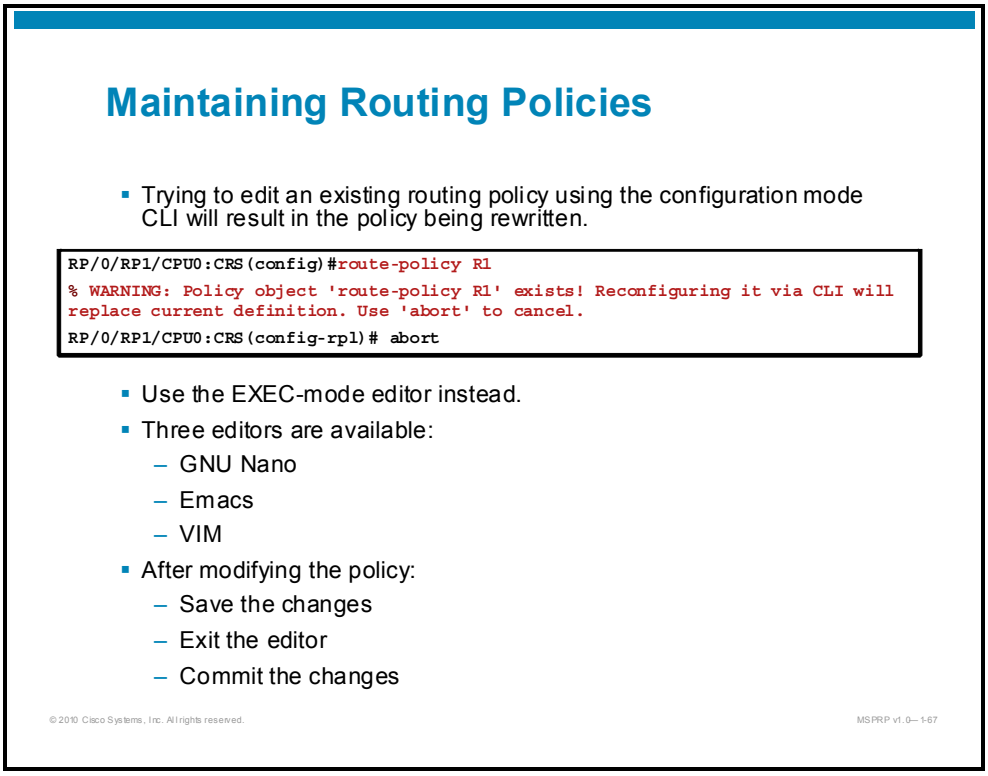

To edit a routing policy, you must use one of three available editors. Using the configuration mode approach will result in the policy being rewritten.

Cisco IOS XR Software comes with three types of editors that are accessible through EXEC mode:

- GNU Nano (the default editor since Cisco IOS XR Release 3.6)
- Micro Emacs
- VIM

Upon exiting from the editor, you will be prompted to save and commit the changes.

The example shows the warning that is displayed if you try to go into policy configuration mode for an already configured policy. Doing so will result in the entire policy being overwritten by the new configuration.

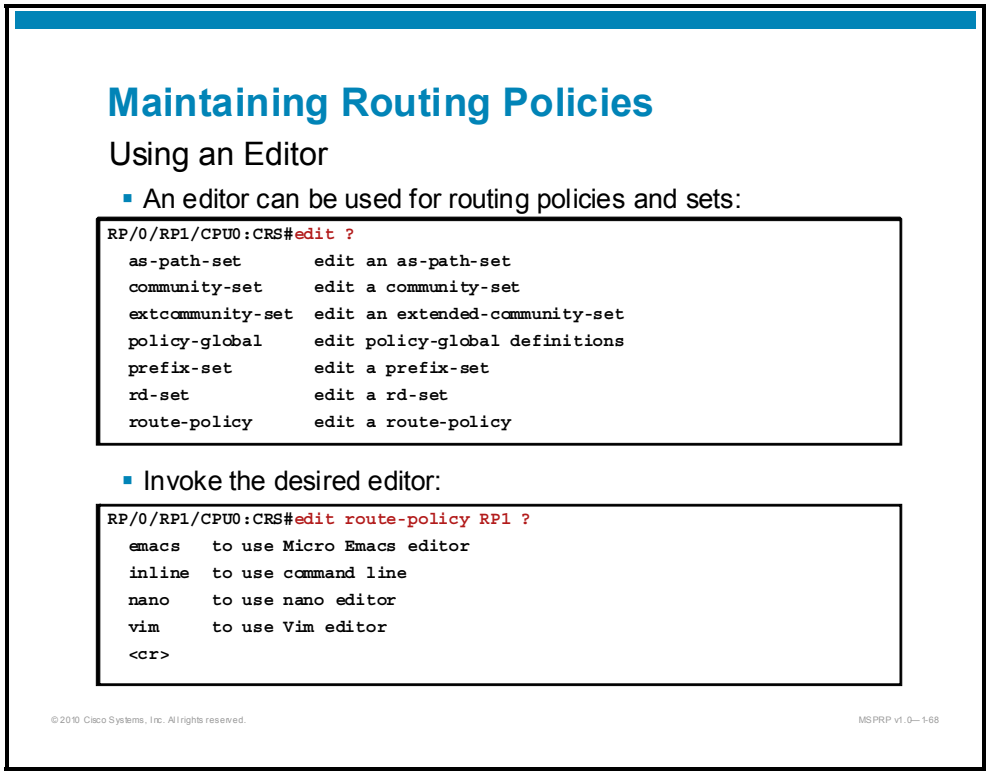

Use the **edit** command in EXEC mode to start editing a configuration object. Select the preferred editor.

The built-in editors are available for route policies and other objects that are covered later in this lesson, such as various sets that are used in combination with route policies.

The following list contains some of the most commonly used commands within the Emacs editor:

- Ctrl-F – move cursor forward (right)
- Ctrl-B – move cursor backward (left)
- Ctrl-N – move cursor to next line (down)
- Ctrl-P – move cursor to previous line (up)
- Ctrl-E – move to the end of the line
- Ctrl-A – move to the start of line
- Backspace – delete character to the left of the cursor
- Ctrl-D – delete character to the right of the cursor
- Ctrl-X followed by Ctrl-S – save changes
- Ctrl-X followed by Ctrl-C – exit and commit saved changes

The following list contains some of the most commonly used commands within the VIM editor:

- ⇦,⇩,⇧,⇨ – move cursor left, down, up, right
- h,j,k,l – move cursor left, down, up, right
- i – start editing at the cursor position
- a – start editing after the cursor position

- Esc – stop editing (return to command mode)
- x – delete character at cursor position
- dd – delete line
- u – undo single action
- Esc followed by :w – save changes
- Esc followed by :q – exit and commit saved changes

After exiting the editor, you will be asked to save and commit the changes:

```
Proceed with commit (yes/no/cancel)? [cancel]: yes
```

Refer to the Cisco IOS XR Software command reference for a detailed list of all commands and options for all the available editors.

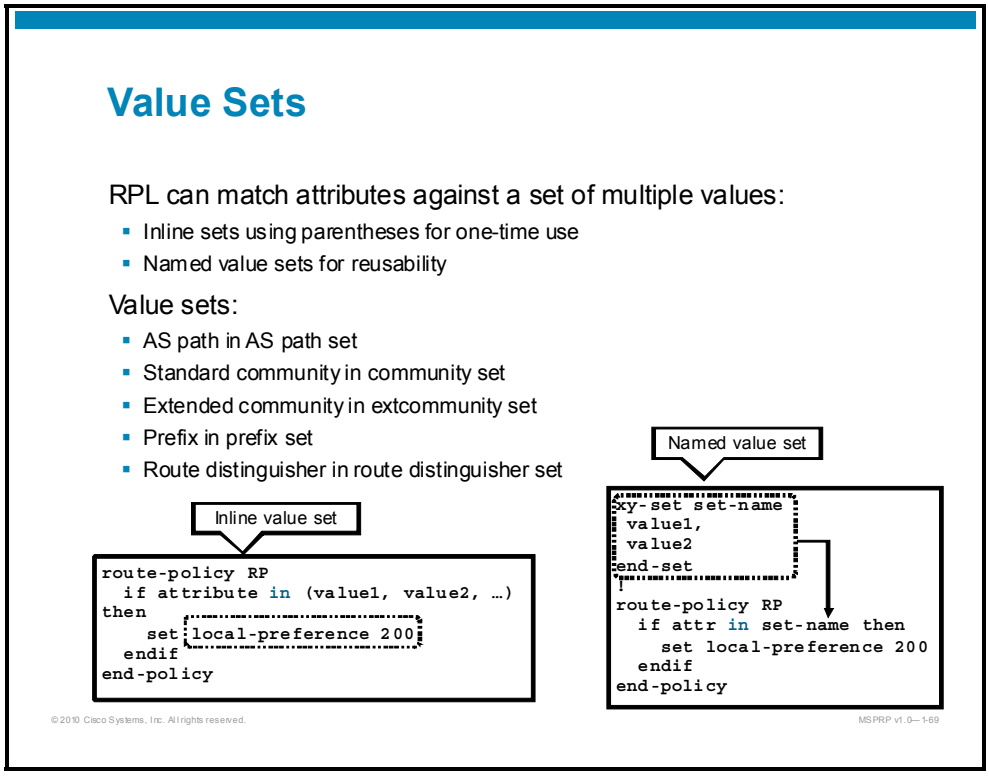

A value set is another object that is used to modularize routing policies. Various types of sets exist for different types of parameters and attributes.

Each set can contain multiple values. The **in** operator can be used for the existence of a value in the set.

The example on the left illustrates a generic condition in which an inline value set is used. The example on the right references a preconfigured value set.

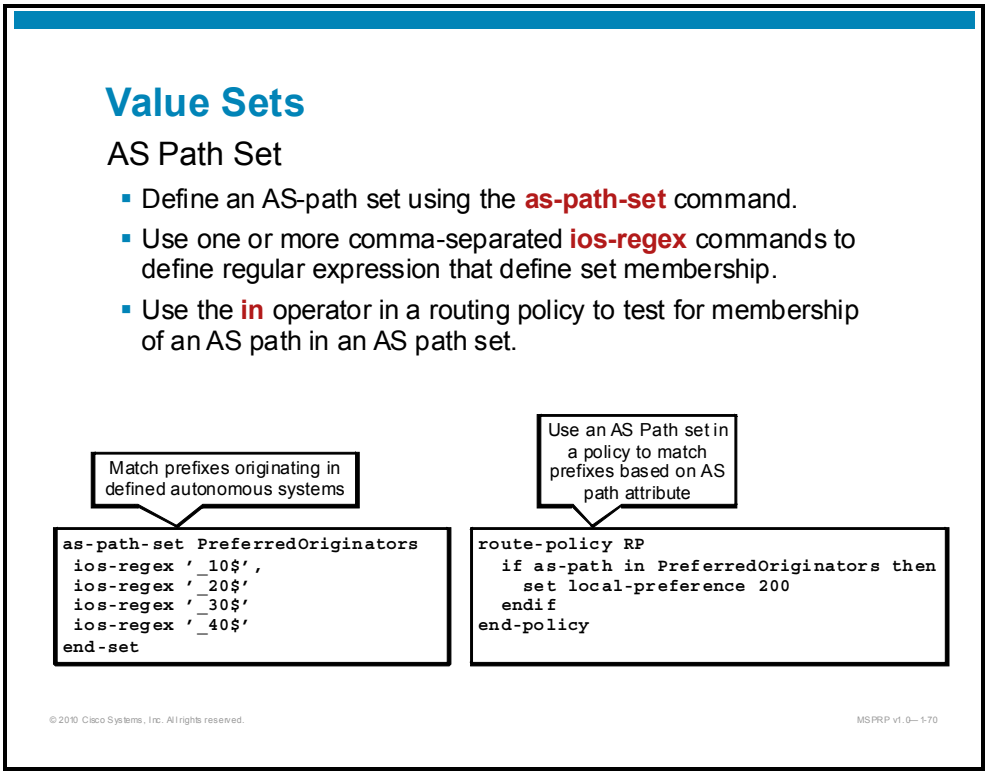

An AS path set can contain one or more regular expressions. A condition can be used to check for an AS path attribute against the set of regular expressions.

The sample configuration uses a policy to set the local preference to 200 for all preferred originating autonomous systems that are listed in the AS path set.

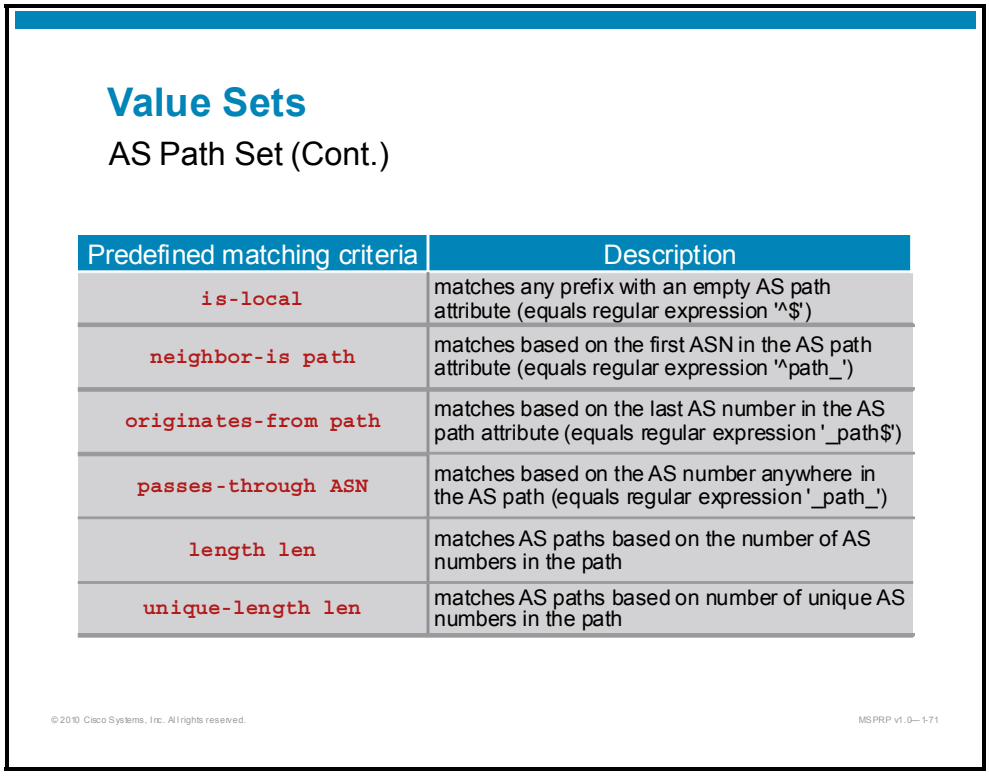

Instead of using regular expressions, you can perform some of the more common AS path checks using built-in conditions:

- "**is-local**" identifies if a prefix is local to the AS; it performs the same function as a regular expression checking for an empty AS path attribute ("^$").
- "**neighbor-is** *path*" identifies if a prefix was received from a neighboring AS; it is equivalent to the regular expression "^path_".
- "**originates-from** *path*" identifies if a prefix was originated by a specified AS; it is equivalent to the regular expression "_path$".
- "**passes-through** *ASN*" identifies if a prefix passed through the specified AS; it is equivalent to the regular expression "_*ASN*_").
- "**length** *len*" matches AS paths based on the number of AS numbers in the path.
- "**unique-length** *len*" matches AS paths that are based on the number of unique AS numbers in the path.

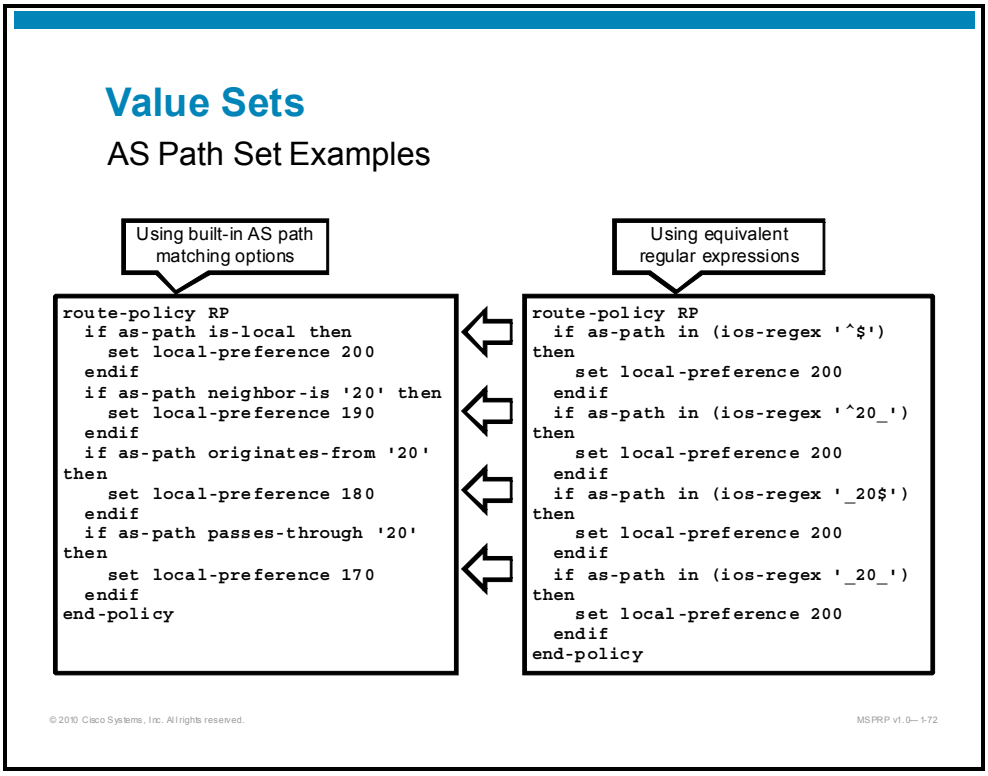

The two samples show configurations equivalent in result but using different approaches:

- The left sample uses built-in conditions.
- The right example uses regular expressions.

In this example:

- The regular expression **'^$'** can be replaced by the built-in operator **is-local**.
- The regular expression **'^20_'** can be replaced by the built-in operator **neighbor-is**.
- The regular expression **'_20$'** can be replaced by the built-in operator **originates-from**.
- The regular expression **'_20_'** can be replaced by the built-in operator **passes-through**.

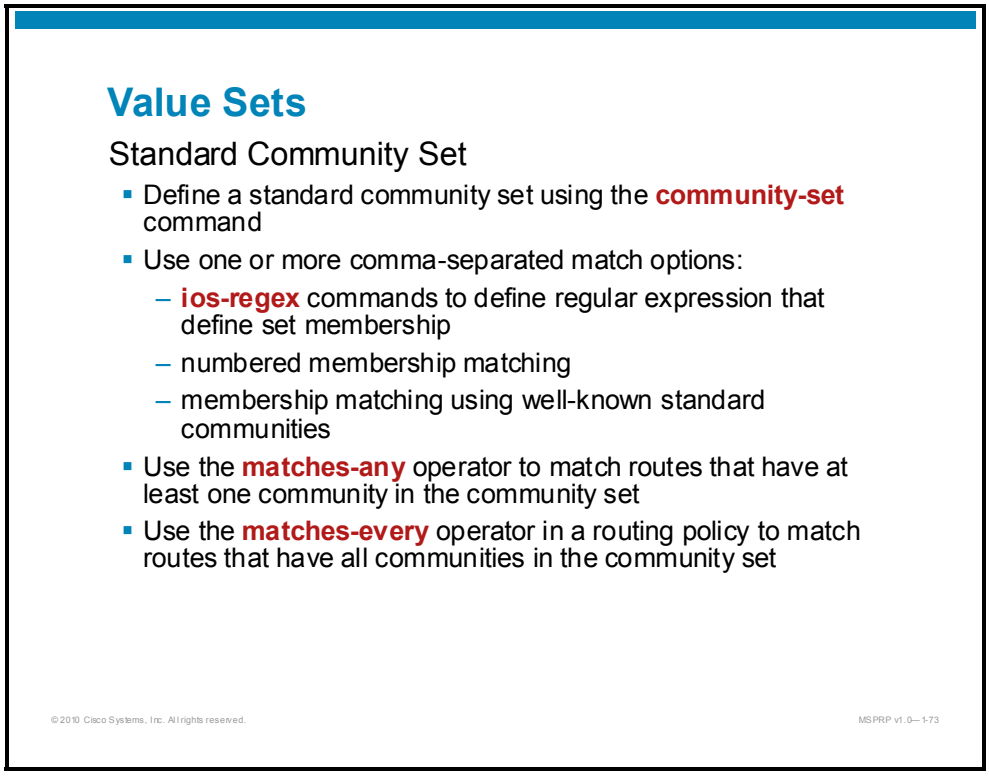

Multiple BGP communities can also be grouped into a community set. The following types of matching can be used with communities:

- **Regular expression matching:** A regular expression is used against an ordered list of communities converted into a string of characters.
- **Numbered matching:** Community attributes are matched against a list of values in a community set.
- **Named matching:** Community attributes are matched against a list of communities including named well-known communities.

Community matching can use modifier that define how the matching is performed:

- The **matches-any** operator should be used to match routes that have at least one community in the community set.
- The **matches-every** operator should be used to match routes that have all communities listed in the community set.

Standard Community Set RegExp Matching

- Use one or more comma-separated **ios-regex** commands to define regular expression that define set membership.

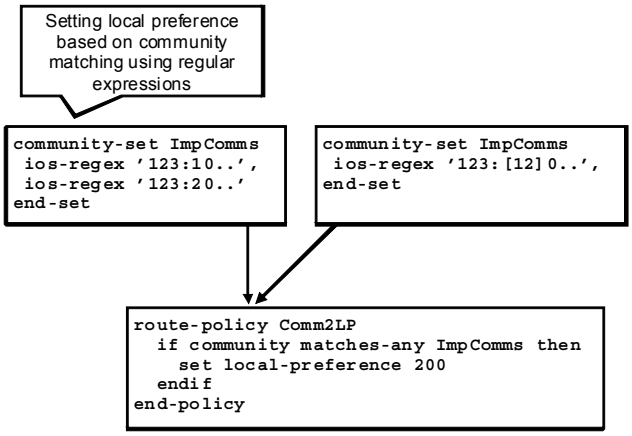

The sample configuration illustrates two community sets based on regular expressions.

The left ImpComms community-set uses two regular expressions. The right ImpComms community-set uses a single regular expression. Either one of them can be used in the Comm2LP route-policy, so that a route will be assigned a local preference of 200 if it contains BGP community 123:10XX or 123:20XX.

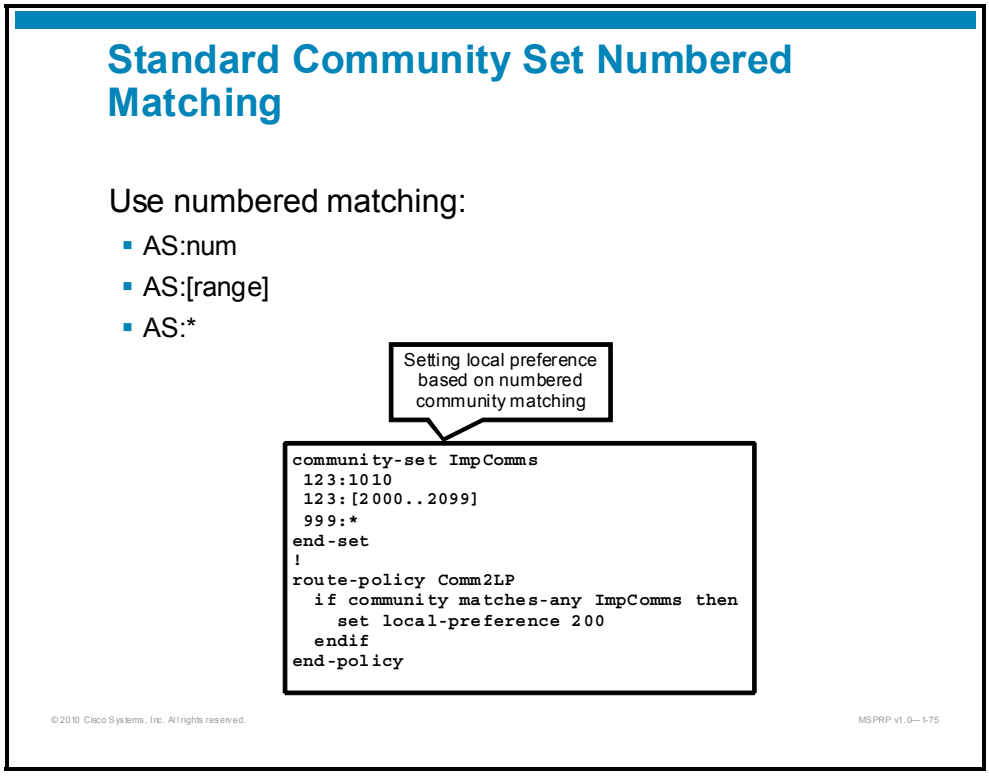

The sample configuration illustrates numbered matching of explicit BGP communities. Additionally, ranges and wildcards can be used in sets:

- "*AS*:*num*" is used to match a specific community.
- "*AS*:[*range*]" is used to match a range of values.
- "*AS*:*" is used to match all values for a give AS.

The sample community set tries to match at least one community from a range defined using different options.

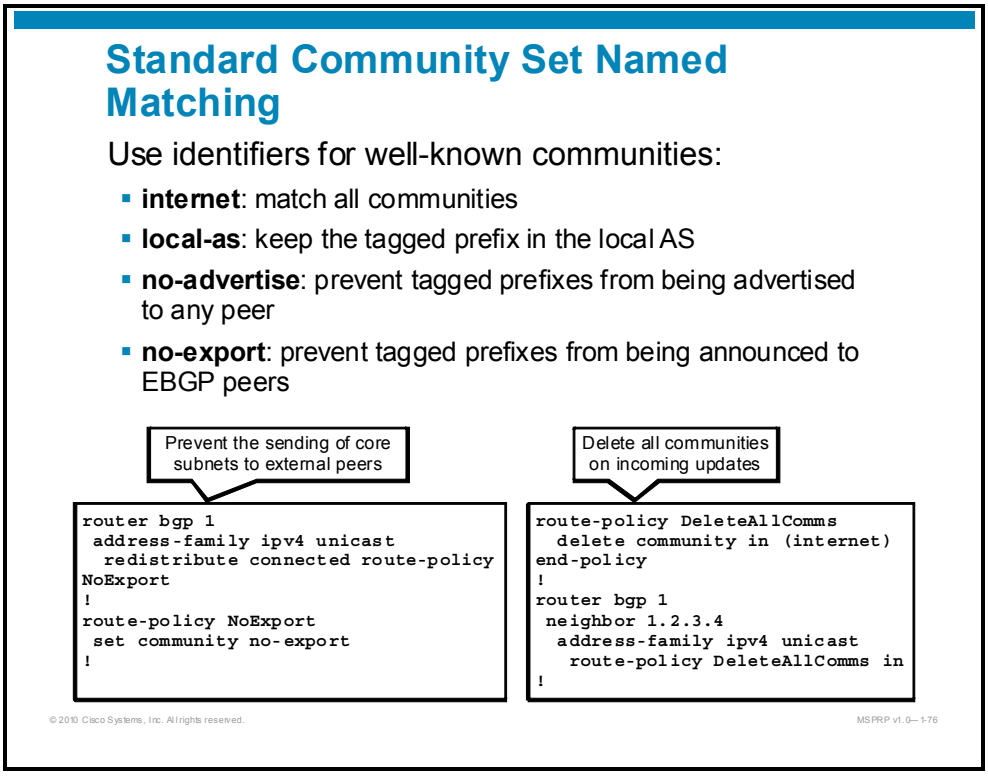

The figure illustrates the third matching option for BGP communities, which is based on the names of well-known communities.

The left example assigns the “no-export” community to all redistributed routes. The right example matches all communities using the **internet** keyword and deletes them.

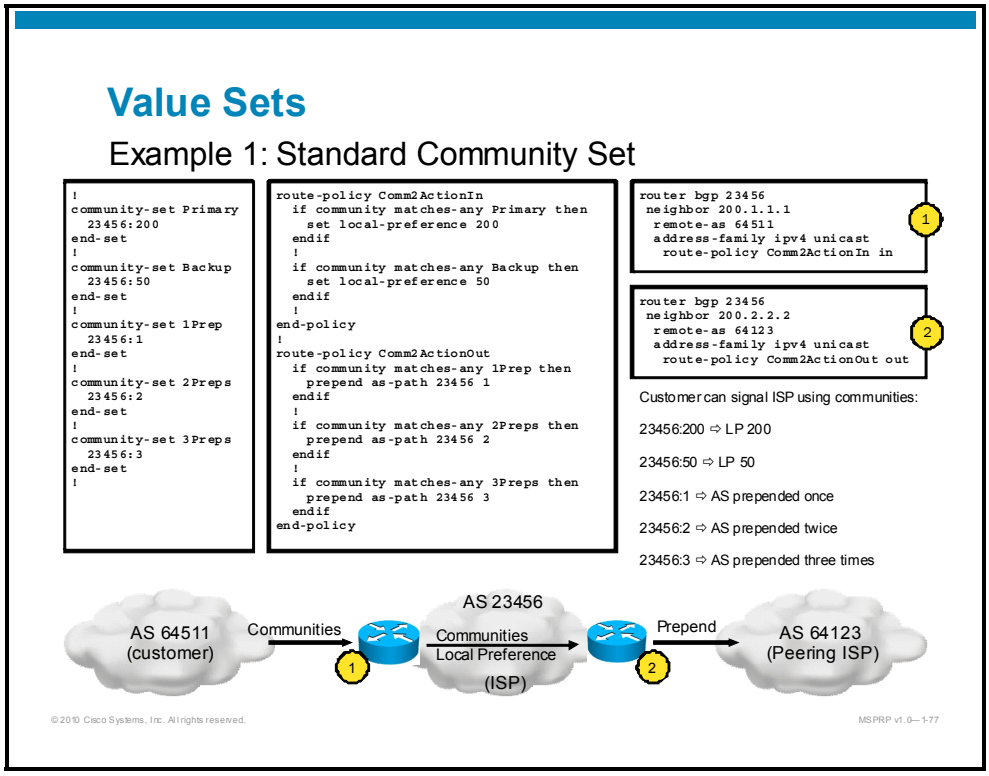

The sample configuration illustrates an AS-wide implementation of a policy:

- Allow customers to signal preference using BGP communities, which are translated into appropriate local preference values.
- Additionally, the egress routers perform prepending on behalf of customers if they have tagged the routes with appropriate BGP communities.

In the example above, five community sets are used to match BGP community attributes coming from external neighbors (for example, a customer). A customer can, for example, signal a desire to use this ISP for its backup connection and may choose to attach two BGP communities to achieve the desired goal: 23456:50 and 23456:3. The route-policy **Comm2ActionIn**, used in the inbound direction on AS edges, will apply an action based on the matched communities. The second **if** statement will match the first community and set the local preference to 50. This will make the first community less desirable compared to other paths with the default local preference of 100. **Comm2ActionOut**, used in the outbound direction on AS edges, will match the second community in the third **if** statement and prepend the AS path attribute three times using its own AS number.

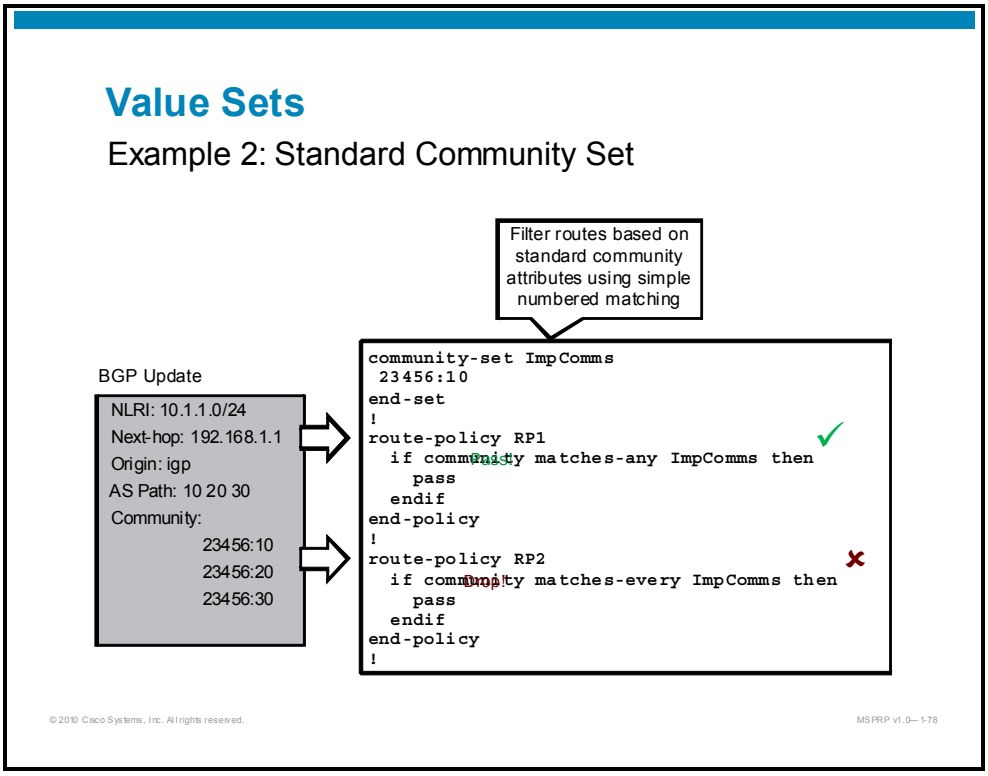

The sample configuration illustrates the difference between the **matches-any** and **matches-every** options:

- RP1: The route with the three community values will match the community set ImpComms, because it contains the 23456:10 community.
- RP2: The route with the three community values will not match the community set ImpComms, because it does not match for two community values (23456:20 and 23456:30).

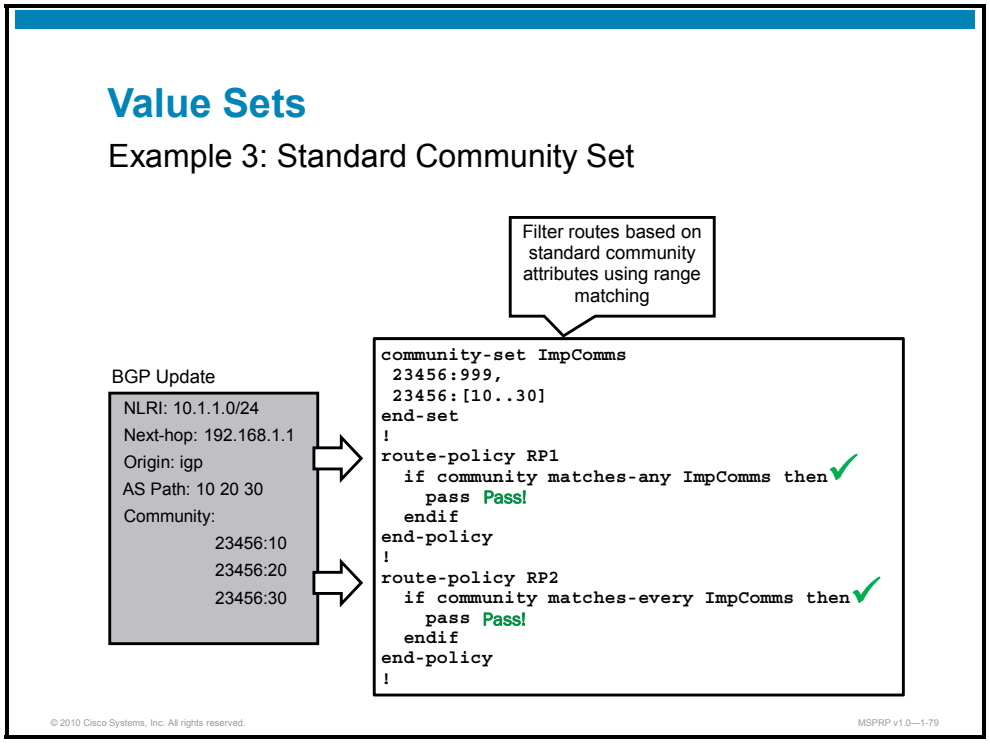

The second example shows that the route with three community values will match in both policies. The reason is that all three community values match the modified community set ImpComms, which now contains range-based matching for 23456:10-23456:30.

The community set in the example uses number-based matching.

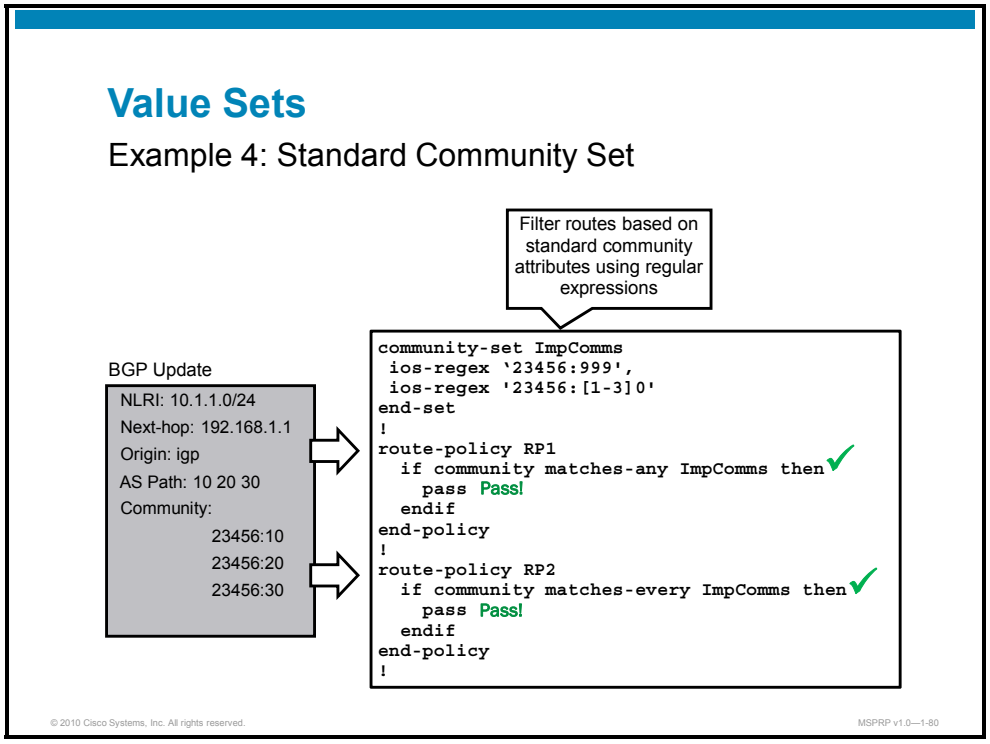

The third example shows the same result as the previous example, except that it uses matching based on regular expressions.

The regular expression 23456:[1-3]0 will match 23456:**1**0, 23456:**2**0, 23456:**3**0

Value Sets

Example 5: Standard Community Set

- Delete all communities on incoming updates that have no meaning in AS 23456.

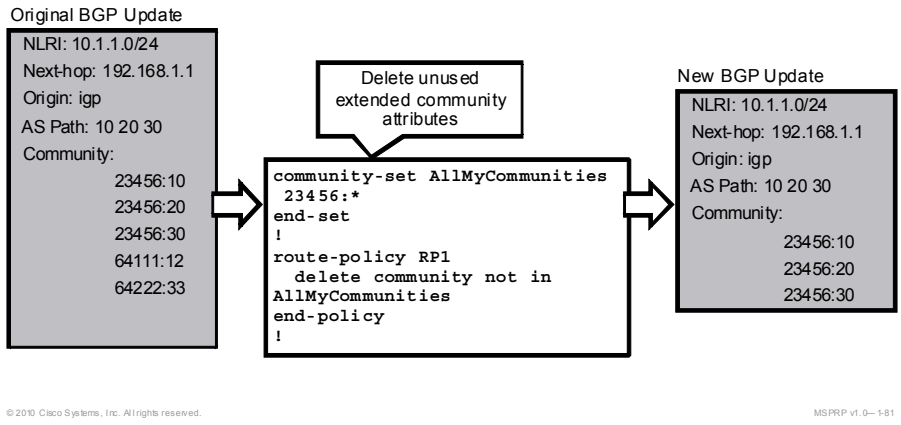

The configuration shows how to delete communities from incoming updates that are outside the desired range (only keep communities that have meaning in the local AS).

This is a common filter that an ISP might use to strip the updates of any BGP communities that have no meaning in its AS. The numbered matching specifies the ISPs AS number 23456 and matches any community value for this AS using the wildcard symbol (“*”). The route policy then deletes all but those communities that are in this range.

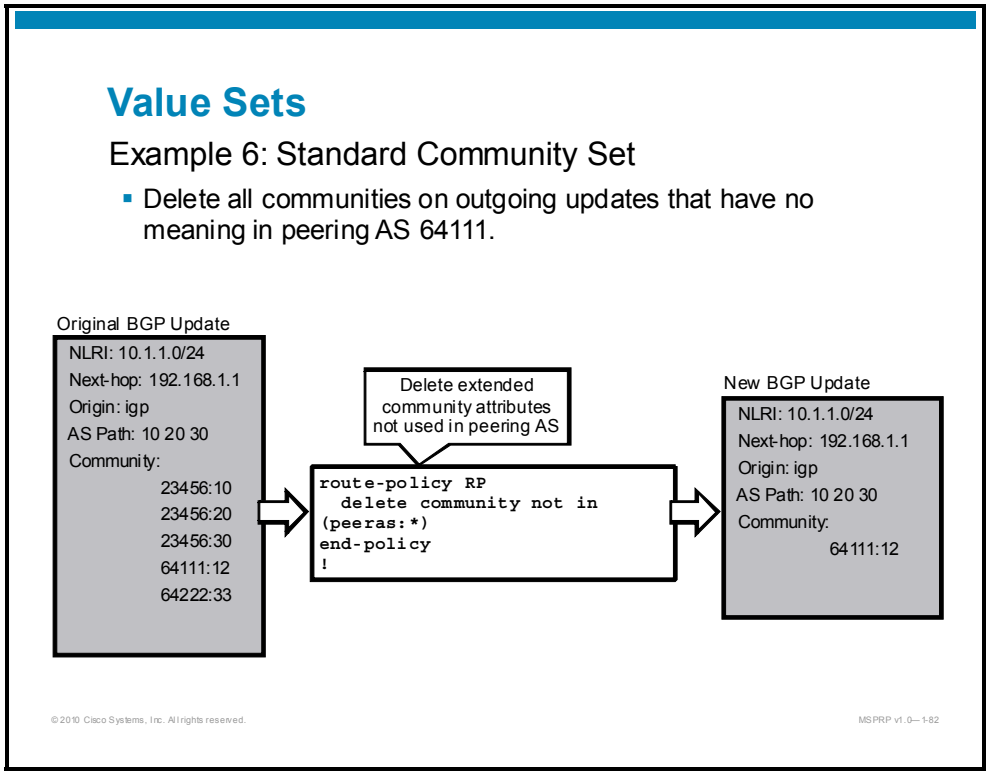

The configuration shows how to delete communities from outgoing updates that are outside the desired range (only keep communities that have meaning in the neighboring AS).

Similarly to the previous example, an ISP can strip out any BGP communities that have no meaning in the neighboring AS. The built-in **peeras** keyword can be used to automatically match the neighbor AS number and the wildcard symbol to match any subsequent value. Instead of using a named community set, the example uses an in-line community set defined within parentheses.

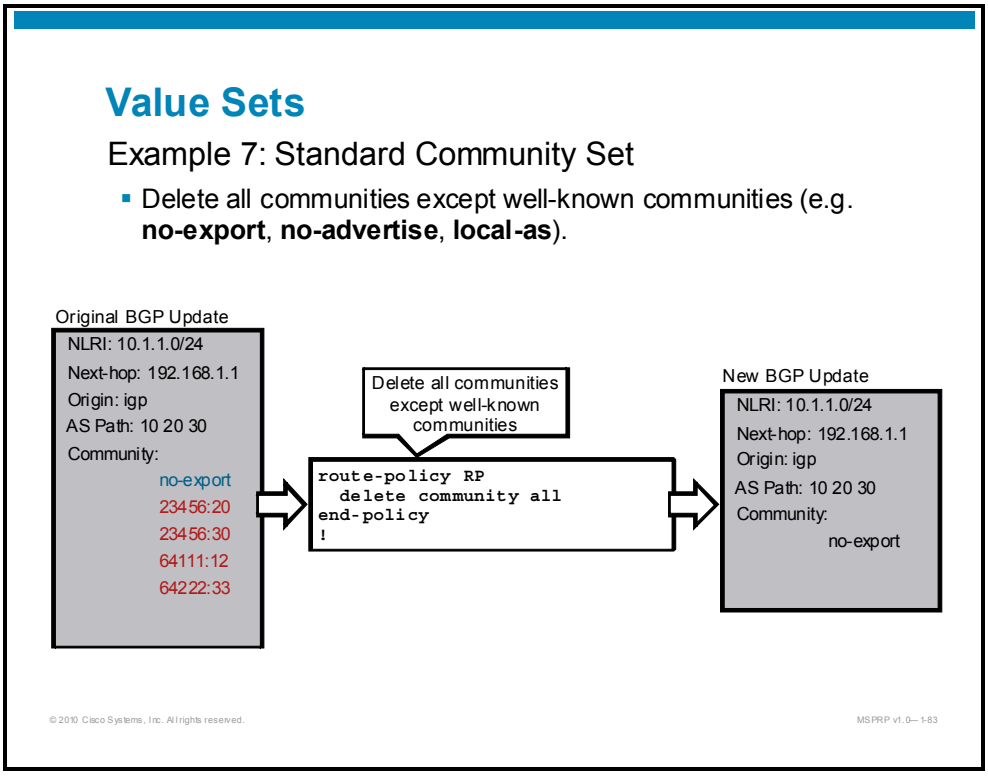

The configuration shows how to delete all communities using the **all** keyword in place of the community set.

Note	This command does not remove the well-known communities (for example, **no-export**), which have predefined actions and must be explicitly deleted if doing so is required.

As shown in the example, all communities except the well-known community **no-export** have been removed from the update.

Value Sets

Prefix Set

- Used to match prefixes in routing protocol updates
  Prefix[/length [{le | ge | eq} mask-len]]

Various prefix sets

```
prefix-set PrivatePrefixes
  10.0.0.0/8 le 32,
  172.16.0.0/12 le 32,
  192.168.0.0/16 le 32
end-set
!
prefix-set CoreLoopbacks
  172.16.1.0/24 eq 32
end-set
!
prefix-set HostRoutes
  0.0.0.0/0 eq 32
end-set
```

Various prefix sets

```
prefix-set DefaultRoute
  0.0.0.0/0
end-set
!
prefix-set AllPrefixes
  0.0.0.0/0 le 32
end-set
!
prefix-set SmallPrefixes
  0.0.0.0/0 ge 24
end-set
!
prefix-set
SmallPrefixesExceptHostRoutes
  0.0.0.0/0 ge 24 le 31
end-set
```


A prefix set is used to match routes based on prefix-list-like criteria in the prefix set. The same syntax is used as with prefix lists.

Monitoring Routing Policies

- Use the **show rpl route-policy** [*policy-name*] **[detail]** commands to display the policies.
- Detailed output also displays all referenced objects (e.g. sets and nested route policies).

Display a policy and all associated objects

```
RP/0/RP1/CPU0:CRS# show rpl route-policy MgmtRTExport detail
extcommunity-set rt MgmtRT
  23456:100,
  23456:200
end-set
!
prefix-set MgmtLoopbacks
  10.1.1.0/24 le 32
end-set
!
route-policy MgmtRTExport
  if destination in MgmtLoopbacks then
    set extcommunity rt MgmtRT
  endif
end-policy
!
```


The **show rpl** command with the **detail** keyword can be used to display the policy configuration including all the dependencies.

In this example, the output shows the configurations of the MgmtRTExport route-policy as well as the configurations of the prefix set and the extended community set referenced within the route policy.

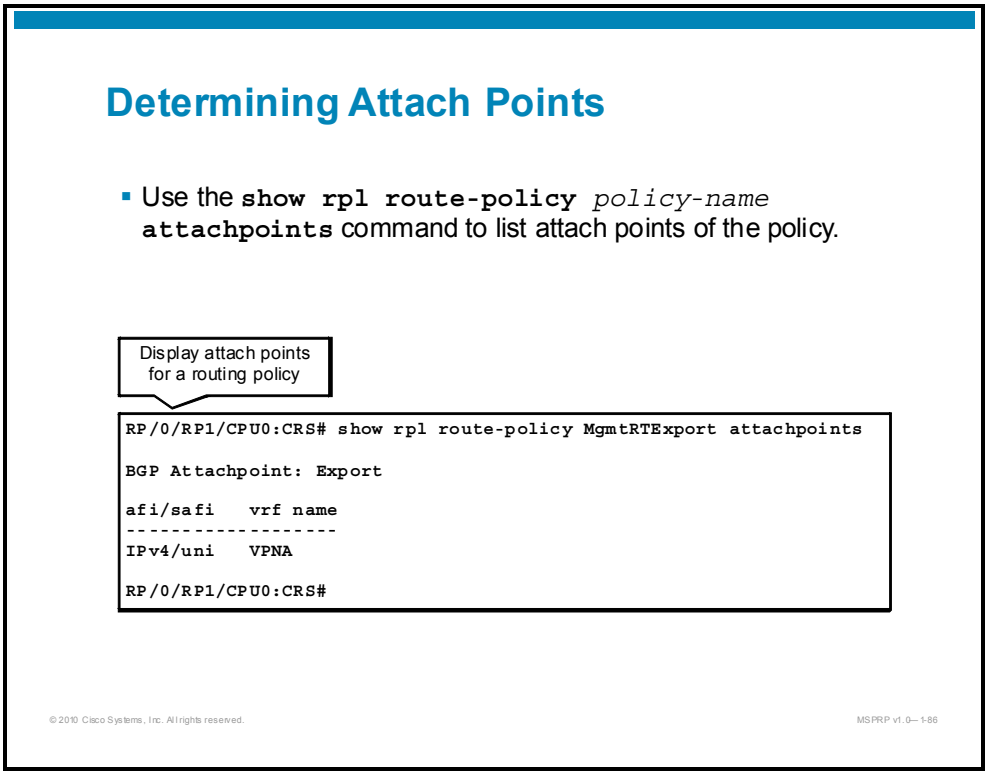

The **attachpoints** option can be used to display all references to the specified policy.

In the example, the **show** command shows that the specified route policy MgmtRTExport is attached to VRF VPNA as an export policy.

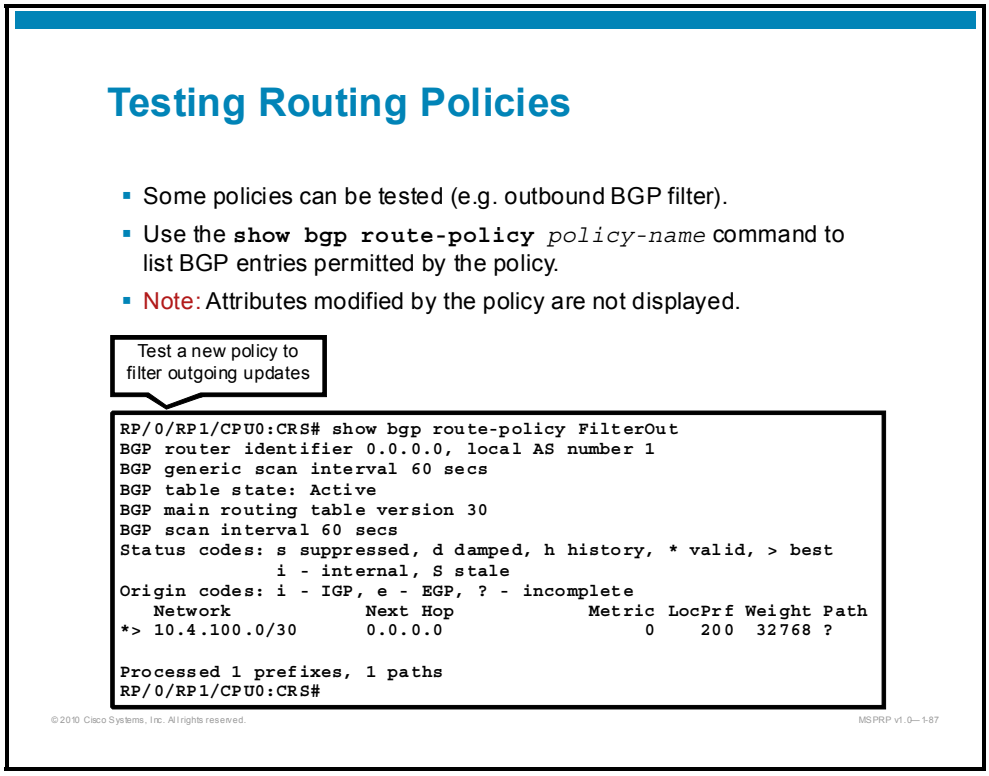

Policies can be combined with the **show bgp** command to display only those BGP entries that are permitted by the policy.

This command can be used to test the performance of a newly configured policy or to limit the display of a large BGP table for troubleshooting purposes.

In the example, the policy FilterOut only displays one entry (10.4.100.0/30) in the BGP table.

Translating Route Maps to Routing Policies

- When migrating from Cisco IOS Software to Cisco IOS XR Software use the following guidelines to translate route maps to policies:
  - Each numbered entry is one if statement.
  - Each match option is one condition:
    - Match conditions of the same type should be joined using the OR logical operator
    - Match conditions of different types should be joined using the AND logical operator
    - Use parentheses to maintain proper precedence

When translating route maps to routing policies it is important to understand the relation between multiple conditions in a single route-map statement. As discussed earlier, multiple conditions of the same type are combined using a logical OR. Hence, you should use the OR operator in the **if** statement of the routing policy. Multiple conditions of different types are combined using a logical AND. Hence, you should use the AND operator in the **if** statement of the routing policy. Make sure that you use parentheses for proper operator precedence.

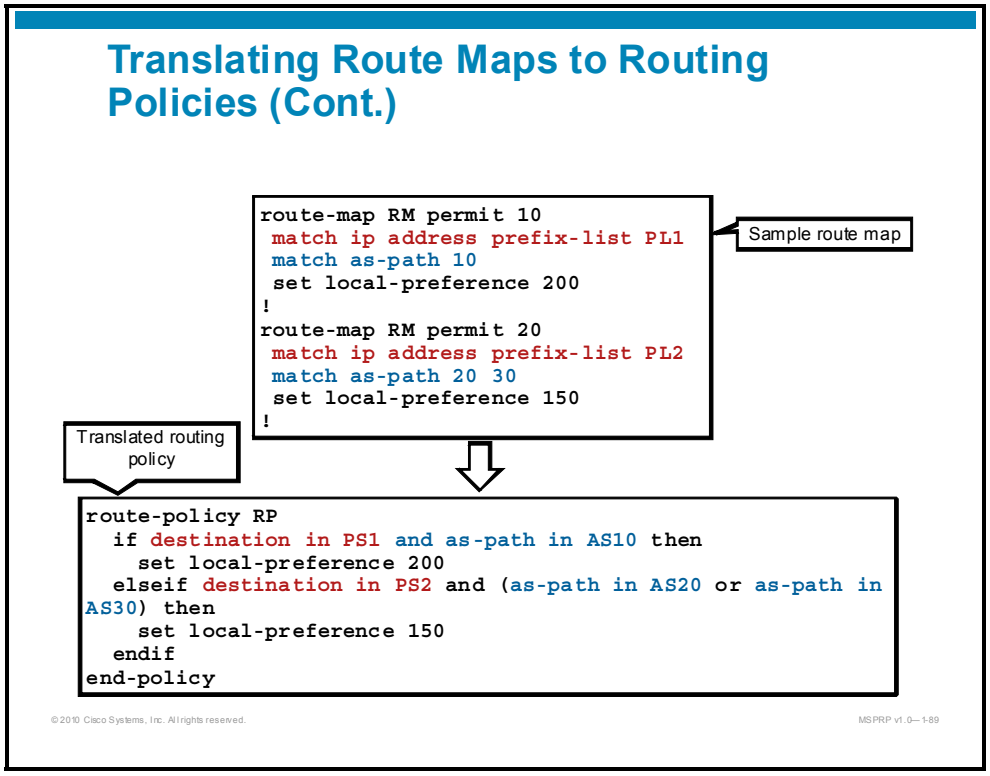

The two sample configurations show how two complex route-map statements can be translated into a routing policy.

Note	For the second route-map statement, you must use parentheses to group the two AS path sets (that is, at least one as-path match is required). You must also combine the route-map statement with a prefix set (PS2) using the AND operator.

Summary

Summary

- Use prefix-based filtering for maximum security.
- Use AS path-based filtering for greater flexibility.
- Use route maps or routing policies to implement complex routing policies.

Lesson 3

Using Management and Monitoring Tools

Overview

This lesson describes the tools that are needed in a service provider environment to successfully manage and monitor network devices and services. Routing protocol performance and issues can be monitored using mechanisms such as Simple Network Management Protocol (SNMP), syslog, and command-line interface (CLI), in combination with various types of central monitoring and provisioning software.

Objectives

Upon completing this lesson, you will be able to identify the main characteristics of routing protocols that are used in service provider environments. This ability includes being able to meet these objectives:

- Describe the characteristics and requirements for management and monitoring tools in service provider environments.
- Describe the characteristics and requirements for event logging using syslog and SNMP traps.
- Describe the characteristics and requirements for using SNMP-based monitoring.
- Describe the characteristics and requirements for using looking glasses.
- Describe the characteristics and requirements for provisioning tools.
- Describe the characteristics and requirements for administration.

Management and Monitoring Overview

This topic provides an overview of management and monitoring tools as they relate to routing protocols.

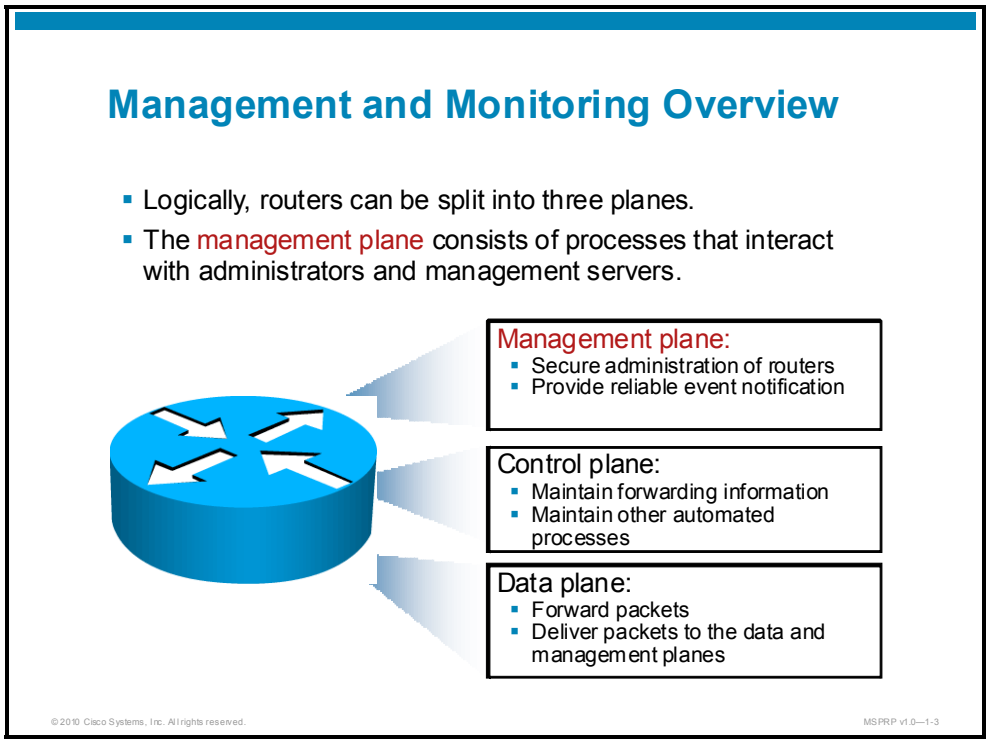

Each network device performs a number of functions, which can be split into three general categories or planes of operation:

- The data plane is the part of the device that is typically implemented in hardware (for example, network interface cards) and is responsible for forwarding of packets or frames.
- The control plane is the part where a collection of protocols ensures that the data plane operates properly. Routing protocols (for example, Open Shortest Path First [OSPF], Intermediate System-to-Intermediate System [IS-IS], Border Gateway Protocol [BGP]), Address Resolution Protocol [ARP], and Network Time Protocol [NTP]) are just some examples of the processes that are typically used on network devices to ensure proper operation.
- The management plane is the part of the device that allows network administrators to interface with the network device directly or indirectly through various types of management servers.

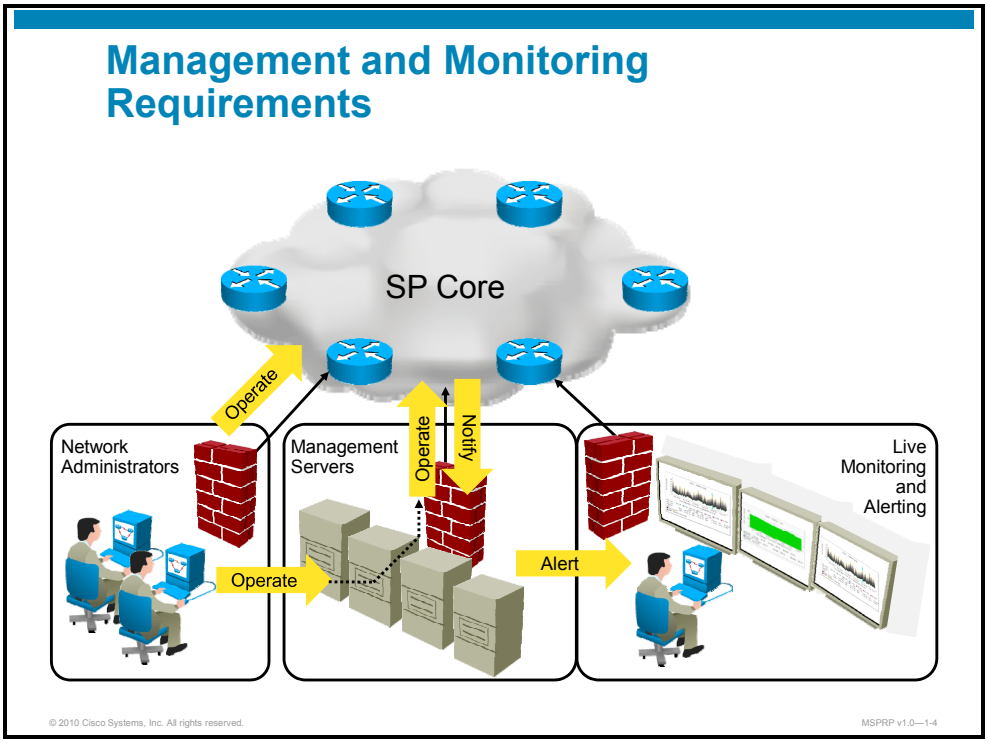

The figure illustrates typical management paths by which network devices interface with network administrators:

- Network administrators can operate devices directly using a CLI or a web-based GUI.
- Network administrators can operate devices indirectly by using an element management system running on a dedicated management server that is isolated and protected from the rest of the network.
- Network devices can notify network administrators and operators of significant events by sending notifications to a dedicated monitoring server.
- Monitoring servers can present significant information and alerts for important events.

Management Tasks

- Operation objectives:
  - Provisioning
  - Optimization
  - Troubleshooting
- Monitoring objectives:
  - Security
  - Availability
  - Performance

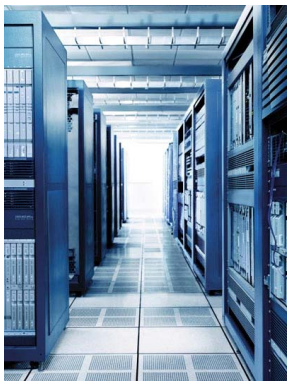

In general, the major requirements of network management include the following:

- **Operation:** managing devices and configuration
- **Monitoring:** processing notifications

Operation can include various tasks, such as provisioning new routers, core links, customer connections, and services. Monitoring is primarily performed for security reasons, to track the availability of resources and help in fault management, or to track the performance of resources (links, CPUs) and help in capacity planning.

Administration of Routing Protocols

- Device administration:
  - Telnet
  - Secure Shell (SSH)
  - HTTP
  - SSL (HTTPS)
- Routing protocols are typically managed through:
  - The command-line interface (CLI) over a secure session (e.g. SSH)
  - Centralized element management systems

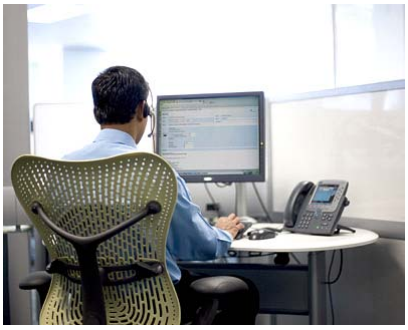

Routing protocols are just one aspect of router configuration and are typically configured in combination with other services. Core routing protocols, such as interior gateway protocol (IGP) and Internal Border Gateway Protocol (IBGP), are often maintained using the CLI. In contrast, edge routing protocols, such as IGP or External Border Gateway Protocol (EBGP) with customers, are often managed through a centralized management system.

Some network devices also support HTTP-based management, although this type of management is not often used in service provider environments. Management systems, on the other hand, will usually utilize a web-based interface.

Routing Protocol Events

Proactive monitoring of routing protocol information:

- Adjacency changes (flaps):
  - More reliable than monitoring link status
  - May result in BGP route flap dampening
  - Sent to a central syslog or SNMP server or both, if configured

```
Edge1#
*Aug 18 15:24:20: %BGP-5-ADJCHANGE: neighbor 10.1.1.1 Down BGP Notification sent
*Aug 18 15:24:20: %BGP-3-NOTIFICATION: sent to neighbor 10.1.1.1 4/0 (hold time expired)
Edge1#
```

- Traffic pattern changes:
  - A flap in BGP can cause traffic to shift from one link to another
  - Traffic statistics collected using SNMP and graphed can be used to detect routing changes

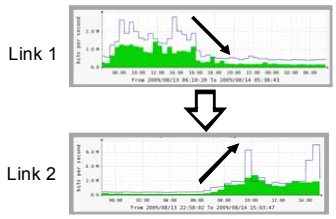

Monitoring of routing protocols is often required to detect failures that would otherwise not be visible. The top example illustrates how a BGP failure is detected even if the underlying physical interface has not failed.

Additionally, link utilization monitoring can be used to detect subtle changes in routing protocol operation. The example shows link-utilization graphs for two links. A change is detected, which indicates that a large amount of traffic has moved from one link to the other. This shift in traffic may be a consequence of a short adjacency flap on Link 1, after which BGP decided to prefer Link 2, because it appeared to be more stable.

Some other examples of routing-protocol monitoring are as follows:

- Monitor the number of routes in complete Internet routing tables coming from upstream and peering service providers
- Monitor CPU utilization for routing processes
- Monitor memory utilization for routing processes

Management and Monitoring Overview Requirements

Enable one or both of the following:

- Syslog logging of adjacency changes in all routing protocols
- SNMP traps for routing protocol events

Use a monitoring system that can send alerts for:

- Significant events (e.g. adjacency changes)
- Traffic pattern changes (e.g. thresholds)

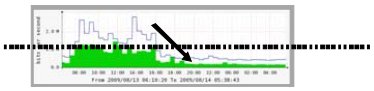

To monitor routing protocols, you use the same two mechanisms as with any other process that needs monitoring:

- Syslog to forward system messages to a central syslog server or pair of servers
- SNMP traps to forward system messages to a central SNMP server or pair of servers

Historically, syslog was used more for security reasons, so it may include many messages that are typically not used with SNMP. In contrast, SNMP was designed to track network performance and utilization.

Today, both protocols can be used to notify a central server of significant network events (for example, a link state change or an adjacency or neighbor state change). However, syslog is still primarily used for security monitoring.

NetFlow can also be used to get a more detailed view of the forwarded traffic and may aid in the monitoring of Internet routing tables. For example, NetFlow supports the mapping of traffic to BGP autonomous system (AS) numbers.

Common Activities in a Service Provider Environment

- Sending syslog messages to a central syslog server to identify security and availability issues
- Sending SNMP traps to identify availability issues
- Using SNMP requests to gather usage and availability statistics
- Using SNMP to manage equipment
- Using the CLI to manage equipment
- Using management (provisioning) servers to manage equipment
- Using a ticketing system for proper process management
- Using and providing BGP looking glass services to ease troubleshooting

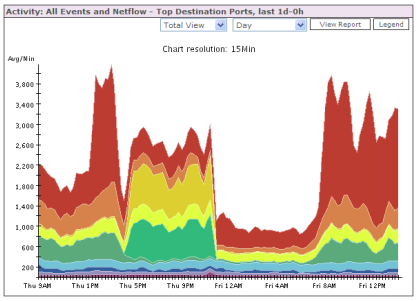

Common activities in a service provider environment are as follows:

- Sending syslog messages to a central syslog server to identify security and availability issues (for example, using the Cisco Security Monitoring, Analysis, and Response System (Cisco Security MARS) appliance, as shown in this screenshot)
- Sending SNMP traps to identify availability and performance issues
- Using SNMP requests to gather usage and availability statistics
- Using SNMP to manage equipment
- Using the CLI to manage equipment
- Using management (provisioning) servers to manage equipment
- Using a ticketing system for proper process management (change and fault management)
- Using and providing BGP looking glass services to ease troubleshooting
- Using NetFlow

Event Logging Using Syslog and SNMP

This topic describes the usage of syslog and SNMP to aid in monitoring of routing protocols.

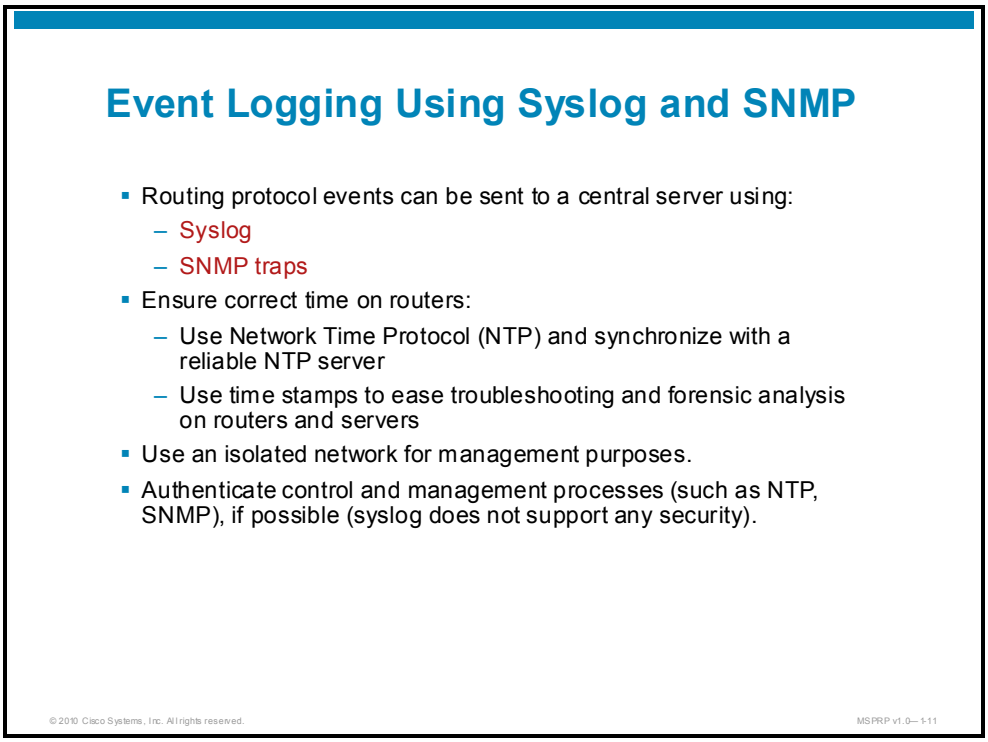

To enable event logging, you need to decide which protocols you will use (which may depend on the monitoring systems you use). Cisco Security MARS, for example, is used for security monitoring and can accept syslog messages, SNMP traps, and NetFlow messages. Tools used to monitor availability, performance, and resource utilization can only use SNMP. To maintain accurate logs, it is recommended that you have accurate time on the servers and network devices. Notifications should be time-stamped at least by the server and optionally also by network devices. NTP is used to synchronize time on all devices—network devices and servers.

To minimize the threat to the network, all management system control and communications should be performed through an isolated environment, such as a VLAN, Multiprotocol Label Switching (MPLS), or a virtual private network (VPN).

Example: Event Logging Using Syslog and SNMP

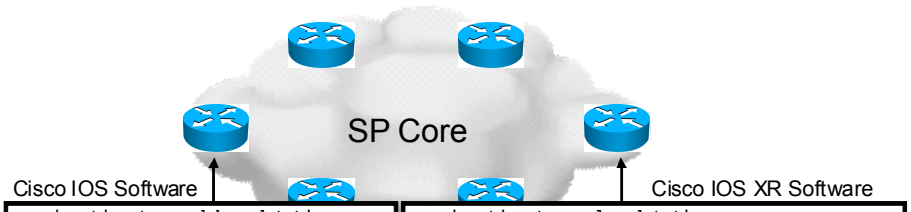


```
service timestamps debug datetime year
service timestamps log datetime year
!
logging host 10.2.2.2
logging trap debugging
!
snmp-server enable traps bgp
snmp-server enable traps ospf
snmp-server enable traps syslog
!
ntp server 172.16.110.1 prefer
ntp server 172.16.110.2
!
router ospf 1
 log-adjacency-changes
!
router bgp 23456
 bgp log-neighbor-changes
!
```

```
service timestamps log datetime msec
service timestamps debug datetime msec
!
logging host 10.2.2.2
logging trap debugging
!
snmp-server traps ospf state-change neighbor-state
snmp-server traps bgp
!
ntp server 172.16.110.1 prefer
ntp server 172.16.110.2
!
router ospf 1
 log adjacency changes detail
!
router bgp 23456
 bgp log neighbor changes
!
```


The two sample configurations show how to configure syslog and SNMP traps for event notification that is related to routing protocols in Cisco IOS and Cisco IOS XR Software.

The first portion of the configuration ensures that all syslog and locally logged messages are time-stamped to help in troubleshooting and forensic analysis when timing information may be required.

The second part of the configuration shows the syslog logging configuration that is configured to send the maximum amount of information (that is, debugging level).

The third part shows the configuration of SNMP to include BGP and OSPF-related events.

The fourth part shows how routing protocol adjacency changes are logged both for OSPF and BGP.

Note that these sample configurations only show the relevant partial configuration of syslog and SNMP.

Availability and Utilization Monitoring Using SNMP

This topic describes the usage of SNMP as it relates to routing protocols.

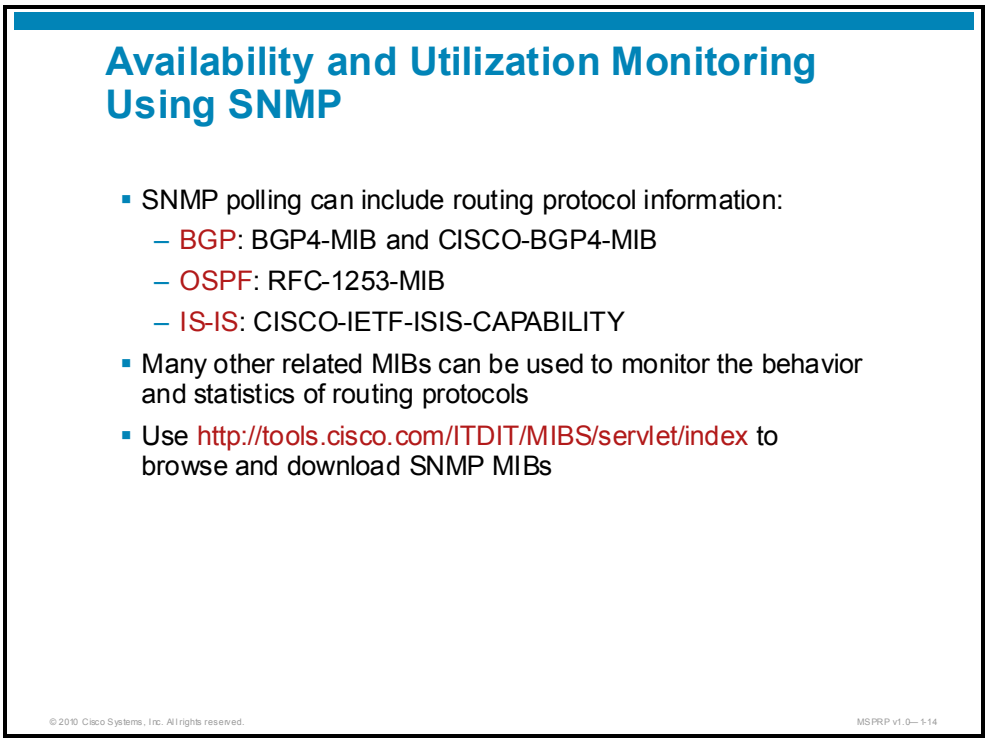

SNMP can also be used to poll network devices for information. Standard and proprietary Cisco SNMP MIBs exist that help in gathering statistics. Refer to the link to get up-to-date information about the available MIBs and support for the MIBs on specific network devices.

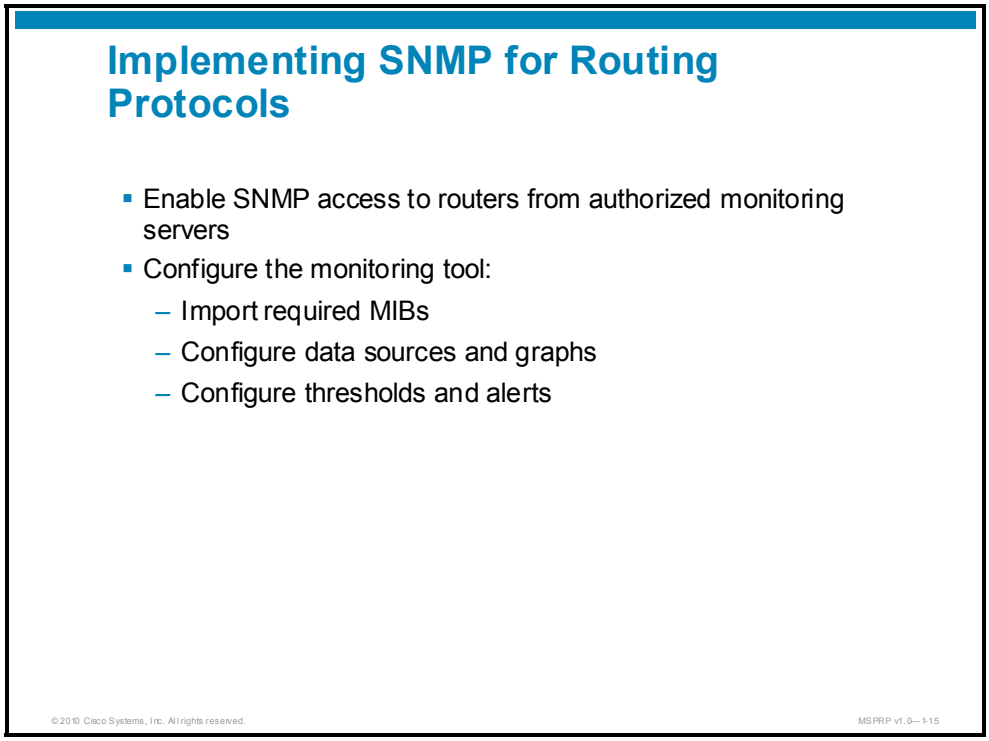

In order to interconnect network devices (routers) and monitoring servers, you need to enable SNMP (for example, read community, access list permitting authorized servers access to routers) on routers. On the monitoring system, you should import the required MIBs, if they are not already there, and start creating data sources and graphs for the required routing protocol monitoring. Thresholds, if supported, can be configured on the monitoring system to alert network operators of any significant events to improve reaction times upon failures in the network.

BGP Looking Glasses

This topic describes the BGP looking glass tools, which are useful in troubleshooting of BGP route propagation and Internet connectivity.

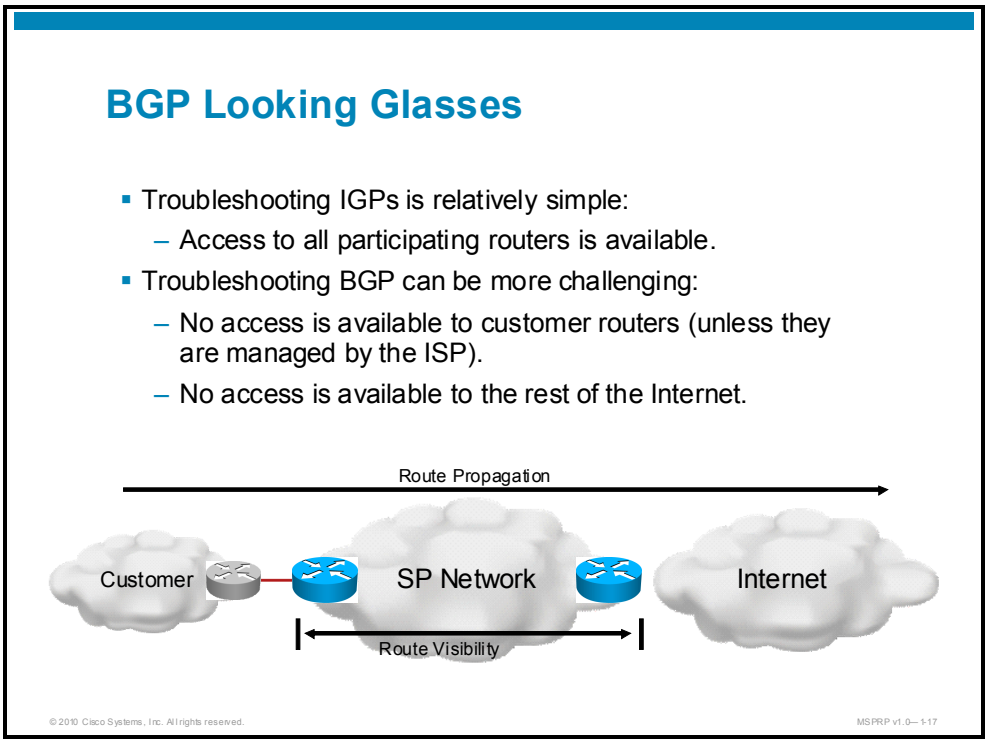

The figure illustrates an environment in which a customer is advertising its address space via BGP. The routes are then propagated throughout the Internet. The problem in this situation is that a service provider can only monitor the route propagation within its administrative domain—within its AS.

If the customer has a connectivity problem, it will want the ability to check not only the received Internet routes, but also how these routes are propagated throughout the Internet.

BGP Looking Glasses (Cont.)

- Many ISPs and exchange points host a looking glass server:
  - Web-based GUI
  - Limited interface to BGP commands in Cisco IOS Software
  - Provide public access
- There are also some routers directly accessible via Telnet (less common)
- Typically available options:
  - BGP commands
  - Ping
  - Traceroute

BGP looking glasses are web-based applications that are used to provide public or private access to some Cisco IOS commands that can be executed on a distant router. These routers are typically dedicated for looking glass purposes and do not do any packet forwarding, only peering with other routers via BGP. Looking glasses typically implement a subset of **show** commands for BGP as well as ping and traceroute.

BGP Looking Glasses (Cont.)

- Compile a list of useful looking glasses:
  - Ask peering providers.
  - Check with registries.
  - Search the Internet.

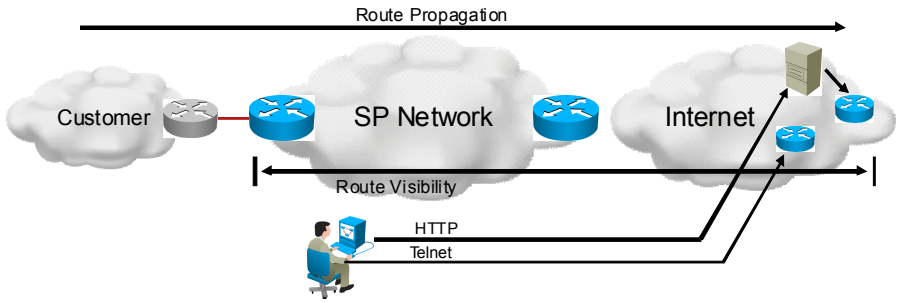

The figure illustrates a network administrator of an AS. The administrator is accessing a BGP looking glass server or publically available router to inspect the contents of the BGP table or perform a ping or traceroute for the investigated prefix.

A number of BGP looking glasses are available throughout the Internet and can be found simply by searching the web. Internet exchange point (IXP), internet registries, or other resources can be used to get a list of BGP looking glass servers as well.

Example: Looking Glasses

RIPE

1. Select from a list of exchange points.
2. Select the command.
3. Optionally, enter an argument.

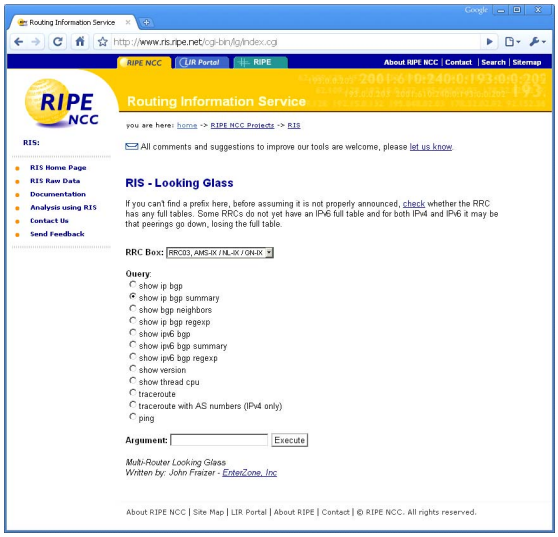

The screenshot illustrates a looking glass user interface as available from Réseaux IP Européens (RIPE). It provides access to some major European IXPs. It allows basic BGP commands, ping, and traceroute.

Example: Looking Glasses

RIPE (Cont.)

- Translates a web request to a CLI command
- Executes the command on a dedicated router
- Retrieves output and displays it on a web page

```
Router# show ip bgp 198.133.219.0
```

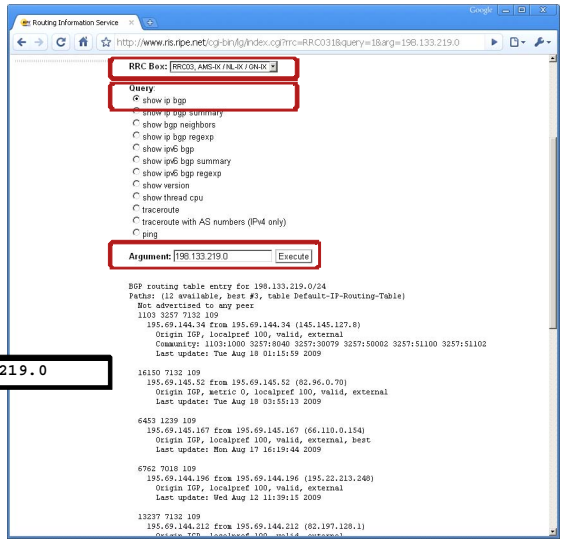

From the RIPE BGP looking glass web page, you can select an IXP, select a command, and optionally enter an argument to execute a monitoring command on a distant router. The sample looking glass operations result in the shown command on a router in an Amsterdam IXP.

Example: Looking Glasses

Oregon Exchange Point

- Connect via Telnet to a public router.
- Log in.
- Execute the required authorized command.

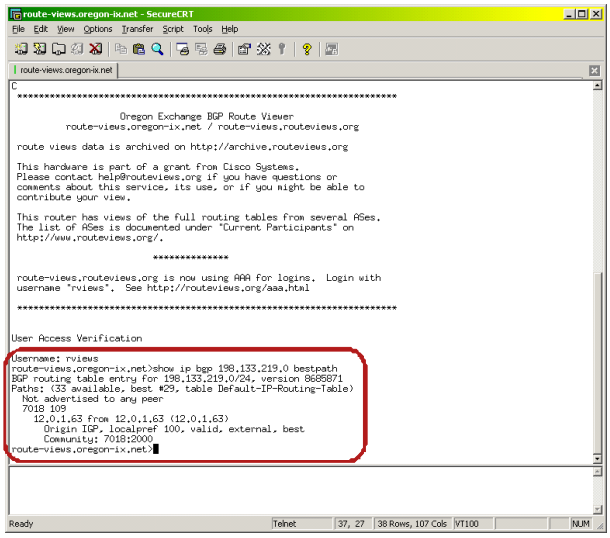

The second example shows a publically available router in the Oregon IXP. Like web-based looking glasses, users of such routers are typically also limited to a small set of commands to prevent any accidental damage or deliberate attacks on these routers.

Provisioning Tools

This topic describes the characteristics and requirements for provisioning tools.

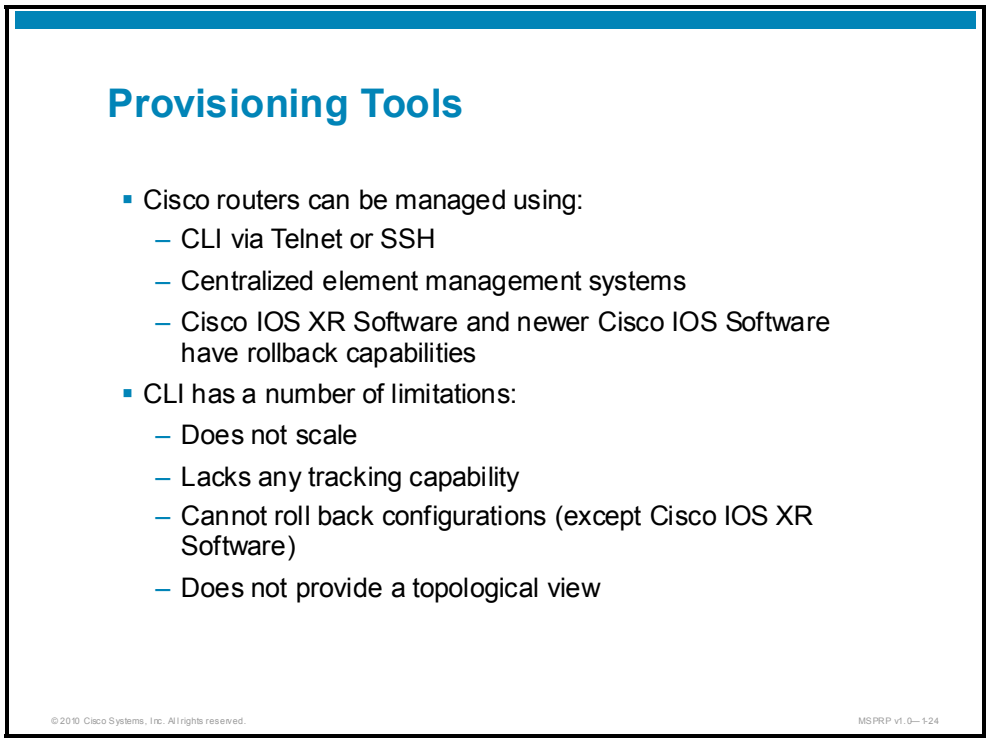

Cisco routers, like most other network devices, can be managed using the CLI that is accessible via Telnet or Secure Shell (SSH). Centralized element management systems are typically used to manage customer services in large environments. Management systems can greatly simplify management processes as well as provide configuration tracking, role-based access control, a topological view of the network, and so on.

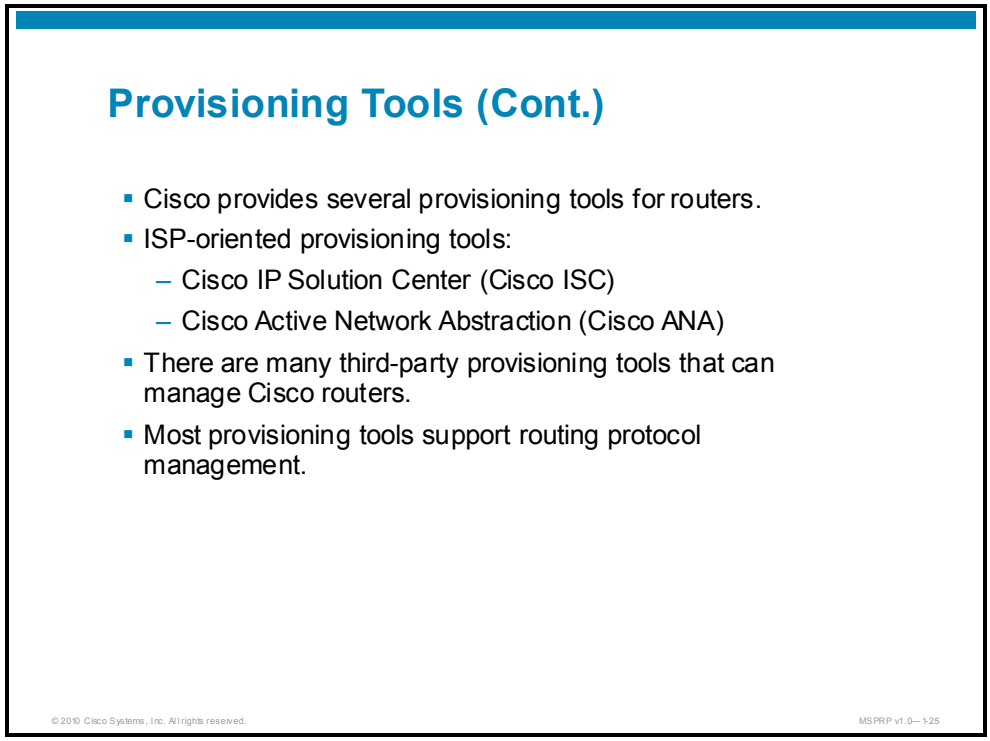

A number of management tools are available from Cisco and third parties. The ISP-oriented provisioning tools such as Cisco IP Solution Center (Cisco ISC) and Cisco Active Network Abstraction (Cisco ANA) are used in large service provider environments to manage customer connectivity and services for IP and MPLS.

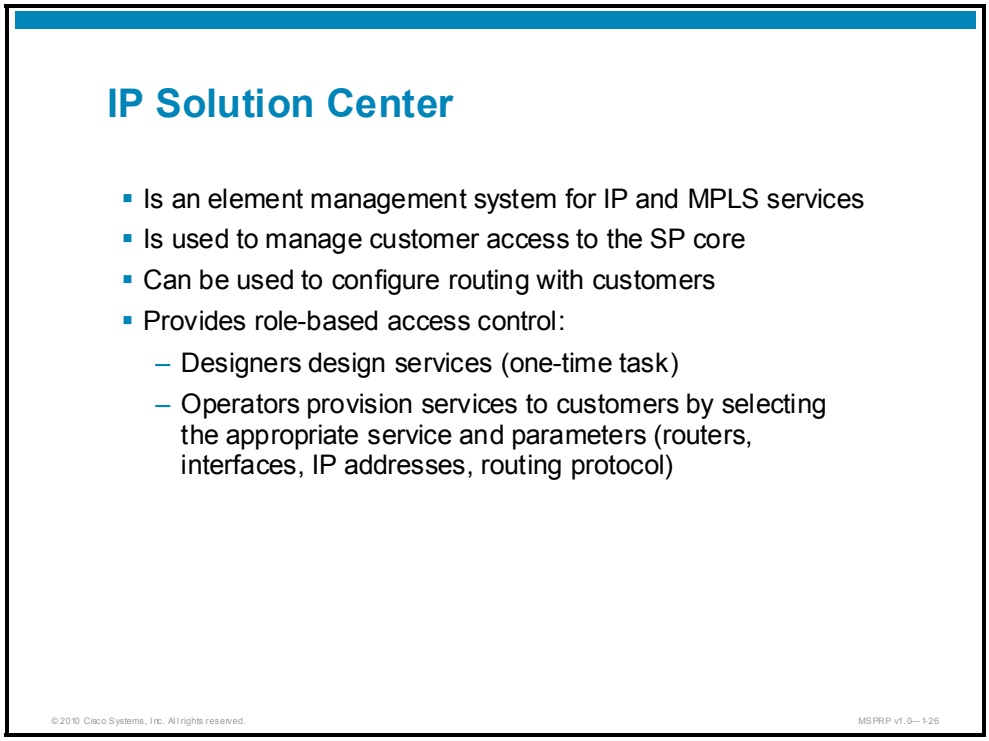

Cisco ISC is a family of intelligent network management applications that help reduce administration, management, and network operational costs:

- Providing automated resource management and rapid profile-based provisioning capabilities that speed deployment and time-to-market for MPLS and Carrier Ethernet technologies
- Working with Cisco MPLS Diagnostics Expert to provide automated, workflow-based troubleshooting and diagnostic capabilities for MPLS VPN networks

Cisco IP Solution Center applications can operate alone or as a suite. Capabilities include the following:

- Provisioning and troubleshooting for MPLS VPNs; ATM, Frame Relay, and Ethernet over MPLS VPNs; and Carrier Ethernet VPNs
- Planning and configuration of MPLS Traffic Engineering (MPLS TE)

Cisco IP Solution Center also offers the following:

- Scalable and reliable architecture for large-scale operations
- Single server solution and Web GUI

Open application programming interfaces (APIs) and operations support system (OSS) interfaces that help:

- Integrate IP VPN services into existing infrastructure
- Integrate Cisco fault-management products with independent software vendor products for VPN-aware performance reporting

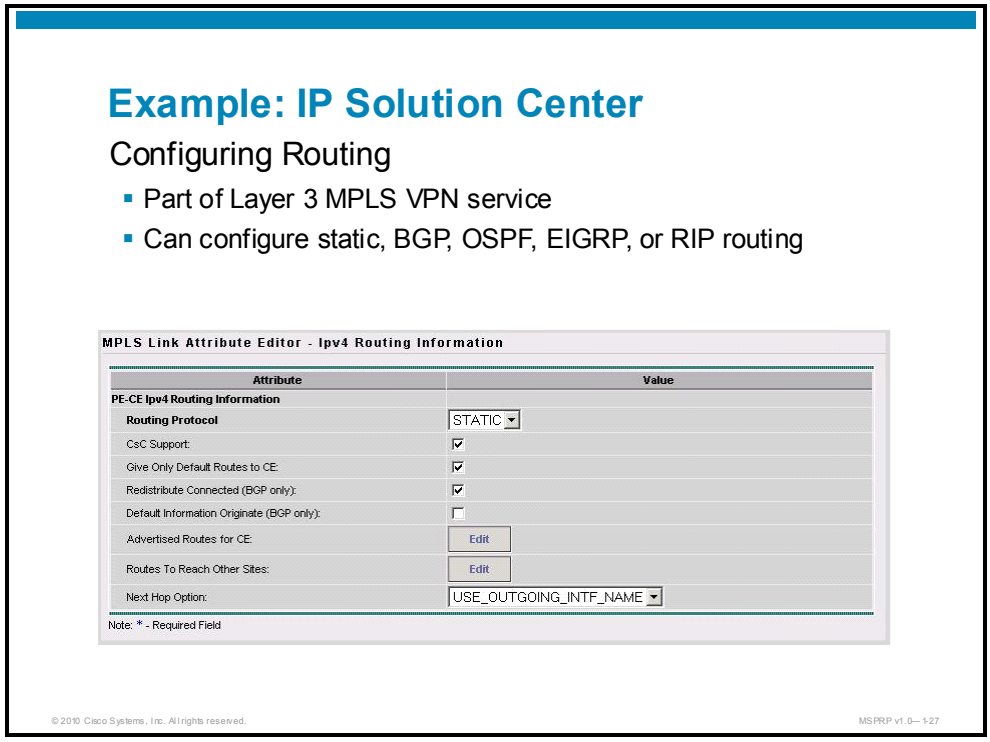

The figure illustrates one step in the workflow-based provisioning of services to end customers that also includes the configuration of a routing protocol. Cisco ISC supports the configuration of static and connected routing, Routing Information Protocol (RIP), OSPF, Enhanced Interior Gateway Routing Protocol (EIGRP) and BGP. A number of routing parameters can be made available to the operator upon service provisioning, or they may be fixed by the designer at the time of service creation.

Refer to http://www.cisco.com/go/isc for more information on Cisco ISC.

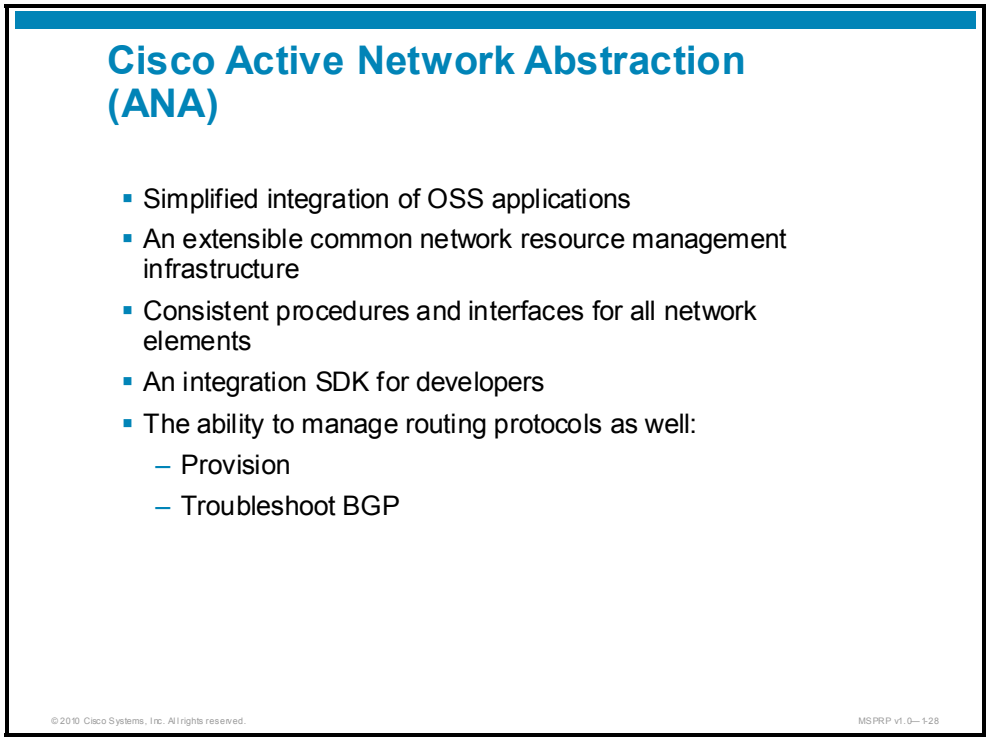

Cisco ANA is a powerful, next-generation network resource management solution. It is designed with a fully distributed OSS mediation platform, which abstracts the network, its topology, and its capabilities from physical elements. Because Cisco ANA is virtual, it gives customers a strong and reliable platform for service activation, service assurance, and network management.

Cisco ANA also provides a flexible, vendor-neutral resource management system. The system supports a multiservice network environment and the management capabilities required to sustain reliable and converged voice, video, and data networks.

A comprehensive developer program is available to help end-user and partner developers integrate and test the interoperability of existing OSS applications with Cisco ANA. This program also includes various go-to-market opportunities for partners to promote their applications to Cisco ANA end users.

Cisco ANA offers advanced features such as:

- Simplified integration of OSS applications with near real-time network information from one authoritative source
- An extensible common network resource management infrastructure
- Consistent procedures and interfaces for all network elements
- An integration software development kit (SDK) for developers

Refer to http://www.cisco.com/go/ana for more information on Cisco ANA.

Administration

This topic summarizes requirements related to management of routing infrastructure in large service provider environments.

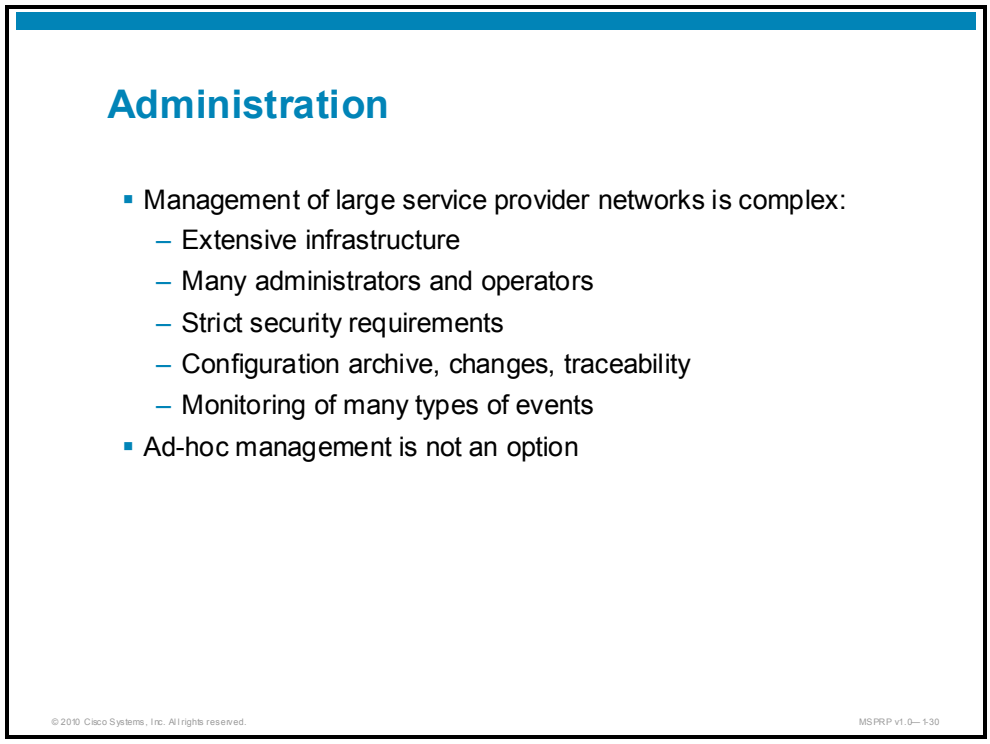

Administration of a service provider environment entails many complex tasks:

- Extensive infrastructure including large numbers of different types of devices using different connectivity methods
- Large groups of network and system administrators and operators working in different regions and departments
- Strict security requirements that are implemented using adherence to security policies and strict implementation of authentication and authorization
- Configuration archives and changes required in order to trace actions
- Monitoring of many types of events for security, availability, and performance reasons

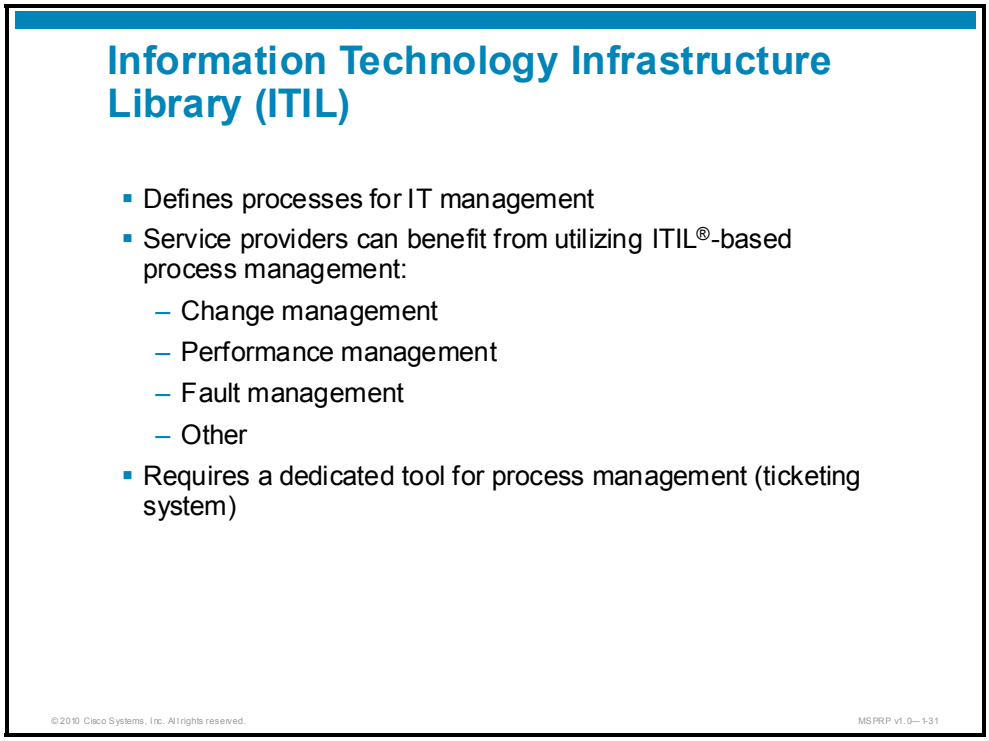

Information Technology Infrastructure Library® (ITIL) defines processes for IT management. It originated in the 1980s, when the British government realized that there was a need to standardize increasingly used IT resources.

In this course, you will focus on change management, performance management, and fault management.

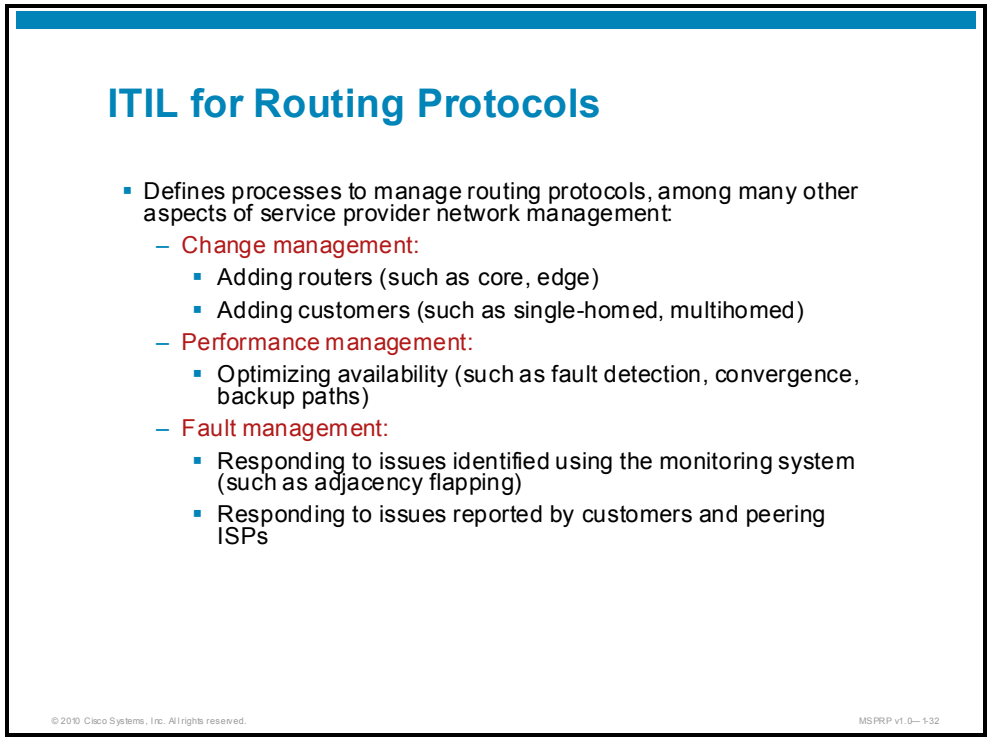

ITIL® defines processes to manage routing protocols, among many other aspects of service provider network management:

Change management for routing protocols can include actions such as these:

- Adding routers to the network core or edge
- Connecting single-homed or multihomed customers (and possibly provisioning a managed customer router)

Performance management for routing protocols can include actions such as these:

- Optimizing availability by improving fault detection, improving routing protocol convergence, providing backup paths, and so on

Fault management for routing protocols can include actions such as these:

- Responding to issues that are identified using the monitoring system (for example, adjacency flapping, lost routes, traffic pattern change)
- Responding to issues that are reported by customers or peering ISPs based on their monitoring systems or based on user reporting of problems

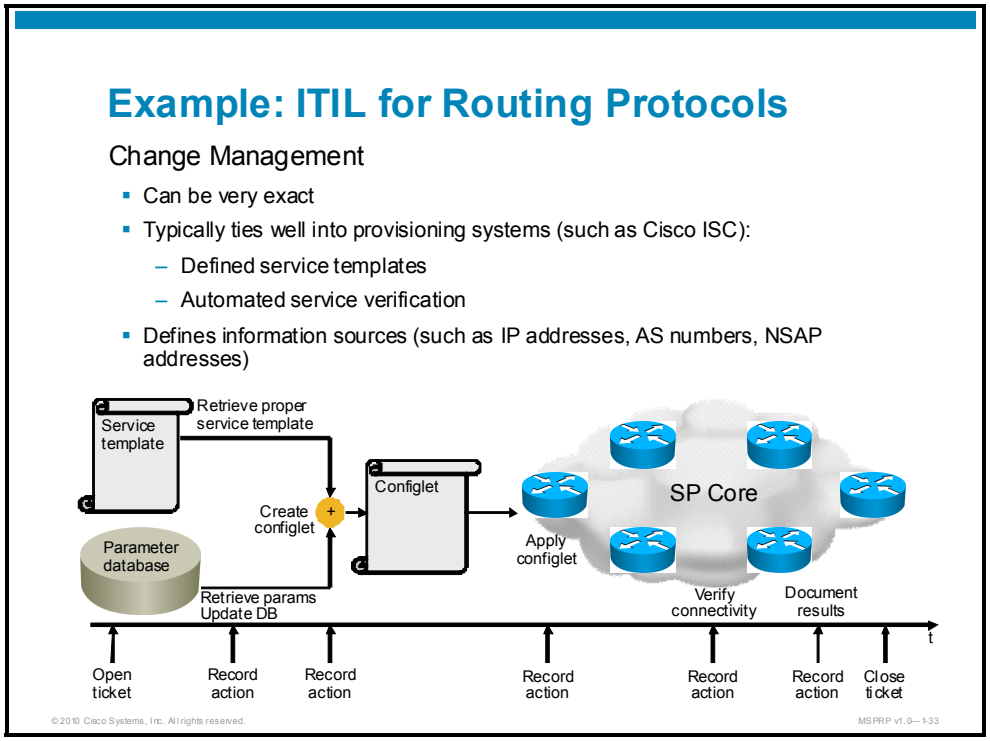

The figure illustrates a sample change management process in which the process starts when a sales administrator creates a change-management ticket in the ticketing system. The operator follows a standardized set of actions to complete the task, such as the following:

1. Identify the type of customer.
2. Identify the service templates that are required to provision a link to the new customer.
3. Identify the required parameters (for example, get a /30 IP subnet for the access link from the available space in the pool of addresses used for access links).
4. Create a configuration template or simply input the information into the provisioning tool.
5. Apply the configuration to the appropriate network device or devices.
6. Verify the operation of the provisioned service or services.
7. Document the process.
8. Close the ticket.

Summary

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

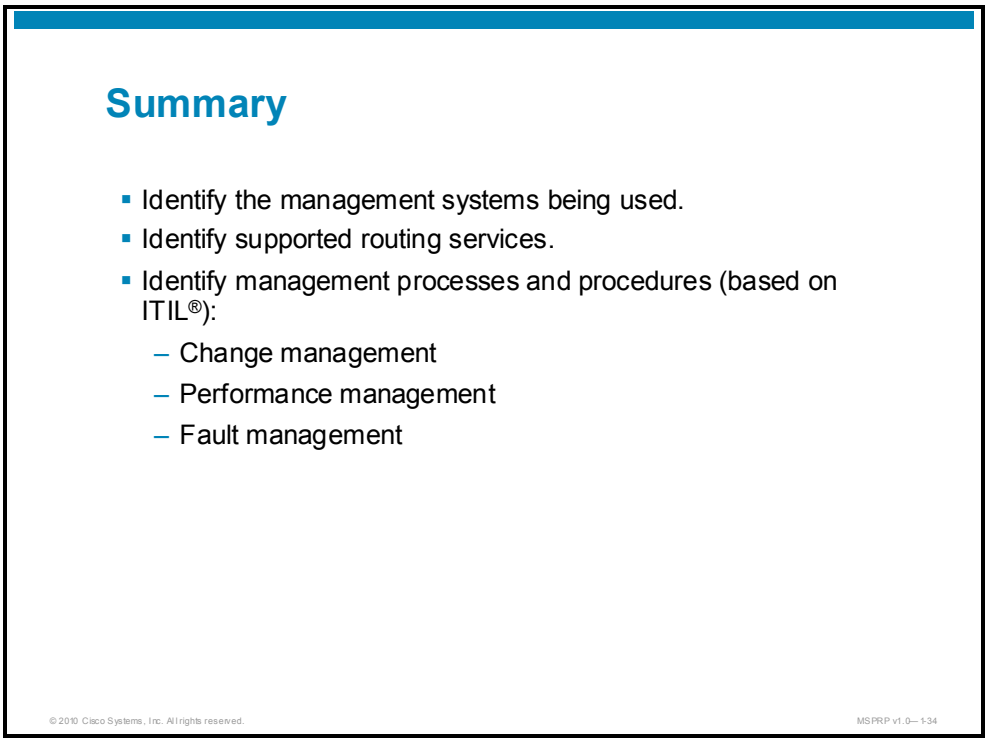

Lesson 4

Applying Service Provider Routing Operation Processes Based on ITIL

Overview

This lesson provides a brief description of the Information Technology Infrastructure Library® (ITIL) standard that is widely used in enterprise and service provider environments. ITIL® is used to govern the processes for managing infrastructure and services. The lesson describes how to apply ITIL® processes to the management of infrastructure for routing operations services.

Objectives

Upon completing this lesson, you will be able to apply ITIL® processes to the management of routing operations technologies and services. This ability includes being able to meet these objectives:

- Identify the main components of ITIL®
- Describe change management procedures
- Describe performance management procedures
- Describe fault management procedures

ITIL Overview

This topic describes the major components of ITIL® and how it applies to routing operations services.

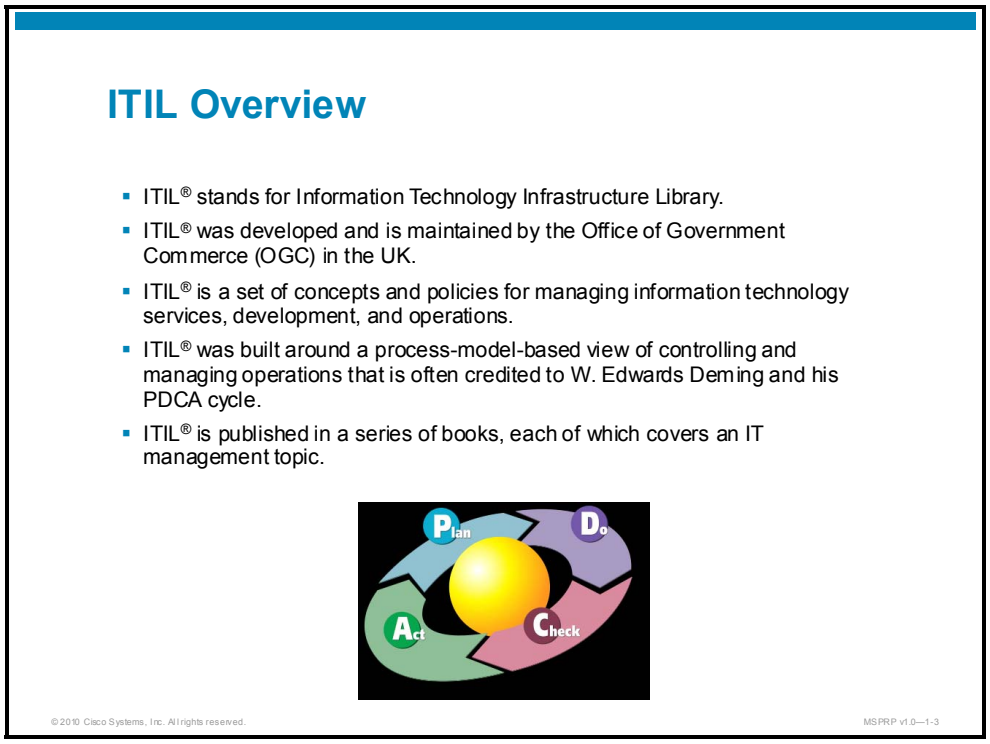

ITIL® was originally developed by the predecessors of the Office of Government Commerce from the United Kingdom. ITIL® was developed in order to standardize the processes for managing IT infrastructure in governmental institutions.

ITIL® gives a detailed description of a number of important IT practices, along with comprehensive checklists, tasks, and procedures that any IT organization can tailor to its needs. ITIL® is published in a series of books, each of which covers a specific management practice within IT service management.

ITIL® was built around a process-model-based view of controlling and managing operations that is credited to W. Edwards Deming and his plan-do-check-act (PDCA) cycle. PDCA is an iterative four-step problem-solving process typically used in business process improvement.

ITIL Versions

- ITIL® version 2 is widely used and primarily focuses on:
  - Service delivery
  - Service support
- In 2009, OGC announced that ITIL® version 2 would be withdrawn.
- ITIL® version 3 is more comprehensive (released in 2007), and is designed to support all modern IT components:
  - Service strategy
  - Service design
  - Service transition
  - Service operation
  - Continual improvement

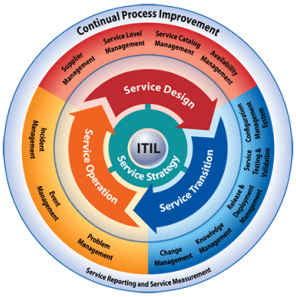

In 2000, ITIL® version 2 consolidated all ITIL® publications into eight books. This consolidation grouped related guidelines to match different aspects of IT management, applications and services. The two books that ITIL® version 2 primarily focused on were as follows:

- **Service Delivery:** Primarily concerned with proactive services delivered by information and communication technology (ICT) to provide adequate support to business users
- **Service Support:** Focused on users of ICT services and is primarily concerned with ensuring that they have access to the appropriate services to support business functions

Enterprises later adopted ITIL®, and ITIL® was gradually refined to version 3, which included the following books:

- **Service Strategy:** Provides guidance on clarification and prioritization of service provider investments in services. More generally, Service Strategy focuses on helping IT organizations improve and develop over the long term.
- **Service Design:** Provides good practice guidance on the design of IT services, processes, and other aspects of the service management effort.
- **Service Transition:** Relates to the delivery of services that are required by the business into operational use.
- **Service Operation:** Describes best practices for achieving the delivery of agreed levels of services both to end users and the customers.
- **Continual Improvement**: Describes how to align and realign IT services to changing business needs by identifying and implementing improvements to the IT services that support the business processes.

The latest version is the most comprehensive and includes the processes of managing services and infrastructure from initial introduction to the IT environment to operation, monitoring, and improvement. ITIL® can also be used in any IT environment, including service provider environments.

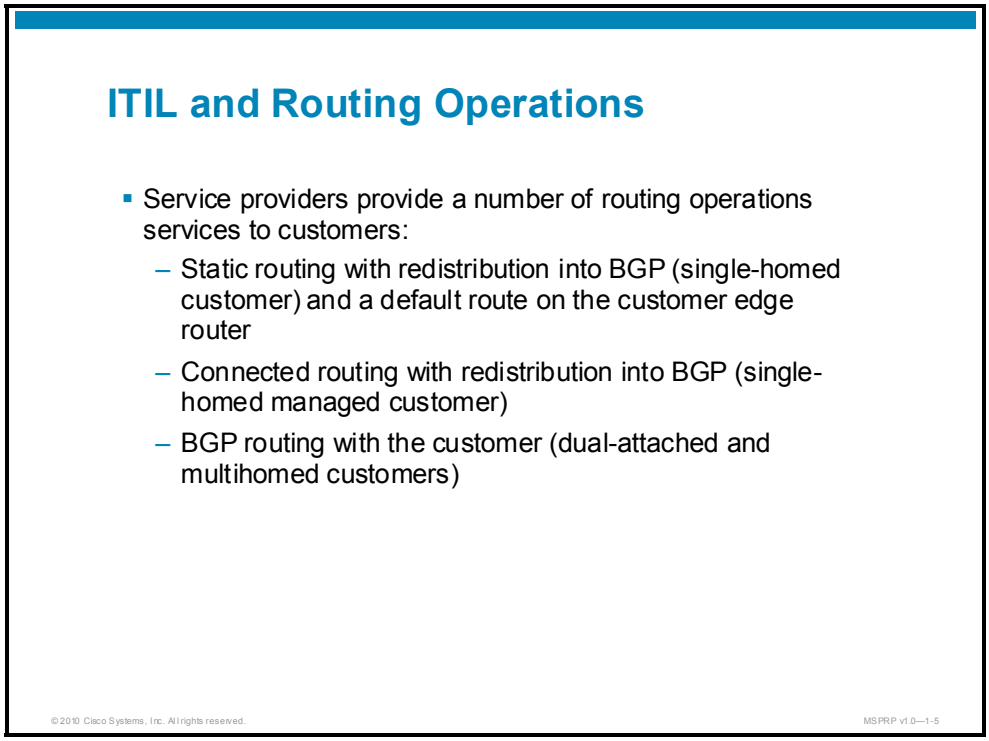

When addressing routing operations services, you can use ITIL® to govern the physical infrastructure and software, as well as individual components, technologies, and services. Service providers offer different routing services to customers. These services can include the following:

- Static routing of customer routes with redistribution into BGP at the provider site and a default route at the customer site. This scenario applies to single-homed customers.
- Connected routing of customer routes with redistribution into BGP at the provider site. This scenario applies to managed single-homed customer.
- Border Gateway Protocol (BGP) routing with the customer, which applies to dual-attached and multihomed customers.

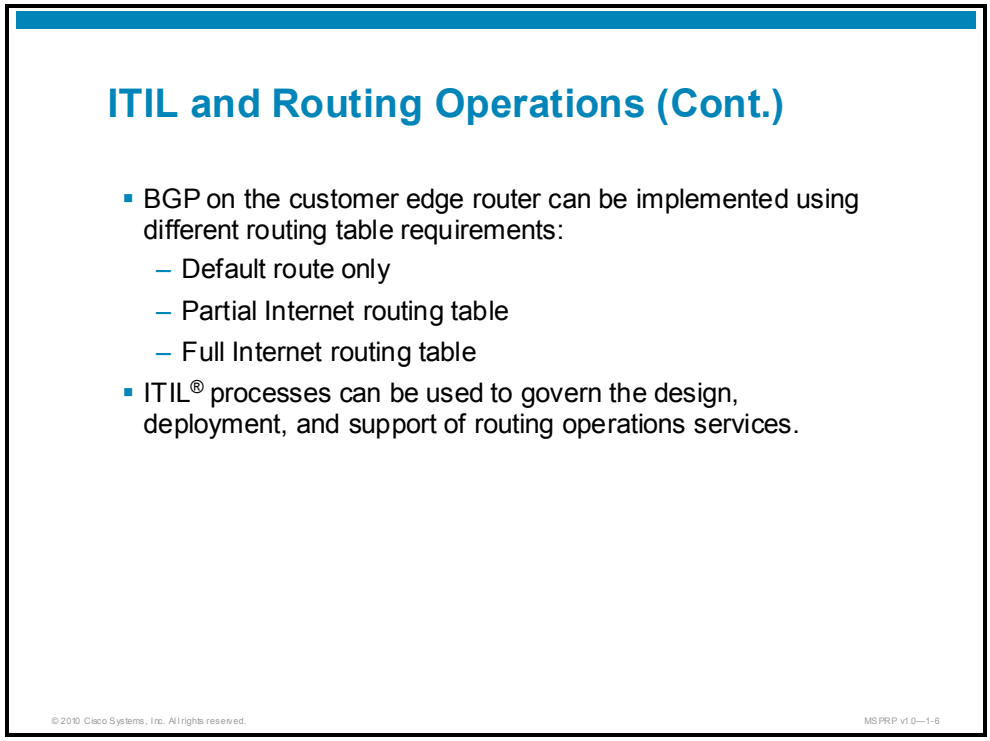

BGP on the customer side can be implemented in three different ways, depending on customer routing table requirements. The customer can receive any of the following:

- Default route only: The customer receives only the default route, which is used to reach all networks.
- Partial Internet routing table: The customer receives routes specific to the service provider and the default route to reach all other networks.
- Full Internet routing table: The customer receives the complete Internet routing table.

Routing services depend on basic connectivity and the number of technologies and protocols, such as BGP and interior routing protocols. All aspects of device configurations are governed by processes that are defined in ITIL® to ensure the required functionality, performance, availability, and security.

ITIL and Routing Operations (Cont.)

- In this course, ITIL® processes are grouped into three major categories for routing operations network services:
  - Change management (infrastructure deployment, service deployment, customer deployment)
  - Performance management (SLAs, capacity, availability)
  - Fault management (faults, issues, and problems)
- The service desk, with level 2 support staff, is responsible for change, performance, and fault management.

This course is divided into three major sections, which encompass various components of ITIL®:

- The **change management** section describes the processes used for infrastructure deployment, new service deployment, and deployment of services to individual customers.
- The **performance management** section describes the methods and processes used to monitor the operation of the network and services. Monitoring of the network and services is needed to ensure compliance with service level agreements (SLA), the availability of adequate capacity (bandwidth, CPU), and the high availability of services (resilience to loss of resources).
- The **fault management** section describes the methods and procedures used to identify faults (physical and logical failures of equipment or software), issues (as reported by customers), and problems (recurring issues).

A service desk is made up of support personnel that address faults, issues, and problems. This course focuses on the level 2 support staff the service desk, which is responsible for most network issues that the first-level support cannot solve.

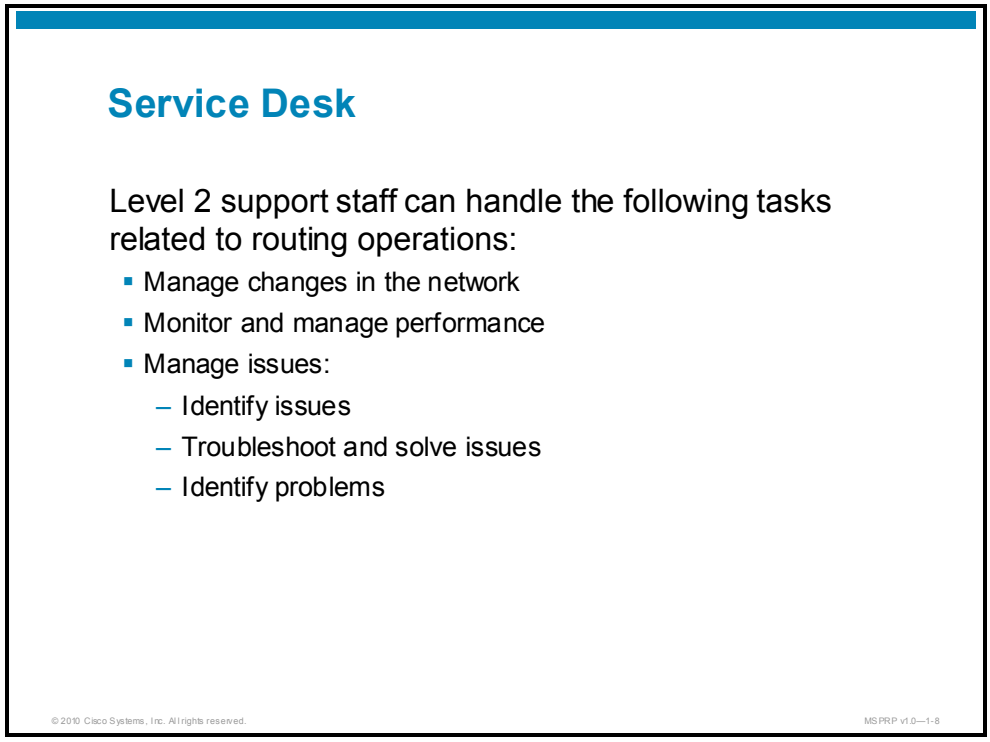

A service provider will often design its internal organization in their own way, possibly following ITIL® or other industry guidelines. The organization also depends on the size of the service provider. In general, level 2 support staff can handle any of the following tasks related to routing operations:

- **Management of changes in the network:** Changes can be anything from minor modifications with little or no impact on performance to complete redesigns or deployments of new services networkwide.
- **Monitoring and management of performance:** Performance monitoring and metering can include the ability to determine the performance characteristics of control plane mechanisms, data plane forwarding performance, the performance of individual services, and compliance with various SLAs.
- **Management of issues:**
  - Identify issues that customers or monitoring servers report. Issues can result from physical or logical faults, or from misconfigurations.
  - Troubleshoot and solve issues using known fixes and workarounds
  - Use standard troubleshooting best practices to find solutions for unknown issues.
  - Identify problems (recurring issues) and devise generic solutions that can be suggested to the network engineering department.

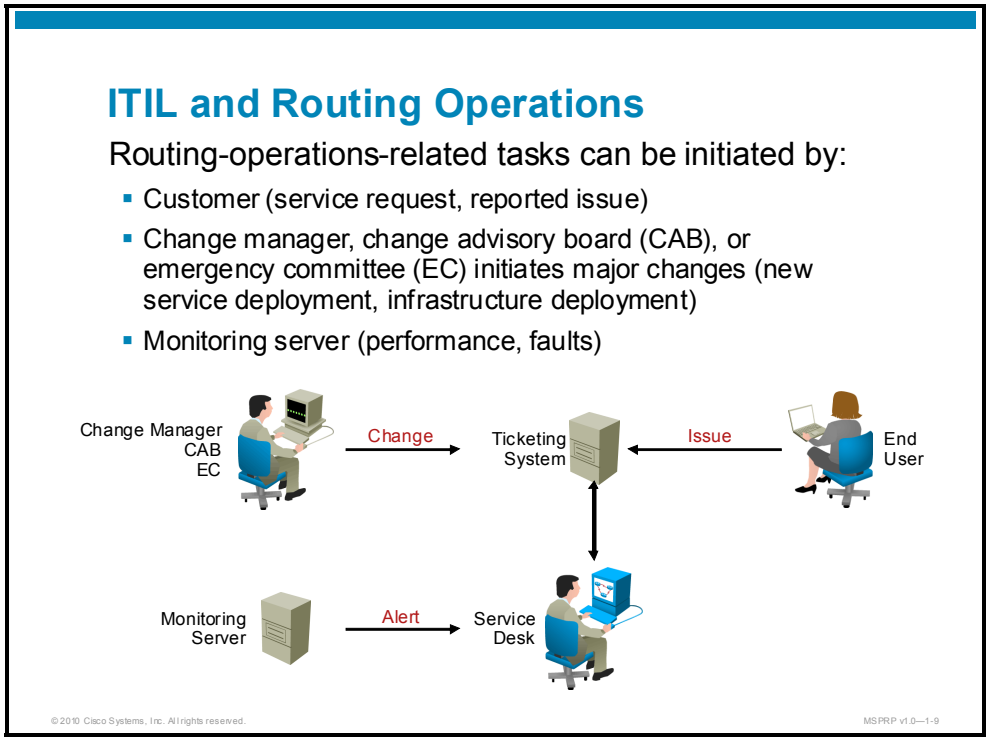

The figure illustrates the users and administrators involved in an ITIL®-based service provider environment. The service desk, with level 2 support staff, is in the center of the management of the infrastructure. A ticketing system (process management software) is used to keep track of individual processes (content, status, resolvers). A “ticket” is used to start a process and keep track of the process until the process is completed and the ticket is closed. The following are some examples of processes that can take place in a service provider environment:

- A customer reports an issue using the publically available customer portal of the ticketing system. The issue is assigned to a support engineer (level 1 or level 2, depending on the type of customer and issue). The support engineer tries to resolve the issue and close the ticket. Alternatively, the issue can be escalated to level 3 support personnel or event to technical assistance of the network device vendor.
- A monitoring server can detect a failure or a monitored value can exceed a threshold, causing an alert about the issue to be sent to the service desk. The issue can then be investigated and rated according to its criticality and a solution can be found to mitigate the issue.
- A change manager can also initiate a task. This might occur when a redesign of a certain service has been completed and the change must be deployed throughout the network. A ticket is created, which requires that the service desk implement the changes. Service providers can have multiple instances of overseers, depending on the criticality of changes that need to be made. A change advisory board and emergency committee are used to decide on major changes that may have a high impact on the network and services.

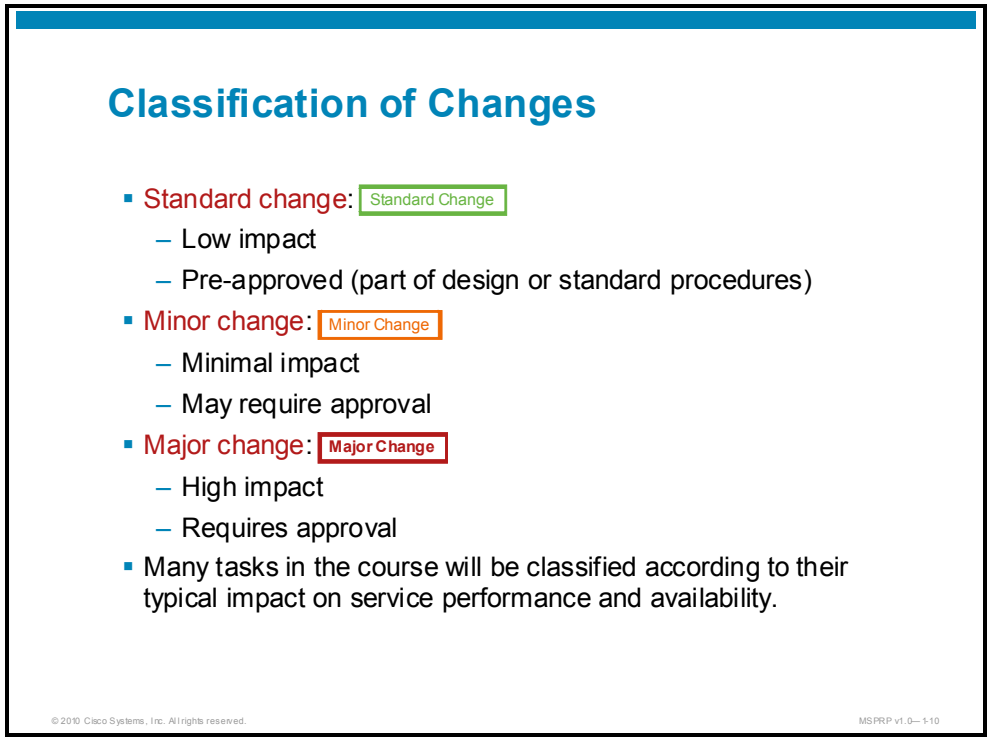

Changes that are being implemented can have different importance, scope, and (most importantly) impact on the network and services.

- Standard changes are typically preapproved and are designed to have a low impact.
- Changes that have a noticeable impact are typically implemented during regularly scheduled maintenance windows and typically require approval from the change manager or a higher-level manager.
- Major changes have a high impact. To implement them, you should use one or multiple maintenance windows.

The course discusses many different types of changes that can result from new deployments or troubleshooting. These three levels are used to categorize changes according to their impact.

Change Management

This topic describes change management, which is related to routing operations services.

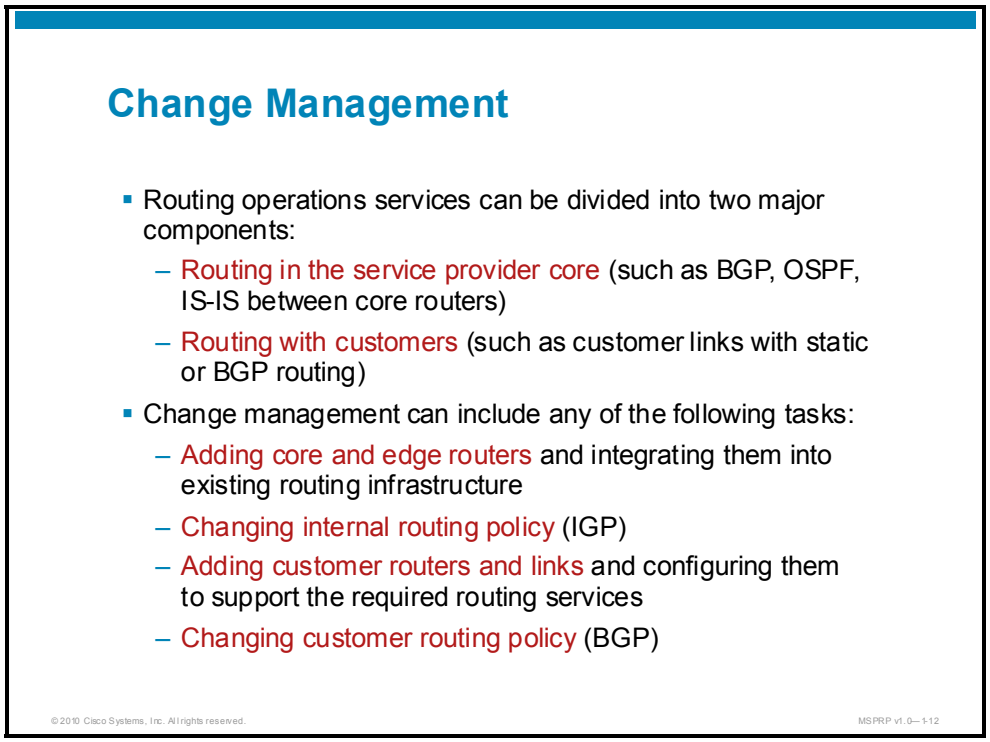

Routing operations can be divided into two major components:

- **Routing in the service provider core network:** Interior gateway protocols (IGPs), such as Open Shortest Path First (OSPF) and Intermediate System-to-Intermediate System (IS-IS) and internal BGP are used.
- **Routing between service provider and customers:** Static or external BGP routing is used.

The major changes that affect these two components are as follows:

- **Adding and configuring core and edge service provider routers**: To satisfy capacity requirements, core devices must be added.
- **Changing internal routing policy**: Path selection inside the service provider network can be influenced using IGP metric.
- **Adding and configuring customer routers and links**: Customers must be added when there is a request.
- **Changing customer routing policy**: The routing policy for customer reachability to and from the Internet and other networks can be influenced by changing BGP attributes.

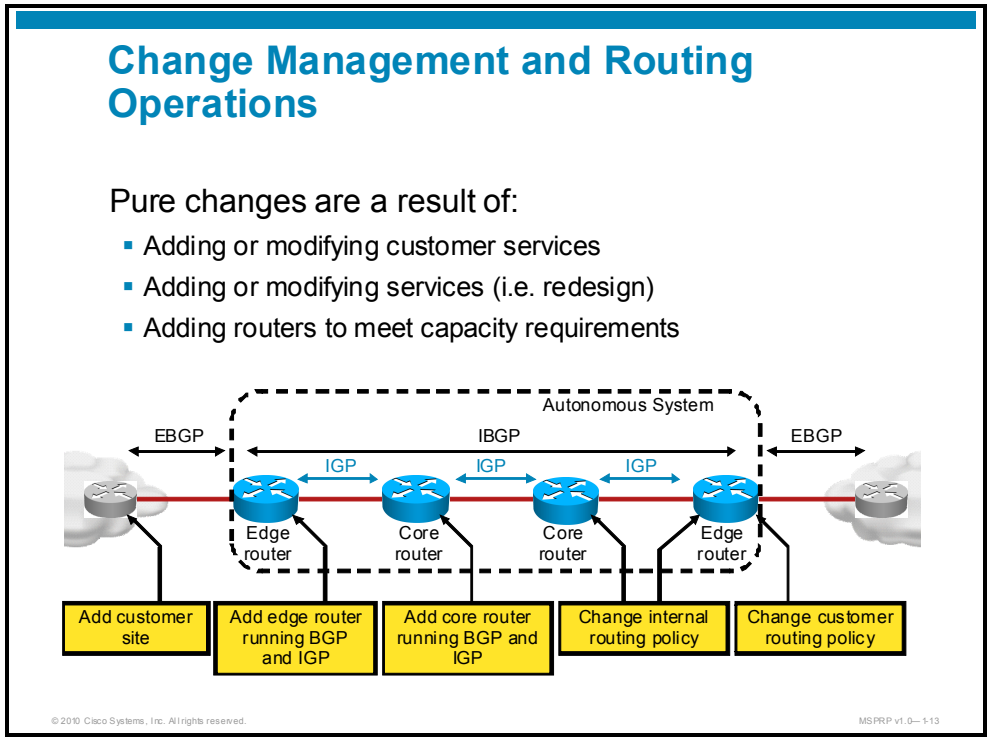

The figure illustrates a service provider network with all protocols that are required to support different customer routing requirements:

- IGP inside the service provider core network
- Internal BGP (IBGP) inside the service provider core network
- External BGP (EBGP) between customers and the service provider

To add a new customer, you must configure the edge router, the routing protocol (BGP or static routing) between the customer and provider edge router, and the customer router. To change a customer routing policy you must change BGP attributes on the provider edge and customer routers. To add a new router in the core network to increase the capacity you must do the following:

- Configure the new router with all the needed protocols (BGP, IGP)
- Integrate the necessary protocols with the existing infrastructure

Changing routing inside service provider core network requires configuration changes on core routers, where the IGP metric must be changed.

Change Management and Routing Operations (Cont.)

- Scheduling changes:
  - Low-impact or minimal-impact changes can be deployed at any time
  - High-impact changes require a maintenance window
- A detailed deployment plan is required for (at least) high-impact changes:
  - Deployment plan
  - Verification procedure
  - Rollback plan

Any change can be analyzed to determine its impact on the network and services. Both minor and major changes may require the use of a scheduled maintenance window during which the impact on end customers will be minimal, such as at 3 a.m. (0300).

Major changes may be composed of many steps and may take longer to implement. The detailed deployment plan should be accompanied by a detailed verification plat to ensure that the service was implemented properly and that it functions as expected. If there is an implementation failure, a rollback plane must be available to simplify reverting to the original setup in the shortest possible time.

Performance Management

This topic describes the components that are used to monitor and measure the performance of the network.

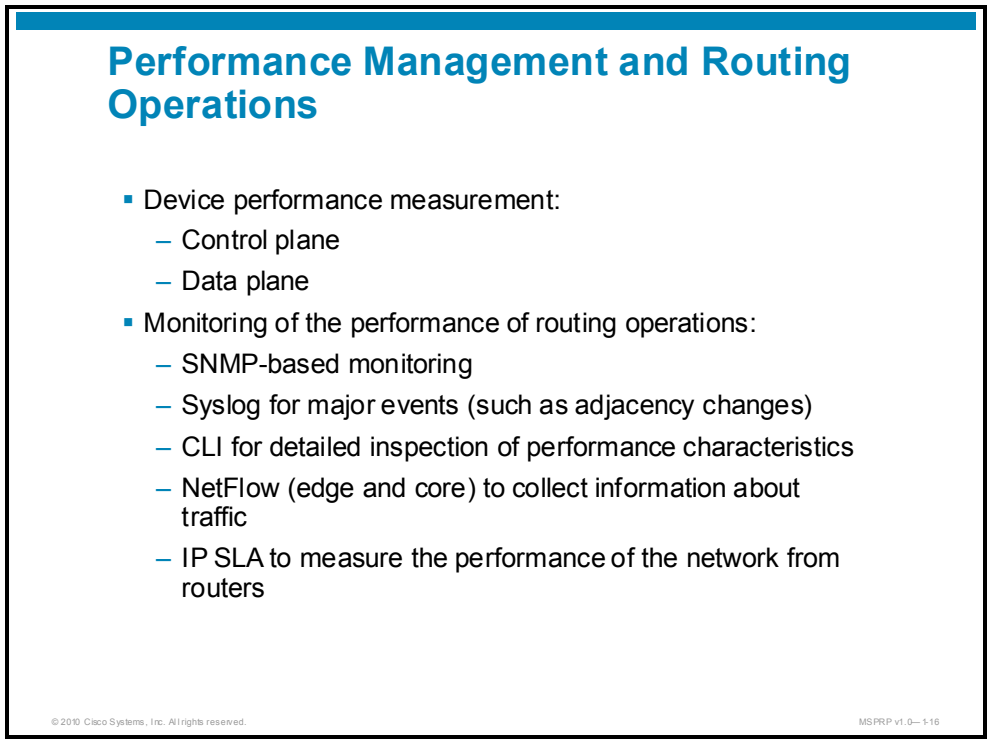

Performance management primarily consists of monitoring network resources. Monitoring of the control and data planes of monitored devices can be achieved using any of the following methods:

- Simple Network Management Protocol (SNMP)-based monitoring can be used, in which case monitoring servers periodically poll routers for the required performance statistics. Additionally, SNMP traps can be used to inform central servers of significant events in the network.
- Syslog can be used to send information about significant events to a central syslog server.
- Various CLI commands are available to investigate performance issues once they have been identified, to determine their causes and possible device actions or plans to improve performance.
- Cisco routers support NetFlow to collect detailed information about traffic in a central location. Based on traffic patterns, traffic analysis can then be used to determine performance characteristics.
- Cisco routers provide the Cisco IOS IP SLA feature, which can be used to measure the performance of the network and applications from routers. Statistics can be collected via SNMP or the CLI.

Performance Management and Routing Operations (Cont.)

- Performance can be measured as:
  - Service performance (SLA, link utilization)
  - Control plane performance (such as CPU utilization)
  - Availability (such as 99.999% uptime; redundancy; convergence)

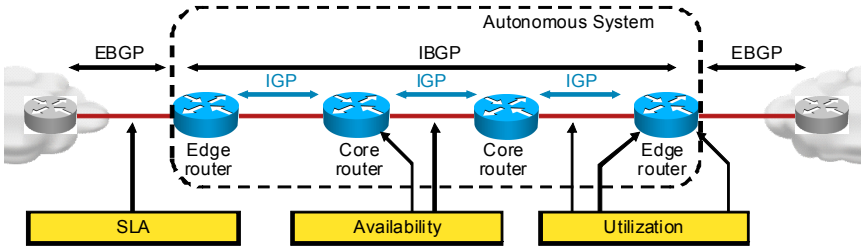

The figure illustrates various performance categories:

- Performance of services for individual customers (SLA assurance)
- Control plane performance, such as CPU utilization due to the use of various protocols or even packet forwarding in the process path
- Data plane performance, such as the utilization of the forwarding capacity of routers, modules, and links
- Availability, as determined by hardware redundancy, link redundancy, and resilience mechanisms

Performance Management and Routing Operations (Cont.)

- Performance monitoring can also be used for:
  - Capacity planning
  - Availability metering
  - Service-level assurance
  - Help with identifying issues and problems

Performance monitoring and analysis can be used to determine the performance characteristics of the network and of services, as well as to aid in capacity planning. Additionally, performance monitoring can be used for SLA assurance and to help identify issues in the network that may not be otherwise detected.

Fault Management

This topic describes fault management processes related to routing operations services.

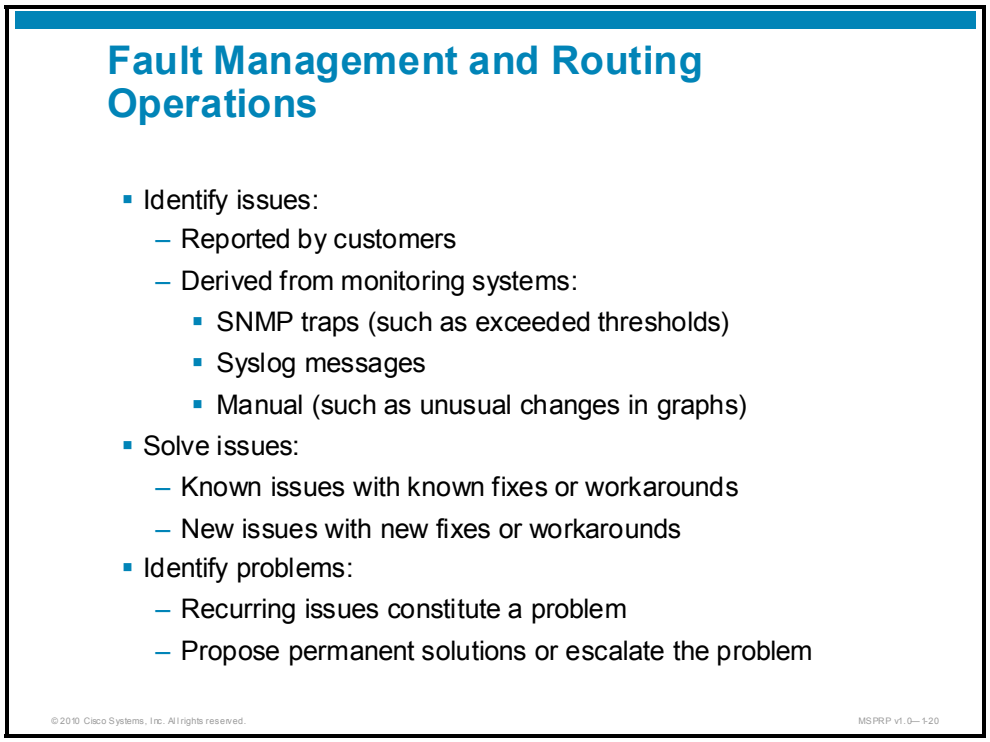

Fault management tasks require that the service desk detect faults and issues and provide fixes or escalation to level 3 support.

Faults and issues can be identified in two ways:

- Customers report issues when their services fail.
- The service desk receives alerts from the monitoring systems indicating a failure or an issue that needs resolution.

Known issues may already have known fixes that can be deployed quickly. Unknown issues may require more troubleshooting. Changes that are required to resolve an issue should be examined for their potential impact on the network and other services.

If the service desk determines that there is a significant number of recurring issues then it should escalate it to level 3 support. This problem needs detailed analysis and a potential change to the design. An upgrade of router software or some other action that will ensure that the problem is resolved may be also needed.

Fault Management and Routing Operations (Cont.)

- Sources of failures:
  - Physical failures (link, device)
  - Software failures (bug, misconfiguration)

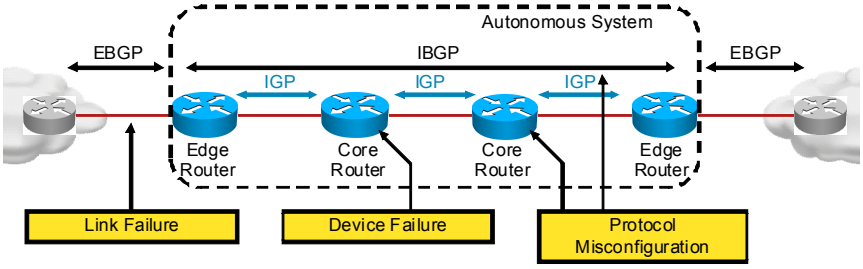

The figure illustrates hardware and software failures that can result in different levels of service degradation or service loss. A customer link failure, for example, will result in the loss of service for a single customer. A core link failure, on the other hand, will result in temporary loss of service to many customers while the network is converging. A sudden protocol failure is typically a result of a software bug that only happens after a certain time and may result in many different issues. Protocol failure can also be a result of a denial of service (DoS) attack if security is not implemented properly.

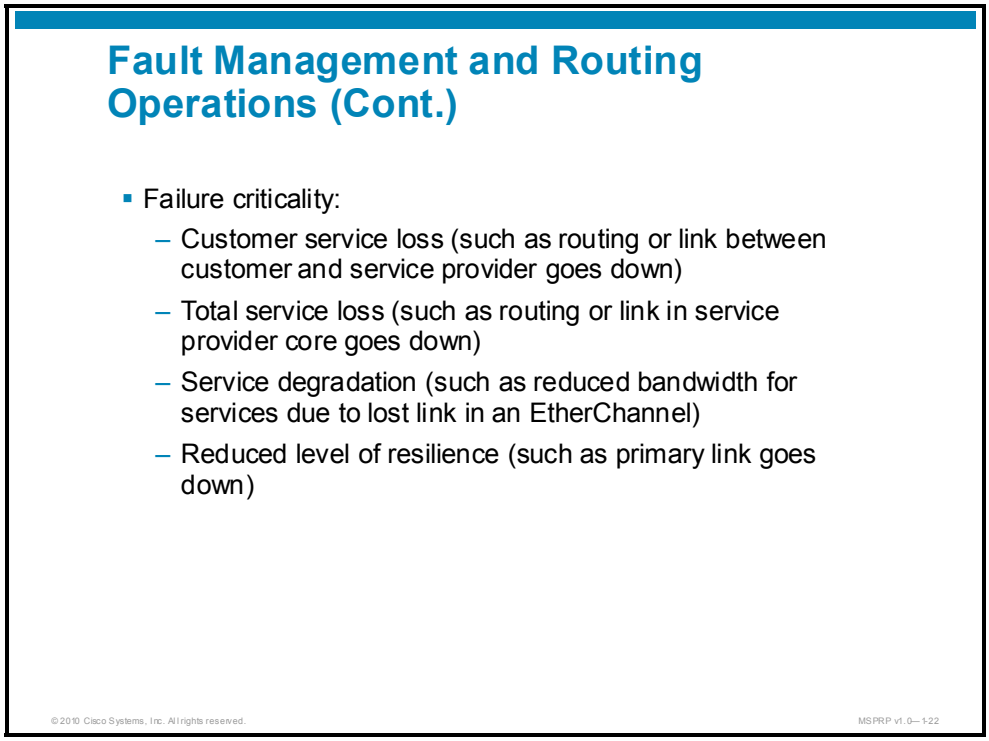

Failures may have different criticalities:

- Customer service loss, such as occurs when routing or a link goes down, may be regarded as highly critical for the customer, but not that critical for the service provider.
- Total service loss, such as occurs when routing protocols, devices, or links inside the provider network go down, may result in total loss of service for many or all customers. This service loss is a highly critical failure. Such failures would first have to be mitigated to restore services and then investigated to determine the reasons for the failure. After that, appropriate actions can be taken to prevent future reoccurrence of the failure.
- Service degradation, such as reduced bandwidth for services due to a lost link in an EtherChannel will typically not be very critical. In this case, service is not lost, but the performance may be degraded.
- Reduced level of resilience is also a result of a failed resource when a redundant resource is available. Failing to mitigate the failure increases the chances of the redundant resource failing sometime in the future, thus causing a complete failure.

Summary

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

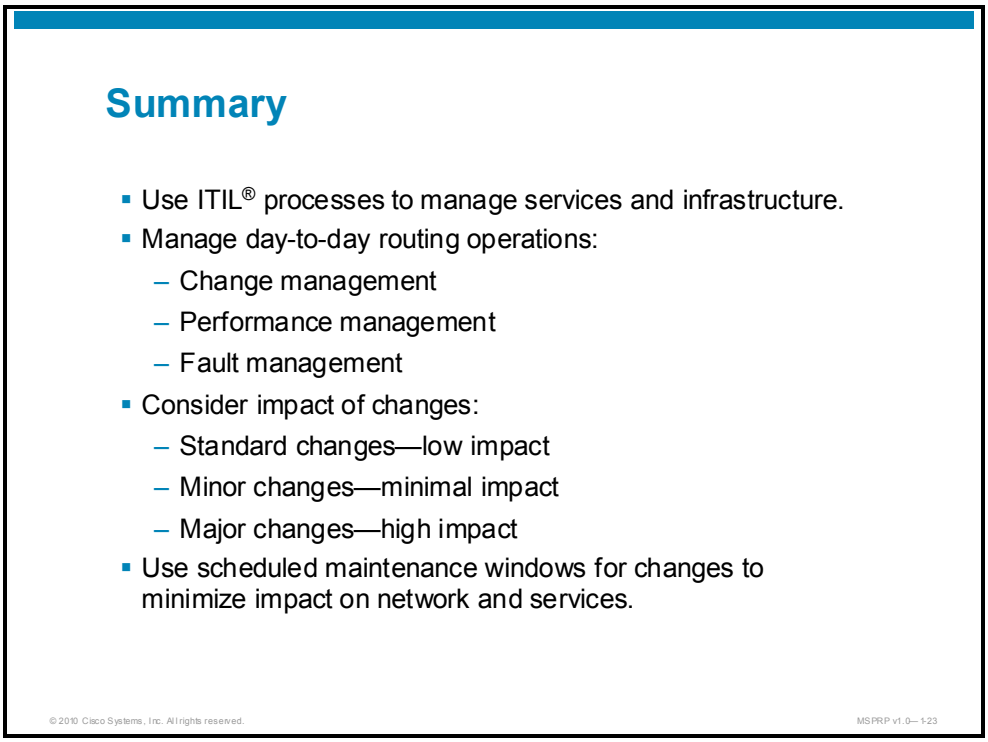

Lesson Self-Check

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

Q1) ITIL® is which two of these? (Choose two.) (Source: ITIL Overview)

A) a library of books with IT content
B) a collection of configuration templates for service provider network devices
C) a set of best practices, concepts, and policies for managing IT
D) a registered trademark of the U.K. OGC

Q2) Classify the following changes with severity levels according to their impact on network services operations. (Source: Change Management)

A) _____ Change IGP inside the service provider network from OSPF to IS-IS
B) _____ Change the IGP metric to influence path selection inside a service provider network.
C) _____ Upgrade network devices with more powerful models.
D) _____ Change customer routing policy to influence path selection.

Severity levels:

_____ 1. standard

_____ 2. minor

_____ 3. major

Q3) Which three of these can be used for performance management? (Choose three.) (Source: Performance Management)

A) SNMP monitoring
B) syslog monitoring
C) SMTP monitoring
D) NetFlow

Lesson Self-Check Answer Key

Q1) C, D

Q2) A-3, B-2, C-1, D-3

Q3) A, B, D

Module Summary

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

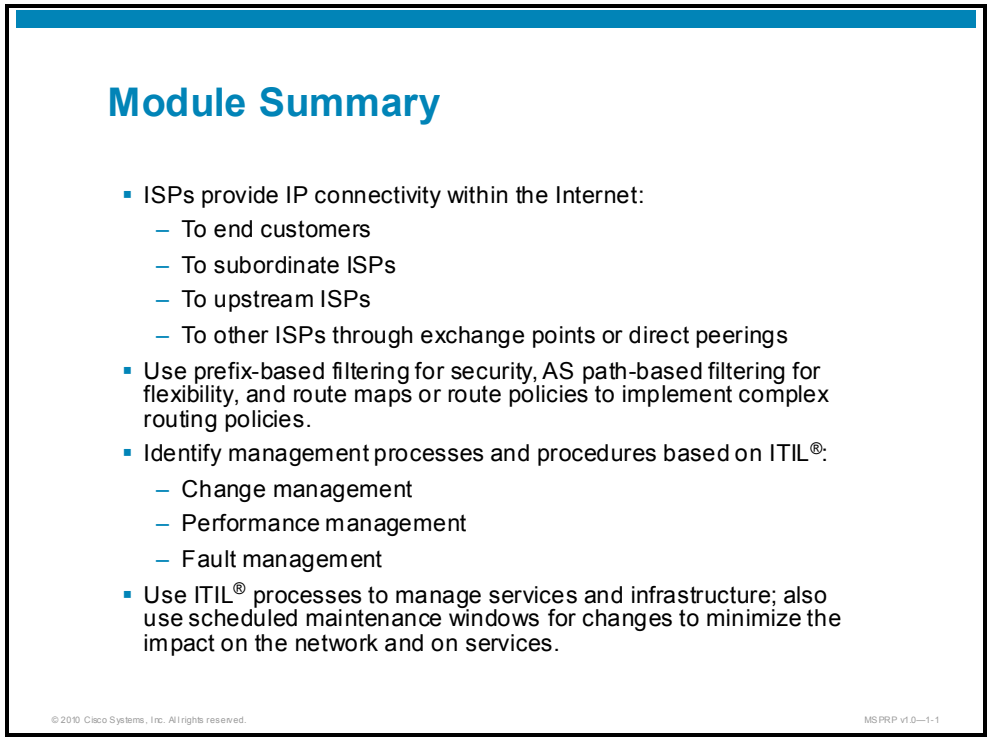

This module identified service provider routing requirements, solutions, and processes for change, performance, and fault management. It discussed typical routing requirements in service provider networks. The module described how an ISP provides IP connectivity within the Internet to customer and other ISPs. Routing solutions were listed and described based on typical examples. This module also presented the basics of prefix-based filtering, AS path-based filtering, and route maps or route policies. The module advised the use of ITIL® processes to manage services and infrastructure. You should also use scheduled maintenance windows for changes to minimize the impact on network and services.

