

ROUTE

Implementing Cisco IP Routing

Volume 2

Version 1.0

Student Guide

Text Part Number: 97-2815-02



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.



Cisco and the Cisco Logo are trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and other countries. A listing of Cisco's trademarks can be found at www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1005R)

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

Table of Contents

Volume 2

<u>Implementing a Scalable Multiarea Network OSPF-Based Solution</u>	3-1
Overview	3-1
Module Objectives	3-1
<u>Planning Routing Implementations with OSPF as the Scalable Routing Protocol</u>	3-3
Overview	3-3
Objectives	3-3
Link-State Routing Protocols	3-4
OSPF Area Structure	3-6
OSPF Adjacency Database	3-9
Calculating the OSPF Metric	3-10
Link-State Data Structures	3-12
Maintaining Link-State Sequence Numbers	3-15
Example: LSA Sequence Numbers and Maximum Age	3-17
Planning for Implementing OSPF	3-18
Summary	3-23
<u>How OSPF Packet Processes Work</u>	3-25
Overview	3-25
Objectives	3-25
OSPF Packet Types	3-26
Establishing OSPF Neighbor Adjacencies	3-29
Exchanging and Synchronizing LSDBs	3-31
Maintaining Network Routes	3-36
Verifying Packet Flow	3-38
Example: debug ip ospf packet	3-38
Summary	3-40
<u>Improving Routing Performance in a Complex Enterprise Network</u>	3-41
Overview	3-41
Objectives	3-41
Introducing OSPF Network Types	3-43
Adjacency Behavior in Point-to-Point Links	3-44
Adjacency Behavior in a Broadcast Network	3-45
Adjacency Behavior in a Metro Ethernet and EoMPLS Network	3-46
Adjacency Behavior in MPLS Networks	3-48
Selecting a DR and BDR	3-50
OSPF over Different Frame Relay Implementations	3-53
OSPF over Frame Relay NBMA	3-58
Example: NBMA Configuration	3-62
Example: show ip ospf neighbor Command	3-63
Using Subinterfaces in OSPF over Frame Relay	3-64
Example: Point-to-Point Subinterface	3-66
Example: Multipoint Subinterface	3-68
Implementing OSPF over a Point-to-Point Frame Relay Network	3-69
Example: Point-to-Point Configuration	3-70
Implementing OSPF over a Point-to-Multipoint Frame Relay Network	3-72
Example: Point-to-Multipoint Configuration	3-74
Example: OSPF over NBMA Topology Summary	3-77
Summary	3-78

Configuring and Verifying OSPF Routing	3-79
Overview	3-79
Objectives	3-79
Initializing Single-Area and Multiarea OSPF	3-80
Activating OSPF Within a Router in Single-Area and Multiarea	3-82
Defining the Router ID	3-85
Verify OSPF Operations	3-90
Identifying LSA Types Within the LSDB	3-97
Type 1	3-97
Type 2	3-97
Types 3 and 4	3-97
Type 5	3-98
Type 6	3-98
Type 7	3-98
Type 8	3-98
Types 9, 10, and 11	3-98
Type 1 LSA Link Types	3-99
Example: LSA Type 4—ASBR Summary LSA	3-102
Example: Types of LSAs	3-105
Limiting Adjacencies in OSPF with the passive-interface Command	3-125
Design Limitations of OSPF	3-127
OSPF Virtual Links and Solutions to Discontiguous Area Problems	3-129
Changing the Cost Metric	3-135
Summary	3-137
Lab 3-1 Debrief	3-139
Overview	3-139
Objectives	3-139
Lab Overview and Verification	3-140
Sample Solution and Alternatives	3-146
Summary	3-149
Lab 3-2 Debrief	3-151
Overview	3-151
Objectives	3-151
Lab Overview and Verification	3-152
Sample Solution and Alternatives	3-156
Summary	3-159
Configuring and Verifying OSPF Route Summarization	3-161
Overview	3-161
Objectives	3-161
OSPF Route Summarization	3-162
Implementing OSPF Route Summarization	3-164
Benefits of a Default Route in OSPF	3-171
Using a Default Route in OSPF	3-172
Summary	3-174
Lab 3-3 Debrief	3-175
Overview	3-175
Objectives	3-175
Lab Overview and Verification	3-176
Sample Solution and Alternatives	3-180
Summary	3-183

Configuring and Verifying OSPF Special Area Types	3-185
Overview	3-185
Objectives	3-185
OSPF Area Types	3-186
Defining and Implementing a Stub Area	3-190
Defining and Implementing a Totally Stubby Area	3-193
Example: Totally Stubby Configuration	3-195
Interpreting Routing Tables for OSPF Area Types	3-196
Defining and Implementing an NSSA and Totally NSSA in OSPF	3-200
Verifying All OSPF Area Types	3-206
Summary	3-207
Lab 3-4 Debrief	3-209
Overview	3-209
Objectives	3-209
Lab Overview and Verification	3-210
Sample Solution and Alternatives	3-215
Summary	3-218
Configuring and Verifying OSPF Authentication	3-219
Overview	3-219
Objectives	3-219
Types of OSPF Authentication	3-220
Implementing Simple Password Authentication for OSPF	3-221
Configuring MD5 Authentication for OSPF	3-226
Troubleshooting Authentication for OSPF	3-230
Summary	3-237
Lab 3-5 Debrief	3-239
Overview	3-239
Objectives	3-239
Lab Overview and Verification	3-240
Sample Solution and Alternatives	3-244
Summary	3-247
Module Summary	3-249
Module Self-Check	3-251
Module Self-Check Answer Key	3-267
Implementing an IPv4-Based Redistribution Solution	4-1
Overview	4-1
Module Objectives	4-1
Assessing Network Routing Performance and Security Issues	4-3
Overview	4-3
Objectives	4-3
Common Network Performance Issues	4-5
How Distribute Lists Work	4-11
Using Distribute Lists to Control Routing Updates	4-13
How Prefix Lists Work	4-17
Using a Prefix List to Control Routing Updates	4-21
How Route Maps Work	4-25
Using Route Maps to Control Routing Updates	4-33
Using Route Maps to Filter Routes	4-36
Suppressing Routing Updates Using a Passive Interface	4-37
Summary	4-40

Operating a Network Using Multiple IP Routing Protocols	4-41
Overview	4-41
Objectives	4-41
Describe a Complex Routing Network	4-42
Defining Route Redistribution	4-45
Default Metrics for Redistributed Routes	4-49
Determining Where to Redistribute	4-53
Caveats of Redistribution	4-57
Summary	4-61
Configuring and Verifying Route Redistribution	4-63
Overview	4-63
Objectives	4-63
Examples of Redistribution	4-64
The Administrative Distance Attribute	4-76
Route Redistribution Using Administrative Distance	4-78
Impact of Using Administrative Distance for Route Manipulation	4-82
Redistribution Using Route Maps	4-87
Summary	4-93
Lab 4-1 Debrief	4-95
Overview	4-95
Objectives	4-95
Lab Overview and Verification	4-96
Sample Solution and Alternatives	4-100
Summary	4-103
Module Summary	4-105
Module Self-Check	4-107
Module Self-Check Answer Key	4-110

Implementing a Scalable Multiarea Network OSPF-Based Solution

Overview

This module examines Open Shortest Path First (OSPF), which is one of the most commonly used interior gateway protocols in IP networking. OSPF is an open-standard protocol that is based primarily on RFC 2328, with some enhancements for IP version 6 (IPv6) that are based on RFC 2740. OSPF is a complex protocol that is made up of several protocol handshakes, database advertisements, and packet types.

Configuration and verification of OSPF in a Cisco router is a primary learning objective of this module. The lessons move from simple to more advanced configuration topics. Each of the important OSPF commands is explained and described in an example, and all the important OSPF **show** commands are defined.

Module Objectives

Upon completing this module, you will be able to build a scalable multiarea network with OSPF. This ability includes being able to meet these objectives:

- Identify, analyze, and match OSPF multiarea routing functions and benefits for routing efficiency in a complex enterprise network
- Demonstrate how OSPF improves packet processes in a multiarea enterprise network
- Configure OSPF to improve routing performance in a complex enterprise network
- Discuss lab results for configuring and verifying OSPF to improve routing performance
- Given a network design, configure OSPF multiarea routing to improve network operations and reach performance expectations
- Discuss lab results for implementing and verifying OSPF multiarea routing
- Configure OSPF route summarization for interarea and external routes
- Discuss lab results for configuring and verifying OSPF route summarization for interarea and external routes

- Configure and verify the OSPF area parameters, including stub areas, not-so-stubby areas, and totally stubby areas
- Discuss lab results for configuring and verifying OSPF special area types
- Implement authentication in an OSPF network
- Discuss lab results for configuring and verifying OSPF authentication

Planning Routing Implementations with OSPF as the Scalable Routing Protocol

Overview

Open Shortest Path First (OSPF) is one of the most commonly used IP routing protocols in networking. It is an open standard that is used by both enterprise and service provider networks.

This lesson introduces each of the major characteristics of the OSPF routing protocol, including a description of link-state routing protocols, the OSPF hierarchical structure, link-state adjacencies, and Shortest Path First (SPF) calculations. The creation of the implementation plan, which is the first step in configuring the OSPF routing protocol, is also described.

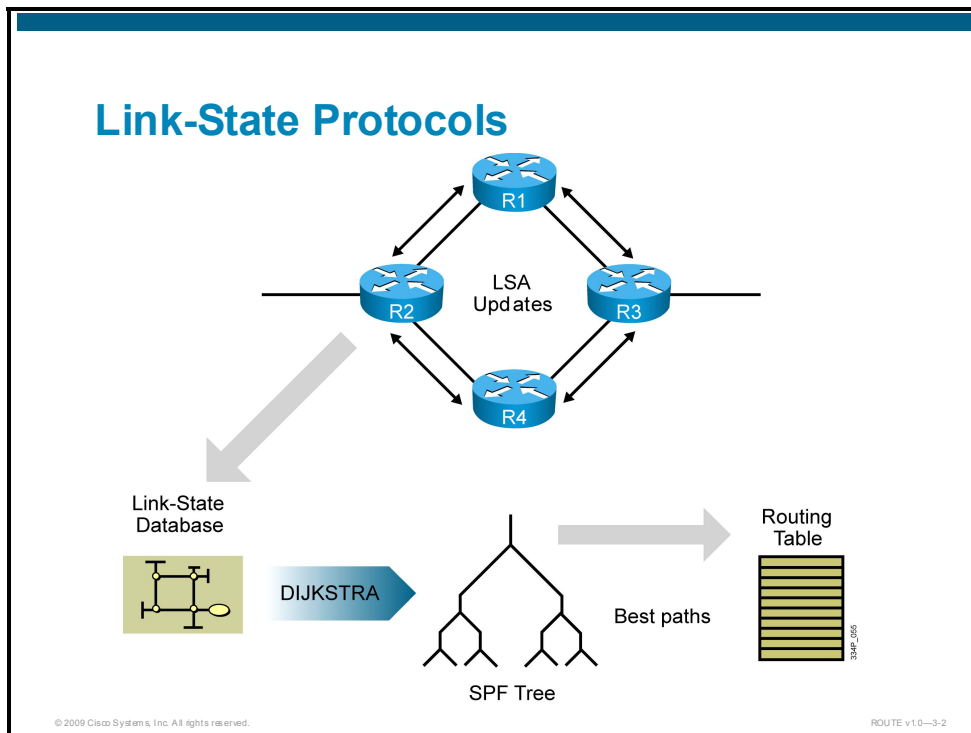
Objectives

Upon completing this lesson, you will be able to explain link-state protocols, components, the metrics of OSPF, the way in which OSPF operates, and the implementation plan creation process. This ability includes being able to meet these objectives:

- Determine link-state routing protocols
- Determine OSPF area structure
- Determine OSPF adjacency database
- Calculate the OSPF metric
- Determine link-state data structures
- Maintain link-state sequence numbers
- Plan for OSPF implementation

Link-State Routing Protocols

This topic describes the features of link-state routing protocols.



The need to overcome the limitations of distance vector routing protocols led to the development of link-state routing protocols. By using link-state advertisements (LSAs), each router builds its own view of the network and also maintains a list of neighbors, a list of all routers in the area, and a list of the best paths to each destination. OSPF is classified as a link-state routing protocol because of the manner in which it distributes routing information and calculates routes.

Link-state routing protocols generate routing updates only when a change occurs in the network topology. When a link changes state, the device that detected the change creates an LSA concerning that link.

The LSA is propagated to all neighboring devices using a special multicast address. Each routing device creates a copy of the LSA, updates its link-state database (LSDB), and forwards the LSA to all neighboring devices (within an area, as described later in this lesson). This flooding of the LSA ensures that all routing devices update their databases before updating routing tables to reflect the new topology.

The LSDB is used to calculate the best paths through the network. Link-state routers find the best paths to a destination by applying Dijkstra's algorithm, also known as SPF, against the LSDB to build the SPF tree. The best paths are then selected from the SPF tree and placed in the routing table.

Note The 224.0.0.5 multicast address is used by OSPF routers, and the 224.0.0.6 multicast address is used by OSPF designated routers (RFC 1583, JXM1).

Link-State Protocol Data Structures

- Link-state routers recognize more information about the network than their distance vector counterparts.
 - **Neighbor table:** also known as the adjacency database
 - **Topology table:** referred as the LSDB
 - **Routing table:** also known as the forwarding database
- Each router has a complete picture of the topology.
- Link-state routers tend to make more accurate decisions.

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-3-3

Link-state routing protocols collect routing information from all other routers in the network or from within a defined area of the network. When link-state routing protocols have collected this information from all routers, each router independently calculates its best path to all destinations in the network, using Dijkstra's algorithm.

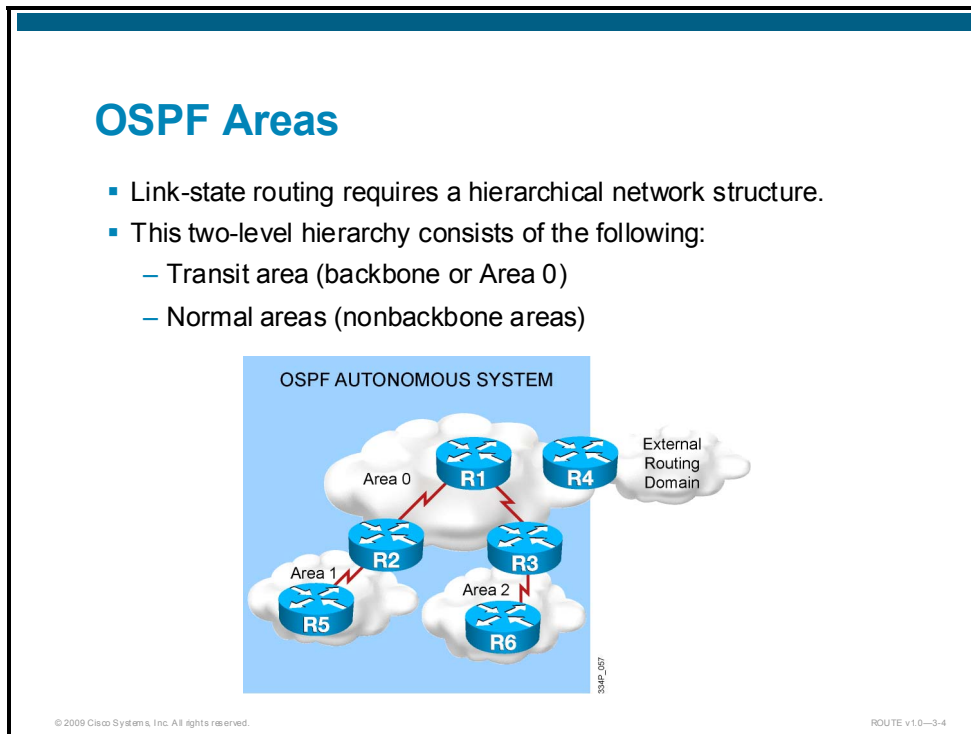
Incorrect information from any particular router is less likely to cause confusion, because each router maintains its own view of the network.

For consistent routing decisions to be taken by all the routers in the network, each router must keep a record of the following information:

- **Its immediate neighbor routers:** If the router loses contact with a neighboring router, within a few seconds, it will invalidate all paths through that router and recalculate its paths through the network. Adjacency information about neighbors is stored in the neighbor table, also known as an adjacency database, in OSPF.
- **All the other routers in the network, or in its area of the network, and their attached networks:** The router recognizes other routers and networks through LSAs, which are flooded through the network. LSAs are stored in a topology table, also called an LSDB. The LSDB is identical for all OSPF routers in an area. The memory resources that are needed to maintain these tables represent one drawback to link-state protocols.
- **The best path to each destination:** Each router independently calculates the best paths to each destination in the network, using Dijkstra's algorithm. The best paths are then offered to the routing table or forwarding database. Packets arriving at the router are forwarded based on the information that is held in the routing table. Each router is able to independently select a loop-free and efficient pathway. This benefit overcomes the "routing by rumors" limitation of distance vector routing.

OSPF Area Structure

This topic describes the two-tier hierarchy structure of OSPF, including the characteristics of transit areas and regular areas, as well as the terminology that is used.



In small networks, the web of router links is not complex, and paths to individual destinations are easily deduced. However, in large networks, the web is highly complex, and the number of potential paths to each destination is large. Therefore, the Dijkstra calculations that are comparing all these possible routes can be very complex and can take a significant amount of time to complete.

Link-state routing protocols usually reduce the size of the Dijkstra calculations by partitioning the network into areas. The number of routers in an area and the number of LSAs that flood within the area are small, which means that the link-state or topology database for an area is small. Consequently, the Dijkstra calculation is easier and takes less time. Routers that are inside an area maintain detailed information about the links and only general or summary information about routers and links in other areas.

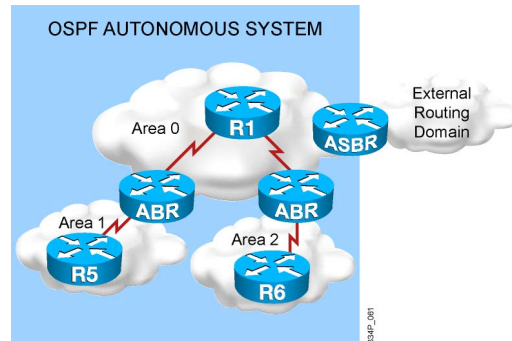
Link-state routing protocols use a two-layer area hierarchy:

- **Backbone or transit area:** The primary function of this OSPF area is to quickly and efficiently move IP packets. Backbone areas interconnect with other OSPF area types. The OSPF hierarchical area structure requires that all areas connect directly to the backbone area. In the figure, links between Area 1 and Area 2 routers are not allowed. Generally, end users are not found within a backbone area, which is also known as OSPF Area 0.
- **Normal or nonbackbone area:** The primary function of this OSPF area is to connect users and resources. Normal areas are usually set up according to functional or geographical groupings. By default, a normal area does not allow traffic from another area to use its links to reach other areas. All traffic from other areas must cross a transit area such as Area 0. Normal areas can be of different types.

The optimal number of routers per area varies based on factors such as network stability, but in the *Designing Large-Scale Internetworks* document, Cisco recommends that there generally be no more than 50 routers per area.

Area Terminology and Router Types

- **ABR:** Area Border Router
- **ASBR:** Autonomous System Boundary Router
- **R5 and R6:** Internal routers
- **R1:** Backbone router



All OSPF areas and routers that are running OSPF routing protocol comprise the OSPF autonomous system.

Routers that make up nonbackbone (normal) areas are known as internal routers; they have all interfaces in one area only.

Routers that make up Area 0 are known as backbone routers (internal routers in backbone). OSPF hierarchical networking defines Area 0 as the core. All other areas connect directly to backbone Area 0.

An Area Border Router (ABR) connects Area 0 to the nonbackbone areas. An OSPF ABR plays a very important role in network design and has interfaces in more than one area. An ABR has the following characteristics:

- It separates LSA flooding zones.
- It becomes the primary point for area address summarization.
- It functions regularly as the source for default routes.
- It maintains the LSDB for each area with which it is connected.

The ideal design is to have each ABR connected to two areas only, the backbone and another area, with three areas being the upper limit.

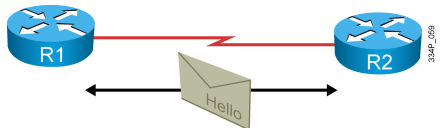
An Autonomous System Boundary Router (ASBR) connects any OSPF area to a different routing administration (such as Border Gateway Protocol [BGP] or Enhanced Interior Gateway Routing Protocol [EIGRP]). The ASBR is the point where external routes can be redistributed into OSPF.

OSPF Adjacency Database

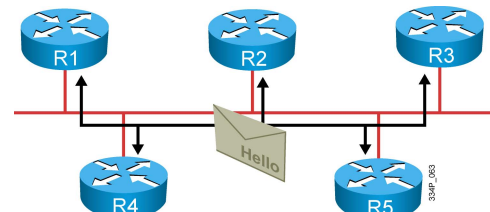
This topic describes how routers that are running a link-state routing protocol establish neighbor adjacencies with their neighboring routers.

OSPF Adjacencies

- Routing updates and topology information are passed only between adjacent routers.



- Forming OSPF adjacencies on point-to-point WAN links:



- Forming OSPF adjacencies on LAN links is different from forming them on point-to-point links.

© 2009 Cisco Systems, Inc. All rights reserved. ROUTE v1.0-3-6

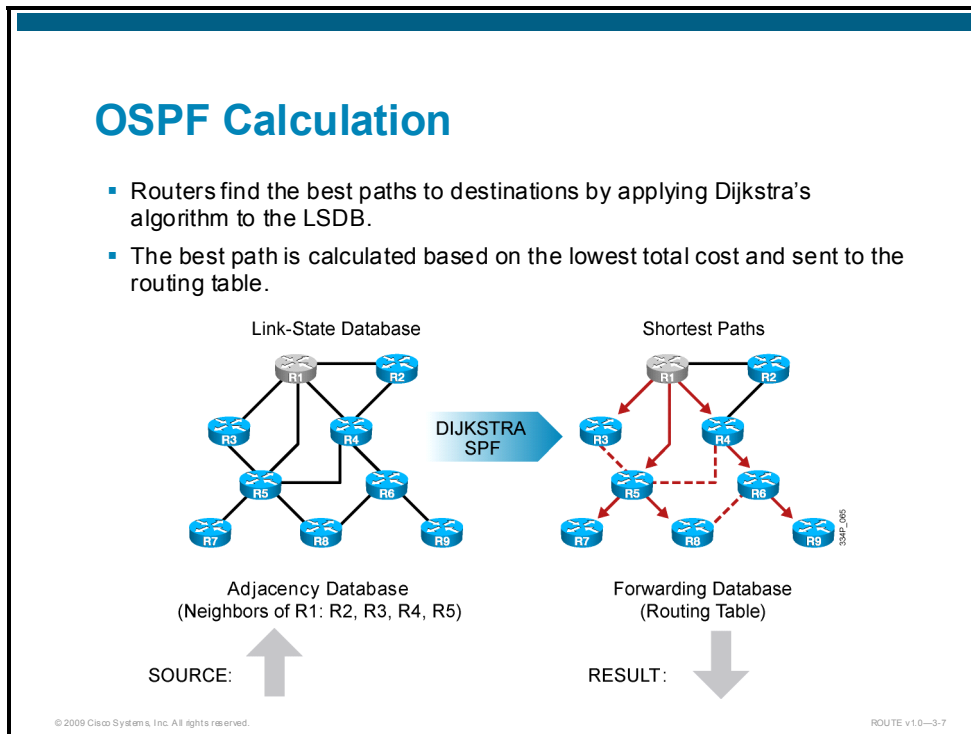
A router that is running a link-state routing protocol must first establish neighbor adjacencies with its neighboring routers. A router achieves this neighbor adjacency by exchanging hello packets with the neighboring routers.

Once routers become adjacent, they begin exchanging the link-state information to synchronize the LSDB. Link-state information must be synchronized between routers; the process is described later in this module. Only by reliably flooding link-state information can you ensure that every router in the area or domain has the latest, most accurate view of the network. Only then can the router make reliable routing decisions that are consistent with the decisions of other routers in the network.

The two OSPF routers on a point-to-point serial link, which are usually encapsulated in High-Level Data Link Control (HDLC) or PPP, form a complete adjacency with each other. OSPF routers on the LAN form adjacencies in a different way; the process is described later in this module.

Calculating the OSPF Metric

This topic describes the SPF algorithm and the calculations that are used by OSPF to find the best path to each destination network. The routing table is then populated with the best paths.



All routers form an adjacency, which is the prerequisite for the LSDB creation process. Once they have all formed adjacencies, they start exchanging LSAs. Router R1 has four neighboring routers: R2, R3, R4, and R5. From these routers, R1 receives the LSAs from all other routers in the network. From these LSAs, it can also deduce the links between all routers and draw the web of routers that are depicted in the figure. In this way, an LSDB is created.

Edsger Dijkstra designed a mathematical algorithm for calculating the best paths through complex networks. Link-state routing protocols use Dijkstra's algorithm to calculate the best paths through a network. They assign a cost to each link in the network and place the specific node at the root of a tree, then sum the costs toward each given destination. In this way, they calculate the branches of the tree to determine the best path to each destination. The best path is calculated with respect to the lowest total cost of links to a specific destination and is put in the forwarding database (routing table). For OSPF, the default behavior is that the interface cost is calculated based on its configured bandwidth. You can also manually define an OSPF cost for each interface, which overrides the default cost value.

The figure illustrates an example of a Dijkstra calculation. Each Ethernet link in the figure is assigned an OSPF cost of 10. By summing the costs to each destination, the router can deduce the best path to each destination. The right side of the figure shows the result of the Dijkstra calculation, in which the SPF tree is defining the best paths. From these best paths, which are shown with solid lines, routes to destination networks that are attached to each router are offered to the routing table. For each route, the next-hop address is the appropriate neighboring router (R2, R3, R4, or R5).

OSPF Metric

- Also called “cost”
- Defined per interface, but may be altered
- Inversely proportional to the bandwidth of that interface
- Cost = 100,000,000 / bandwidth in b/s

Link Type	Default Cost
64-kb/s serial link	1562
T1 (1.544-Mb/s serial link)	64
E1 (2.048-Mb/s serial link)	48
Ethernet	10
Fast Ethernet	1
ATM	1

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v10-3-8

The cost (also called the metric) of an interface in OSPF is an indication of the overhead that is required to send packets across a certain interface. The cost of an interface is inversely proportional to the bandwidth of that interface. A higher bandwidth indicates a lower cost. There is more overhead (higher cost) and time delay in crossing a 56-kb/s serial line than in crossing a 10-Mb/s Ethernet line. The formula that is used to calculate the cost is:

$$\text{Cost} = 100,000,000 / \text{bandwidth in b/s}$$

For example, it will cost $10^8/10^7 = 10$ to cross a 10-Mb Ethernet line and will cost $10^8/1544000 = 64$ to cross a T1 line.

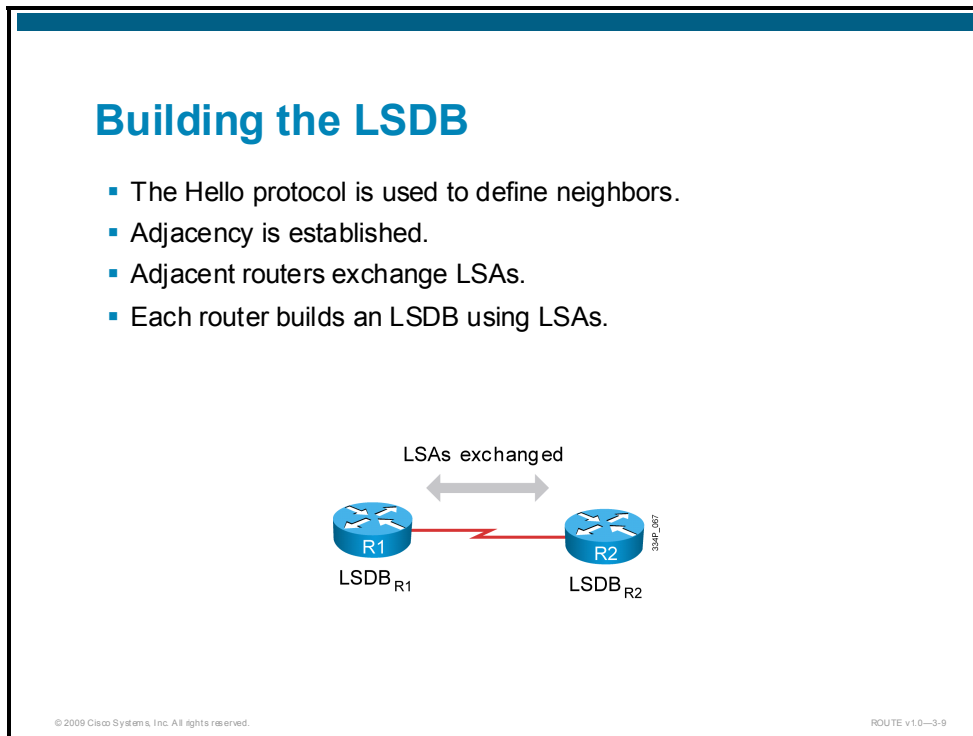
By default, the cost of an interface is calculated based on the bandwidth and can be changed by using the OSPF configuration command. If you change the link bandwidth, it will also indirectly change the cost. Only one cost can be assigned per interface, and it is advertised as the link cost in the router link advertisements.

Links have different default costs, as follows:

- **A 56-kb/s serial link:** The default cost is 1785.
- **A 64-kb/s serial link:** The default cost is 1562.
- **T1 (1.544-Mb/s serial link):** The default cost is 64.
- **E1 (2.048-Mb/s serial link):** The default cost is 48.
- **Ethernet:** The default cost is 10.
- **Fast Ethernet:** The default cost is 1.
- **FDDI:** The default cost is 1.
- **ATM:** The default cost is 1.

Link-State Data Structures

This topic describes how routers use link-state updates (LSUs) to verify that links are still active.



Routers that share a common segment become neighbors on that segment. Neighbors are elected via the Hello protocol. Hello packets are sent periodically out of each interface; routers become neighbors as soon as they see themselves listed in the neighbor hello packet. Therefore, a two-way communication is guaranteed. Two routers will not become neighbors unless they agree on the following:

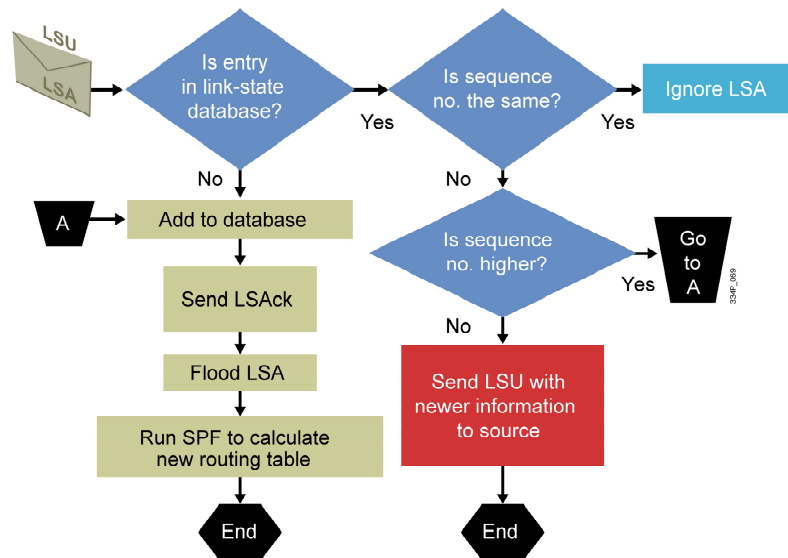
- **Area ID:** Two routers having a common segment; their interfaces have to belong to the same area on that segment.
- **Authentication:** OSPF allows for the configuration of a password for a specific area. Routers that want to become neighbors have to exchange the same password on a particular segment.
- **Hello and dead intervals:** OSPF exchanges hello packets on each segment. This is a form of keepalive that is used by routers to acknowledge their existence on a segment and to elect a designated router (DR) on multiaccess segments. OSPF requires these intervals to be exactly the same between two neighbors.
- **Stub area flag:** Two routers have to agree on the stub area flag in the hello packets to become neighbors. Stub areas will be discussed in a later section. Keep in mind that defining stub areas will affect the neighbor election process.

Adjacency is the next step after the process of becoming neighbors. Adjacent routers are routers that go beyond the simple hello exchange and proceed into the database exchange process. LSU packets are exchanged, and acknowledgement is required. Selection and processing of LSUs is based on the sequence numbers. At the end, the adjacent routers will have the same link-state database.

The following is a brief summary of the states that an interface passes through before becoming adjacent to another router:

- Down
- Attempt
- INIT
- Two-way
- Exstart
- Exchange
- Loading
- Full

Link-State Data Structures: LSA Operation



Each LSA entry has its own aging timer, which the link-state age field carries. The default timer value for OSPF is 30 minutes (expressed in seconds in the link-state age field).

After an LSA entry ages, the router that originated the entry sends the LSA (with a higher sequence number, in an LSU) to verify that the link is still active. The LSU can contain one or more LSAs. This LSA validation method saves on bandwidth compared with distance vector routers, each of which sends its entire routing table at short intervals.

When each router receives the LSU, it does the following:

- If the LSA does not already exist, the router adds the entry to its LSDB, sends back a link-state acknowledgment (LSAck), floods the information to other routers, runs SPF, and updates its routing table.
- If the entry already exists and the received LSA has the same sequence number, the router ignores the LSA entry.
- If the entry already exists but the LSA includes newer information (it has a higher sequence number), the router adds the entry to its LSDB, sends back an LSAck, floods the information to other routers, runs SPF, and updates its routing table.
- If the entry already exists but the LSA includes older information, it sends an LSU to the sender with its newer information.

Maintaining Link-State Sequence Numbers

This topic describes the process of maintaining a database of only the most recent link-state sequence numbers.

Defining the “More Recent” LSA

An LSA is more recent if it has:

- A higher sequence number
- A higher checksum number
- An age that is equal to the maximum age (poisoning)
- A significantly smaller link-state age (the LSA is significantly younger)

A combination of the maximum age (max-age) and refresh timers, as well as link-state sequence numbers, helps OSPF maintain a database of only the most recent link-state records.

An LSA is more recent if it has:

- A higher sequence number
- A higher checksum number
- An age that is equal to the maximum age (poisoning)
- A significantly smaller link-state age (the LSA is significantly younger)

LSA Sequence Numbering

Each LSA in the LSDB maintains a sequence number:

- A 4-byte number
- Begins with 0x80000001; ends with 0x7FFFFFFF

OSPF floods each LSA every 30 minutes:

- Each time, the sequence number is incremented by one.
- The LSA with the higher (newer) sequence number is more recent.

Ultimately, a sequence number will wrap around to 0x80000001:

- The existing LSA was prematurely aged to the maximum age (one hour) and flushed.

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-3-12

The link-state sequence number field in an LSA header is 32 bits long (4 bytes). Beginning with the leftmost bit set, the first legal sequence number is 0x80000001 and the last one is 0x7FFFFFFF. The number is incremented in each new LSA until it reaches 0xFFFFFFFF. The next number is then 0x00000000, followed by 0x00000001, 0x00000002, and so on. 0x80000000 is reserved and never used. This number field is used to detect old or redundant LSAs.

To ensure an accurate database, OSPF floods (refreshes) each LSA every 30 minutes. Each time that a record is flooded, the sequence number is incremented by one. An LSA record will reset its maximum age when it receives a new LSA update. An LSA will never remain in the database longer than the maximum age of one hour without a refresh.

It is possible for an LSA to remain in the database for long periods of time, getting refreshed every 30 minutes. At some point, the sequence number will need to wrap back to the starting sequence number. When this occurs, the existing LSA will be prematurely aged out (the maximum age timer is immediately set to one hour) and flushed. The LSA will then begin its sequencing at 0x80000001 again.

When a router encounters two instances of an LSA, it must determine which is more recent. The LSA with the higher link-state sequence number is the more recent LSA.

LSA Sequence Numbers and Maximum Age

- Every OSPF router announces a router LSA for those interfaces that it owns in that area.
- Router with link ID 192.168.1.67 has been updated eight times; the last update was 48 seconds ago.

```
R1#show ip ospf database
      OSPF Router with ID (192.168.1.67) (Process ID 10)
      Router Link States (Area 1)
Link ID      ADV Router      Age  Seq#      Checksum  Link count
192.168.1.67 192.168.1.67    48  0x80000008  0xB112    2
192.168.2.130 192.168.2.130  212 0x80000006  0x3F44    2
<output omitted>
```

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-3-13

Example: LSA Sequence Numbers and Maximum Age

The **show ip ospf database** command displays lists of information that are related to the OSPF database for a specific router. The output of the command that is shown in the figure provides an example of how the link-state age and LSA sequence numbers are kept in the database.

Every OSPF router has interfaces in one or more areas and announces a router LSA for those interfaces that it owns in those areas. The link ID is the ID of the router that created the router LSA. The advertising router (shown as “ADV Router” in the output) is the router ID of the OSPF router that announced the router LSA. Generally, the link ID and advertising router for a router LSA are the same.

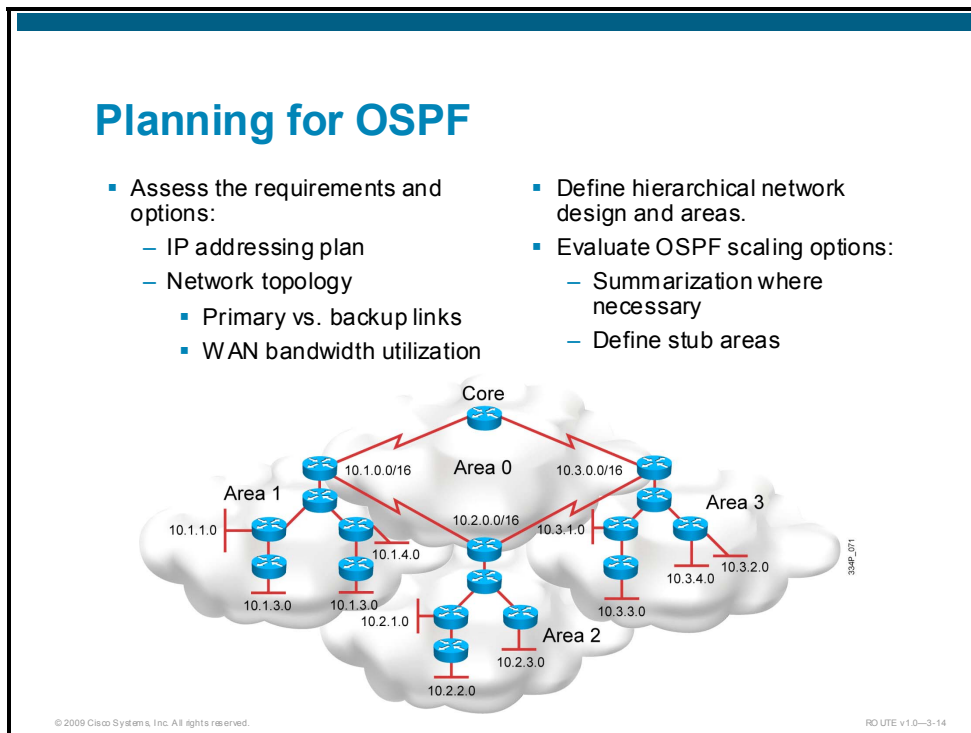
The first router LSA entry in the OSPF database indicates that the router LSA with link ID 192.168.1.67 has been updated eight times (because the sequence number is 0x80000008) and that the last update occurred 48 seconds ago.

For more details about the **show ip ospf database** command, go to the Cisco IOS IP Routing: OSPF Command Reference via the following link:

http://www.cisco.com/en/US/docs/ios/iproute_ospf/command/reference/iro_book.html

Planning for Implementing OSPF

This topic describes how to plan, implement, and document the OSPF deployment.



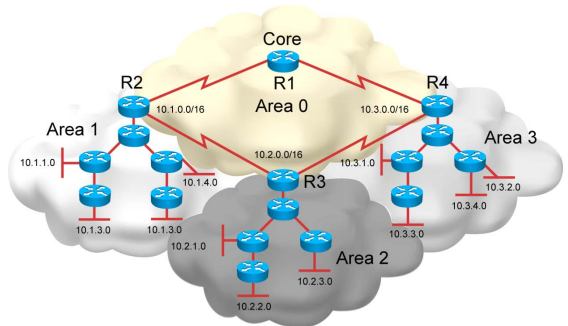
The OSPF routing protocol implementation depends on specific needs and topologies. When preparing to deploy OSPF routing in a network, you must gather the existing state and requirements and also consider different deployment considerations as follows:

- The IP addressing plan governs how OSPF can be deployed and how well the OSPF deployment might scale. Thus, a detailed IP addressing plan, along with IP subnetting information, must be defined. A solid IP addressing plan should enable usage of OSPF summarization to scale the network and to optimize OSPF behavior more easily.
- A network topology consists of links connecting the network equipment (routers, switches, and so on). A detailed network topology plan should be presented to assess OSPF scalability requirements and determine which OSPF features might be required (for example, OSPF summarization and redistribution).
- OSPF can control the size of the LSDB by using different areas, which decrease the LSDB size and limit the propagation of LSUs in case of a topology change within one area.

OSPF Implementation Plan

- Verify and configure IP addressing.
- Enable OSPF for the correct interfaces.
- Enable OSPF for the correct areas.
- Define special metric to influence path selection.
- Verify the configuration.

Area	Router	Interface
0	R1	S0/0, S0/1
0	R2	S0/0, S0/1
1	R2	S0/2
0	R3	S0/0, S0/1
2	R3	S0/2
0	R4	S0/0, S0/1
3	R4	S0/2



Once you have assessed the requirements, you can create the implementation plan. To implement OSPF routing, you must do the following:

- Learn IP addressing—or more precisely, the networks that need to be included and advertised by OSPF.
- Enable the OSPF process for the correct interfaces or the correct network statements under the OSPF routing process configuration mode.
- The OSPF process must be enabled for the correct areas on the interfaces.
- Any specific metric that needs to be applied to certain interfaces to influence the default best path selection by OSPF routing protocol must be known. The table should include the required metric along with the interface where the metric needs to be applied.

When the implementation plan is created, the list of tasks for each router in the network must be defined as follows:


- Enable the OSPF routing protocol on an interface, or use the correct network statements under the OSPF routing process configuration mode.
- Assign the correct area ID to the interface through the OSPF configuration under the interface or under the OSPF routing process configuration mode.
- You can also apply the metric to proper interfaces.

After implementation, you should perform verification to confirm that OSPF has been deployed properly on each router:

- Verify the setup of OSPF adjacency.
- Verify that the OSPF LSDB is populated with the proper information.
- Verify that the IP routing table is populated with the proper information.
- Verify that there is connectivity in the network.
- Verify that OSPF behaves as expected when a topology change occurs (test link failure and router failure events).

Documenting OSPF

- Topology
- Areas and IP addressing
- Networks and interfaces included in OSPF per router
- Default and nondefault metrics applied
- Configuration and verification results




© 2009 Cisco Systems, Inc. All rights reserved.

RO UTE v1.0—3-16

After a successful OSPF deployment, you should document the solution and then put the information about the deployment in a safe place. The implementation plan itself is half of the information. To complete the documentation, you must also include the verification process and results.

Example: Planning for Basic OSPF

- Define the network requirements.
- Gather the required parameters.
- Define the OSPF areas and routing.
- Configure basic OSPF.
- Verify the OSPF configuration.
- Complete the documentation.



The example in the figure illustrates the steps that are necessary to prepare an implementation plan to configure basic OSPF.

When you are planning for the basic OSPF configuration, you should ensure that the implementation plan includes the following steps:

- Define the network requirements.
- Gather the required parameters.
- Define OSPF routing.
- Configure basic OSPF.
- Verify the OSPF configuration.

Summary

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

Summary

- Link-state routing protocols respond quickly to changes, send triggered updates when changes occur, and send periodic updates every 30 minutes.
- A two-tier hierarchical network structure is used by OSPF in which the network is divided into areas. This area structure is used to separate the LSDB into more manageable pieces.
- Adjacencies are built by OSPF routers using the Hello protocol. LSUs are sent over these logical adjacencies to exchange database information between adjacent OSPF routers.
- Dijkstra's algorithm, also known as SPF, is used to calculate best paths for all destinations. SPF is run against the LSDB, and the result is a table of best paths, known as the routing table.

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—3-18

Summary (Cont.)

- After an LSA entry ages, the router that originated the entry sends an LSU about the network to verify that the link is still active. The LSU can contain one or more LSAs.
- Each LSA in the LSDB has a sequence number, which is incremented by one each time that the LSA is flooded. When a router encounters two instances of an LSA, it must determine which is more recent. The LSA with the higher LSA sequence number is the more recent.
- When planning an OSPF deployment, define the network requirements, gather the required parameters, and define the OSPF routing.

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—3-19

How OSPF Packet Processes Work

Overview

The Open Shortest Path First (OSPF) protocol has several functions and five packet types to support them: hello, database description (DBD), link-state request (LSR), link-state update (LSU), and link-state acknowledgement (LSAck). The five OSPF packet types enable all OSPF information flow between routers. This lesson defines each packet type and explains where and how these packets interact to build OSPF neighbor adjacencies and maintain the OSPF topology database.

Objectives

Upon completing this lesson, you will be able to explain how information flows between routers to maintain OSPF links and which packets are used. This ability includes being able to meet these objectives:

- Determine OSPF packet types
- Establish OSPF neighbor adjacencies
- Exchange and synchronize LSDBs
- Maintain network routes
- Verify packet flow

OSPF Packet Types

This topic describes the five OSPF packet types.

OSPF Functions

High-level functions of OSPF include the following:

- Discover neighbors and form adjacencies.
- Flood link-state database (LSDB) information.
- Compute the shortest path.
- Install routes in the route-forwarding table.

Additional functions of OSPF include the following:

- Detect changes in the link state.
- Propagate changes to maintain LSDB synchronization.

Several OSPF packet types are involved.

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0--9-2

For OSPF to operate properly, several processes must occur:

- Neighbor discovery to form adjacencies
- Flooding of the link-state information to build a link-state database (LSDB)
- Computation of Shortest Path First (SPF) to find out the shortest path to all known destinations
- Populating of the route forwarding table with the best routes to known destinations

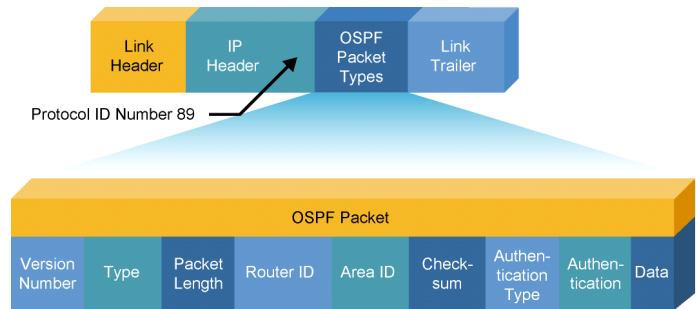
Once OSPF initially populates the router forwarding table, the state of the links that are around the OSPF autonomous system may change. OSPF is able to detect these changes and respond by flooding this information throughout the OSPF autonomous system, or at least in the area where the change was detected. The flooding of new information is needed to maintain the LSDBs in all neighboring routers.

For all these functions, several OSPF packet types are involved and will be described in this lesson.

OSPF Packet Header Format

```
R1#debug ip ospf packet
OSPF packet debugging is on
R1#
*Feb 16 11:03:51.206: OSPF: rcv. v:2 t:1 l:48 rid:10.0.0.12
aid:0.0.0.1 chk:D882 aut:0 auk: from Serial0/0/0.2
```

- This debug output shows fields in the OSPF header.



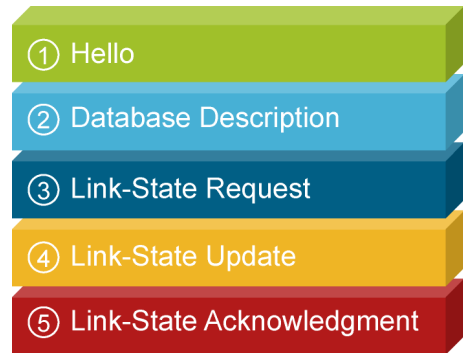
All five OSPF packets are encapsulated directly into an IP payload, as shown in the figure. The OSPF packet does not use TCP or User Datagram Protocol (UDP). OSPF requires a reliable packet transport scheme, and because it does not use TCP, it has defined its own acknowledgment routine using an acknowledgment packet (OSPF packet type 5).

In the IP header, a protocol identifier of 89 defines all OSPF packets. Each of the OSPF packets begins with the same header format. This header has the following fields:

- **Version number:** Version 2 for OSPF with IPv4 and version 3 for OSPF with IPv6
- **Type:** Differentiates the five OSPF packet types
- **Packet length:** The length of the OSPF packet in bytes
- **Router ID:** Defines which router is the source of the packet
- **Area ID:** Defines the area where the packet originated
- **Checksum:** Used for packet-header error detection to ensure that the OSPF packet was not corrupted during transmission
- **Authentication type:** An option in OSPF that describes no authentication, cleartext passwords, or encrypted Message Digest 5 (MD5) formats for router authentication
- **Authentication:** Used in the authentication scheme
- **Data:** Each of the five packet types includes different data:
 - **Hello packets:** Contains a list of known neighbors
 - **DBD packet:** Contains a summary of the LSDB, which includes all known router IDs and their last sequence numbers, among a number of other fields
 - **LSR packet:** Contains the type of LSU that is needed and the router ID that has the needed LSU
 - **LSU packet:** Contains the complete link-state advertisement (LSA) entries. Multiple LSA entries can fit in one OSPF update packet.
 - **LSAck packet:** Empty

OSPF Packet Types

- OSPF uses five types of routing protocol packets.



© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—3-4

OSPF uses five types of routing protocol packets, which share a common protocol header. The protocol ID number inside the IP header is 89, and all five packet types are used in the normal operation of OSPF.

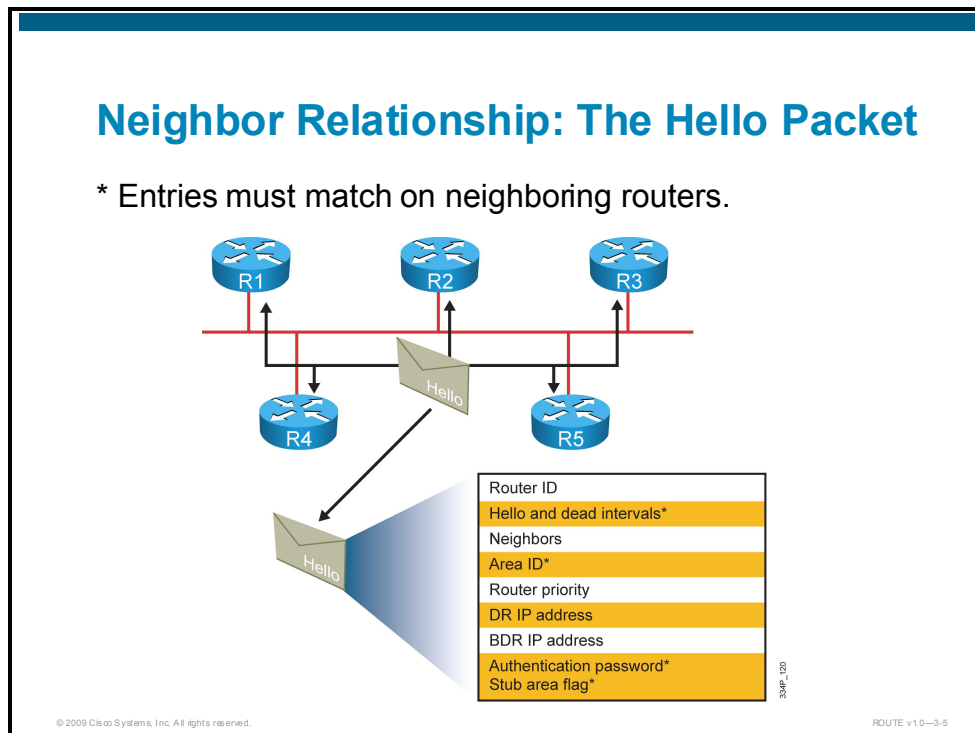
OSPF Packets

The table contains descriptions of each type.

Type	Packet Name	Description
1	Hello	Discovers neighbors and builds adjacencies between them
2	DBD	Checks for database synchronization between routers
3	LSR	Requests specific link-state records from another router
4	LSU	Sends specifically requested link-state records
5	LSAck	Acknowledges the other packet types

Establishing OSPF Neighbor Adjacencies

This topic describes how OSPF neighbor adjacencies are established.



Each interface that is participating in OSPF uses IP multicast address 224.0.0.5 to send hello packets periodically. A hello packet contains the following information:

- **Router ID:** The router ID is a 32-bit number that uniquely identifies the router. The highest IP address on an active interface is chosen by default, unless a loopback interface or the router ID is configured; for example, IP address 172.16.12.1 would be chosen over 172.16.1.1. This identification is important in establishing neighbor relationships and coordinating LSU exchanges. Also, the router ID breaks ties during the designated router (DR) and backup designated router (BDR) selection processes if the OSPF priority values are equal.
- **Hello and dead intervals:** The hello interval specifies the frequency, in seconds, with which a router sends hello packets (10 seconds is the default on multiaccess networks). The dead interval is the time in seconds that a router waits to hear from a neighbor before declaring the neighboring router out of service (four times the hello interval, by default). These timers must be the same on neighboring routers; otherwise, an adjacency will not be established.
- **Neighbors:** The neighbors field lists the adjacent routers with established bidirectional communication. This bidirectional communication is indicated when the router recognizes itself listed in the neighbors field of the hello packet from the neighbor.
- **Area ID:** To communicate, two routers must share a common segment, and their interfaces must belong to the same OSPF area on that segment. (They must also share the same subnet and mask.) These routers will all have the same link-state information.
- **Router priority:** The router priority is an 8-bit number that indicates the priority of a router. Priority is used when selecting a DR and BDR.

- **DR and BDR IP addresses:** These are the IP addresses of the DR and BDR for the specific network, if they are known.
- **Authentication password:** If router authentication is enabled, two routers must exchange the same password. Authentication is not required, but if it is enabled, all peer routers must have the same password.
- **Stub area flag:** A stub area is a special area. Two routers must agree on the stub area flag in the hello packets. Designating a stub area is a technique that reduces number of routing updates by replacing many of them with a default route.

Note The following fields must match when hello packets are exchanged between the neighboring routers: hello and dead intervals, area ID, authentication password, and stub area flag.

Exchanging and Synchronizing LSDBs

Once a bidirectional adjacency is formed, OSPF must exchange and synchronize the LSDBs between routers. This topic describes the process of exchanging and synchronizing the LSDBs between routers.

OSPF Routing Update Packets

- Use of multicast and unicast IP address
- Four types of update packets
- LSDB synchronization process:
 - Discover neighbor.
 - Establish bidirectional communication.
 - Elect a DR, if desired.
 - Form an adjacency.
 - Discover the network routes.
 - Update and synchronize LSDBs.

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-3-6

All routing updates are sent in IP packets (protocol type 89), for which OSPF does not do any fragmentation and reassembly on the IP.

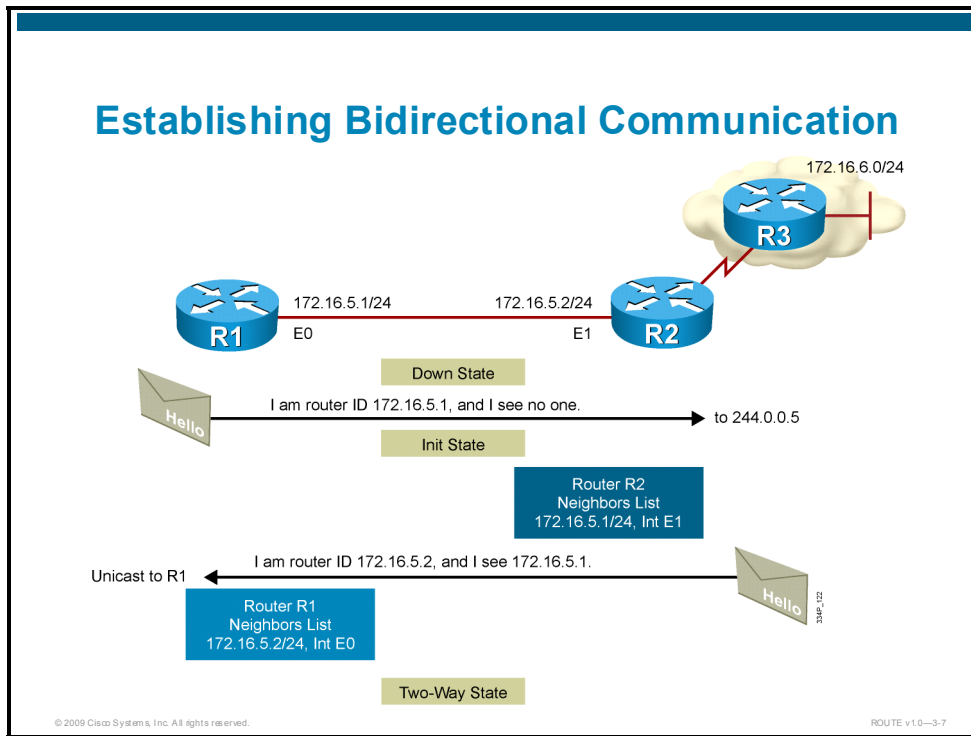
Four types of update packets are used when exchanging and synchronizing LSDBs:

- **Type 2 DBD packet:** Used to describe the network routes of each neighbor.
- **Type 3 LSR packet:** After database description packets are exchanged, the routers request missing information by using request packets.
- **Type 4 LSU packet:** All missing information is sent to the neighbors by sending update packets, which contain different LSAs.
- **Type 5 LSAck packet:** Every packet is acknowledged, to ensure reliable transport and reliable exchange of information.

Type 4 and type 5 packets are sent to multicast IP addresses, except when retransmitting, when sent across a virtual link, and when sent on nonbroadcast networks. All other packets are sent to unicast IP addresses.

The LSDB synchronization process starts with neighbor discovery. After a neighbor has been discovered, bidirectional communication commences. A designated router may also be elected (for example, if the OSPF protocol is active on the LAN). Then, routers make a decision whether an adjacency is formed. If an adjacency is to be formed, neighbors have to synchronize LSDBs. First, network routes are discovered, and then missing information is exchanged. Finally, LSDB is synchronized, as described in the following pages.

Establishing Bidirectional Communication



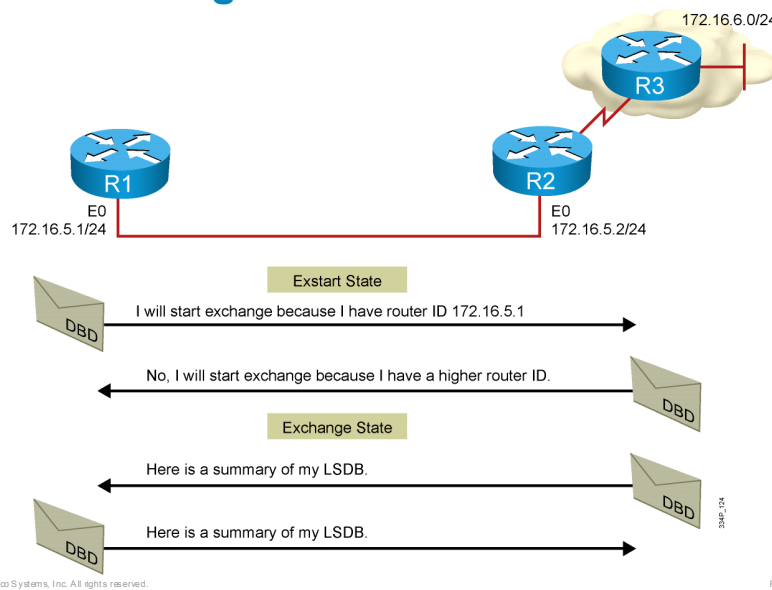
When routers that are running OSPF are initialized, an exchange process using the Hello protocol is the first procedure. The exchange process that happens when routers appear on the network is illustrated in the figure:

1. Router R1 is enabled on the LAN and is in a down state, because it has not exchanged information with any other router. It begins by sending a hello packet through each of its interfaces that are participating in OSPF, even though it does not know the identity of the DR or of any other routers. The hello packet is sent out using the multicast address 224.0.0.5.
2. All directly connected routers that are running OSPF receive the hello packet from router R1 and add R1 to their lists of neighbors. After adding R1 to the list, other routers are in the initial state (INIT state).
3. Each router that received the hello packet sends a unicast reply hello packet to R1 with its corresponding information. The neighbor field in the hello packet includes all neighboring routers and R1.
4. When R1 receives these hello packets, it adds all the routers that had its router ID in their hello packets to its own neighbor relationship database. After this process, R1 is in the two-way state. At this point, all routers that have each other in their lists of neighbors have established bidirectional communication.

If the link type is a broadcast network—generally, a LAN link like Ethernet—a DR and BDR must first be selected. The DR forms bidirectional adjacencies with all other routers on the LAN link. This process must occur before the routers can begin exchanging link-state information. Periodically (every 10 seconds on broadcast networks, by default), the routers within a network exchange hello packets to ensure that communication is still working. The hello updates include the DR, BDR, and the list of routers for which hello packets have been received by the router. Remember that “received” means that the receiving router recognizes its own name as one of the entries in the received hello packet.

Note The concepts of DR and BDR will be described later in this module.

Discovering the Network Routes



After the DR and BDR have been selected (applied to LAN link types), the routers are considered to be in the exstart state. The routers are then ready to discover the link-state information about the internetwork and create their LSDBs. The exchange protocol is used to discover the network routes, and it gets all the routers from exchange state to a full state of communication. The first step in this process is for the DR and BDR to establish adjacencies with each of the other routers. When adjacent routers are in a full state, they do not repeat the exchange protocol unless the full state changes.

As shown in the figure, the exchange protocol operates as follows:

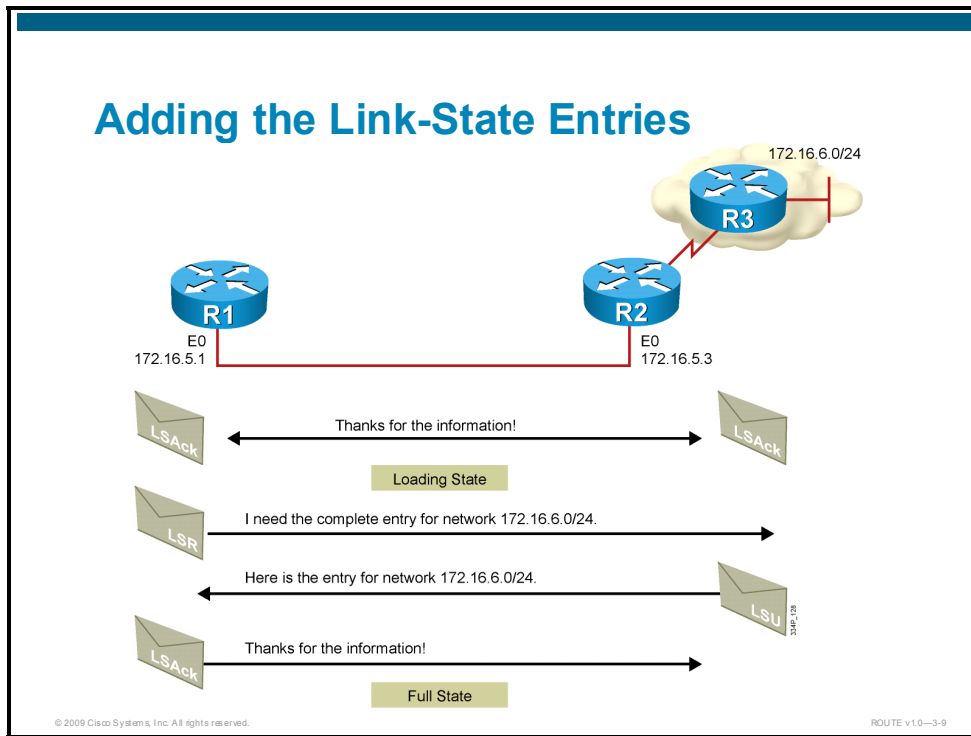
Step 1 In the exstart state, the DR and BDR establish adjacencies with each router in the network. During this process, a master-slave relationship is created between each router and its adjacent DR and BDR. The router with the higher router ID acts as the master during the exchange process—router R2 becomes the DR.

Note Only the DR exchanges and synchronizes link-state information with the routers to which it has established adjacencies. Having the DR represent the network in this capacity reduces the amount of routing update traffic.

The master and slave routers exchange one or more DBD packets. The routers are in the exchange state.

A DBD includes information about the LSA entry header that appears in the LSDB of the router. The entries can be about a link or a network. Each LSA entry header includes information about the link-state type, the address of the advertising router, the cost of the link, and the sequence number. The router uses the sequence number to determine the “newness” of the received link-state information.

Adding the Link-State Entries



Step 2 When the router receives the DBD, it performs these actions, as shown in the figure:

1. It acknowledges the receipt of the DBD using the LSAck packet.
2. It compares the information that it received with the information that it has. If the DBD has a more up-to-date link-state entry, the router sends an LSR to the other router. When routers start sending LSRs, they are in the loading state.
3. The other router responds with the complete information about the requested entry in an LSU packet. Again, when the router receives an LSU, it sends an LSAck.

Step 3 The router adds the new link-state entries to its LSDB.

When all LSRs have been satisfied for a given router, the adjacent routers are considered synchronized. They are in a full state. The routers must be in a full state before they can route traffic. At this point, all the routers in the area should have identical LSDBs.

OSPF Neighbor States

OSPF routers progress through seven states:

- **Down:** no active neighbor detected
- **INIT:** hello packet received
- **Two-way:** own router ID in received hello
- **Exstart:** master and slave roles determined
- **Exchange:** database description packets sent
- **Loading:** exchange of LSRs and LSUs
- **Full:** neighbors fully adjacent

© 2009 Cisco Systems, Inc. All rights reserved.

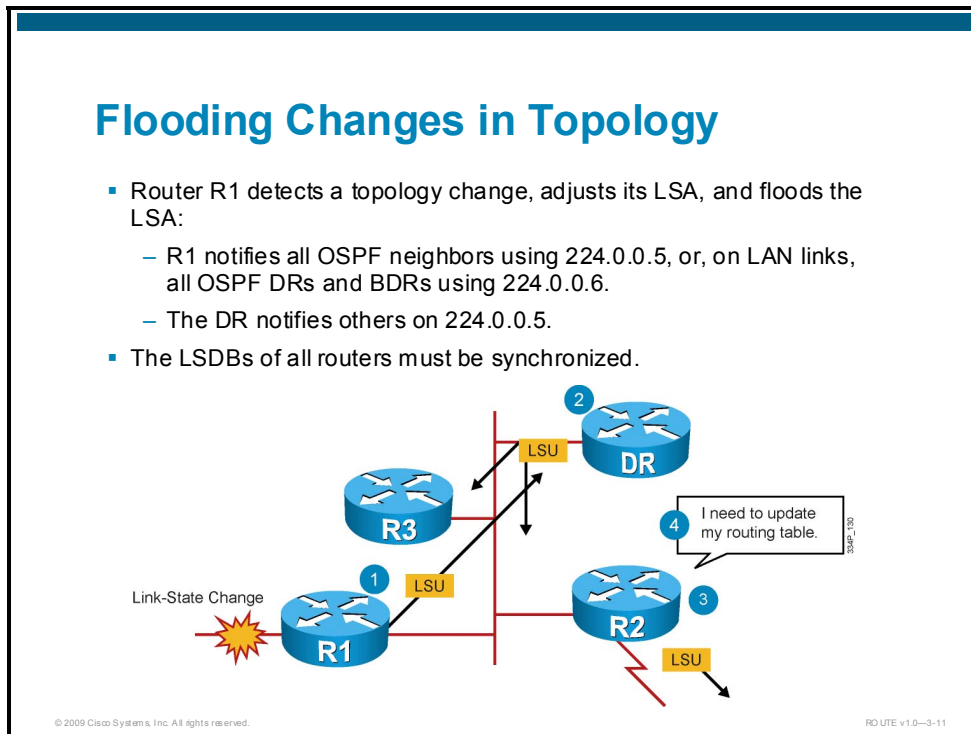
RO UTE v1.0-3-10

The following is a brief summary of the states that an interface passes through before becoming adjacent to another router:

- **Down:** No information has been received on the segment.
- **Attempt:** In nonbroadcast multiaccess clouds such as Frame Relay and X.25, this state indicates that no recent information has been received from the neighbor. An effort should be made to contact the neighbor by sending hello packets at the reduced rate poll interval.
- **INIT:** The interface has detected a hello packet coming from a neighbor, but bidirectional communication has not yet been established.
- **Two-way:** There is bidirectional communication with a neighbor. The router has seen itself in the hello packets coming from a neighbor. At the end of this stage, the DR and BDR election would have been done. When routers are in the two-way state, they must decide whether to proceed in building an adjacency. The decision is based on whether one of the routers is a DR or BDR or the link is a point-to-point or a virtual link.
- **Exstart:** Routers are trying to establish the initial sequence number that is going to be used in the information exchange packets. The sequence number ensures that routers always get the most recent information. One router will become the primary and the other will become secondary. The primary router will poll the secondary for information.
- **Exchange:** Routers will describe their entire LSDB by sending database description packets. In this state, packets may be flooded to other interfaces on the router.
- **Loading:** In this state, routers are finalizing the information exchange. Routers have built a link-state request list and a link-state retransmission list. Any information that looks incomplete or outdated will be put on the request list. Any update that is sent will be put on the retransmission list until it gets acknowledged.
- **Full:** In this state, adjacency is complete. The neighboring routers are fully adjacent. Adjacent routers will have similar LSDBs.

Maintaining Network Routes

This topic describes how OSPF maintains synchronization of the LSDBs (topology tables) of all routers in the network.



In a link-state routing environment, it is very important for the LSDBs (topology tables) of all routers to stay synchronized. When there is a change in a link state, the routers use a flooding process to notify the other routers in the network of the change. LSUs provide the mechanism for flooding LSAs.

A router that detects a topology change adjusts its LSA, increments its LSA sequence number, and floods the LSA. Only the LSA originator can adjust or regenerate the LSA. A router that receives an LSA update packet compares the received LSA to its own copy. If the received LSA is more recent, it is installed in the database and flooded. If the database copy is newer, it is sent back to the sending neighbor. This is an exception to the rule, because a router is not supposed to reoriginate the LSAs of other routers.

Periodic LSA refreshment is another mechanism to maintain the LSDB. Each LSA has an age field. LSAs are aged when flooded and in the topology database. LSAs that reach the maximum age are purged from the topology database (garbage collection). Each router has to resend its LSA (the LSA that was originated from that router) before it ages out in the topology databases. LSAs with an age that is equal to the maximum age are used to immediately remove LSAs from the database (route poisoning).

In general, the flooding process steps in a multiaccess network are as follows:

- Step 1** A router notices a change in a link state and multicasts an LSU packet (which includes the updated LSA entry) to all OSPF neighbors (on point-to-point links) at 224.0.0.5 or to all OSPF DRs and BDRs (on LAN links like the one in the figure) at 224.0.0.6. An LSU packet may contain several distinct LSAs.

- Step 2** The DR acknowledges receipt of the change and floods the LSU to others on the network using the OSPF multicast address 224.0.0.5. After receiving the LSU, each router responds to the DR with an LSAck. To make the flooding procedure reliable, each LSA must be acknowledged separately.
- Step 3** If a router is connected to other networks, it floods the LSU to those other networks by forwarding the LSU to the DR of the multiaccess network (or to the adjacent router if it is in a point-to-point network). The DR, in turn, multicasts the LSU to the other routers in the network.
- Step 4** The router updates its LSDB using the LSU that includes the changed LSA. It then recomputes the SPF algorithm against the updated database after a short delay (the SPF delay) and updates the routing table as necessary.

Note Although it is not shown in this figure, all LSUs are acknowledged.

When is the SPF algorithm run? A change in the topology database is a necessary, but not sufficient, condition for SPF recalculation. The SPF algorithm is triggered if:

- The LSA Options field has changed.
- The age of the LSA instance is set to MaxAge.
- The length field in the LSA header has changed.
- The contents of the LSA (excluding the LSA header) have changed.

An SPF calculation is performed separately for each area in the topology database.

OSPF simplifies the synchronization issue by requiring only adjacent routers to remain synchronized.

Summaries of individual link-state entries, not the complete link-state entries, are sent every 30 minutes to ensure proper LSDB synchronization. Each link-state entry has a timer to determine when the LSA refresh update must be sent.

Each link-state entry also has a maximum age of 60 minutes. If a link-state entry has not been refreshed within 60 minutes, it is removed from the LSDB.

Note On a Cisco router, if a route already exists, the routing table is used at the same time that the SPF algorithm is calculating a new best route to the destination. However, if the SPF is calculating a new route, the new route is used only after the SPF calculation is complete.

Verifying Packet Flow

This topic describes how to verify that OSPF packets are flowing properly between two routers.

The debug ip ospf packet Command

Debugging a single packet:

```
R1#debug ip ospf packet
OSPF packet debugging is on
R1#
*Feb 16 11:03:51.206: OSPF: rcv. v:2 t:1 l:48 rid:10.0.0.12
      aid:0.0.0.1 chk:D882 aut:0 auk: from Serial0/0/0.2
```

- This debug output shows the fields in the OSPF header.

© 2009 Cisco Systems, Inc. All rights reserved. ROUTE v1.0-3-12

The **debug ip ospf packet** command is used to verify that OSPF packets are flowing properly between two routers as well as for troubleshooting.

Example: debug ip ospf packet

The output of this debug command is shown in the figure. Notice that the output shows the fields in the OSPF header, but they are not described in any detail. After each field, there is a parameter for it. The values are described in the table that follows.

From the following sample output and table, you can say the following:

```
OSPF: rcv. v:2 t:1 l:48 rid:10.0.0.12
      aid:0.0.0.1 chk:D882 aut:0 auk: from Serial0/0/0.2
```

OSPF:	OSPF packet
rcv.	was received
v:2	OSPF version 2
t:1	Hello packet
l:48	48 bytes is the OSPF packet length.
rid:10.0.0.12	The OSPF router ID is 10.0.0.12.
aid:0.0.0.1	The OSPF area ID is 0.0.0.1.
chk:D882	The OSPF checksum is D882.
aut:0	Authentication is not being used.
auk:	There is no key; authentication is not being used.

from Serial0/0/0.2

The packet was received via interface Serial0/0/0.2.

OSPF Packet Header Fields

The table lists the OSPF packet header fields that are represented in this output.

Field	Description
v	Provides the version of OSPF
t	Specifies the OSPF packet type: 1: hello 2: DBD 3: LSR 4: LSU 5: LSAck
l	Specifies the OSPF packet length in bytes
rid	Provides the OSPF router ID
aid	Shows the OSPF area ID
chk	Displays the OSPF checksum
aut	Provides the OSPF authentication type: 0: No authentication 1: Simple password 2: MD5
auk	Specifies the OSPF authentication key, if used
keyid	Displays the MD5 key ID; only used for MD5 authentication
seq	Provides the sequence number; only used for MD5 authentication

For more details about the **debug ip ospf packet** command, go to the Cisco IOS Debug Command Reference via the following link:

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

Summary

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

Summary

- There are five OSPF packet types: hello, DBD, LSU, LSR, and LSAck.
- The Hello protocol forms logical neighbor adjacency relationships. A DR may be required to coordinate adjacency formations.
- The exchange protocol passes through several states (down, INIT, two-way, exstart, exchange, and loading) before finally reaching the goal of the full state. When the protocol is in the full state, its databases are synchronized with adjacent routers.
- LSAs are sent when a change occurs, but are also sent every 30 minutes to ensure database integrity. The maximum time that an LSA will stay in the database, without an update, is 1 hour.
- Use the **debug ip ospf packet** command to verify that OSPF packets are flowing properly between two routers.

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-3-13

Improving Routing Performance in a Complex Enterprise Network

Overview

Open Shortest Path First (OSPF) protocol areas may be made up of different types of network links. You must know which types of network links are being used in order to properly configure OSPF to work with all of them and their different adjacency behaviors, as well as over certain network types.

It is important to note that OSPF pays special attention to different network types, such as point-to-point and broadcast networks, and that the OSPF default settings do not always work properly with some network topologies.

This lesson describes each OSPF network type, how the adjacencies are formed for these OSPF network types, and how link-state advertisements (LSAs) are flooded on each.

Objectives

Upon completing this lesson, you will be able to describe the features of various OSPF network architectures. This ability includes being able to meet these objectives:

- Introduce OSPF network types
- Determine adjacency behavior in point-to-point links
- Determine adjacency behavior in a broadcast network
- Determine adjacency behavior in a Metro Ethernet and EoMPLS network
- Determine adjacency behavior in MPLS networks
- Select a DR and BDR
- Implement OSPF over different Frame Relay implementations
- Implement OSPF over Frame Relay NBMA
- Use subinterfaces in OSPF over Frame Relay

- Implement OSPF over a point-to-point Frame Relay network
- Implement OSPF over a point-to-multipoint Frame Relay network

Introducing OSPF Network Types

This topic describes the three types of networks that are defined by OSPF.

OSPF Network Types

- **Point-to-point:** A network that joins a single pair of routers
- **Broadcast:** A multiaccess broadcast network, such as Ethernet
- **Nonbroadcast multiaccess (NBMA):** A network that interconnects more than two routers but has no broadcast capability
 - Examples: Frame Relay, ATM, and X.25
 - Five modes of OSPF operation available for NBMA networks

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-3-2

OSPF defines distinct types of networks, based on their physical link types. OSPF operation on each type is different, including how adjacencies are established and which configuration is required.

There are three types of networks that are defined by OSPF:

- **Point-to-point:** A network that joins a single pair of routers.
- **Broadcast:** A multiaccess broadcast network, such as Ethernet.
- **Nonbroadcast multiaccess (NBMA):** A network that interconnects more than two routers but has no broadcast capability. Frame Relay, ATM, and X.25 are examples of NBMA networks. There are five modes of OSPF operation that are available for NBMA networks, as described later in this lesson.

Adjacency Behavior in Point-to-Point Links

This topic describes the adjacency behavior for point-to-point serial links.

Point-to-Point Links

- Usually, a serial interface running either PPP or HDLC
- May also be a point-to-point subinterface running Frame Relay or ATM
- Does not require DR or BDR election
- Is automatically detected by OSPF
- Sends OSPF packets using multicast 224.0.0.5



© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0--3-3

A point-to-point network joins a single pair of routers. A T1 serial line that is configured with a data link layer protocol such as PPP or High-Level Data Link Control (HDLC) is an example of a point-to-point network.

On point-to-point networks, the router dynamically detects its neighboring routers by multicasting its hello packets to all OSPF routers, using the address 224.0.0.5. On point-to-point networks, neighboring routers become adjacent whenever they can communicate directly. No designated router (DR) or backup designated router (BDR) election is performed; there can be only two routers on a point-to-point link, so there is no need for a DR or BDR.

Usually, the IP source address of an OSPF packet is set to the address of the outgoing interface on the router. It is possible to use IP unnumbered interfaces with OSPF. On unnumbered interfaces, the IP source address is set to the IP address of another interface on the router.


The default OSPF hello and dead intervals on point-to-point links are 10 seconds and 40 seconds, respectively.

Adjacency Behavior in a Broadcast Network

This topic describes the adjacency behavior for a broadcast network link.

Multiaccess Broadcast Network

- This generally applies to LAN technologies like Ethernet.
- DR and BDR selection are required.
- All neighbor routers form adjacencies with the DR and BDR only.
- Packets to the DR and the BDR use 224.0.0.6.
- Packets from DR to all other routers use 224.0.0.5.



© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-3-4

An OSPF router on a multiaccess broadcast network such as Ethernet forms an adjacency with its DR and BDR. Adjacent routers have synchronized link-state databases (LSDBs). A common media segment is the basis for adjacency, such as an Ethernet segment that is connecting two routers. When routers first come up on the Ethernet, they perform the hello process and then elect the DR and BDR. The routers then attempt to form adjacencies with the DR and BDR.

The routers on a segment must elect a DR and a BDR to represent the multiaccess broadcast network. The BDR does not perform any DR functions when the DR is operating. Instead, the BDR receives all the information, but the DR performs the LSA forwarding and LSDB synchronization tasks. The BDR performs the DR tasks only if the DR fails. If the DR fails, the BDR automatically becomes the DR, and a new BDR election occurs.

The DR and BDR improve network functioning in the following ways:

- **Reducing routing update traffic:** The DR and BDR act as central point of contact for link-state information exchange on a multiaccess broadcast network; therefore, each router must establish a full adjacency with the DR and the BDR only. Each router, rather than exchanging link-state information with every other router on the segment, sends the link-state information to the DR and BDR only. The DR represents the multiaccess broadcast network in the sense that it sends link-state information from each router to all other routers in the network. This flooding process significantly reduces the router-related traffic on a segment.
- **Managing link-state synchronization:** The DR and BDR ensure that the other routers on the network have the same link-state information about the internetwork. In this way, the DR and BDR reduce the number of routing errors.


Note After a DR and BDR have been selected, any router that is added to the network establishes adjacencies with the DR and BDR only.

Adjacency Behavior in a Metro Ethernet and EoMPLS Network

This topic describes the adjacency behavior in a Metro Ethernet and Ethernet over Multiprotocol Label Switching (EoMPLS) network link.

OSPF Adjacency Over Metro Ethernet and EoMPLS

- EoMPLS and Metro Ethernet service does not participate in STP, nor does it learn MAC addresses.
- Customer routers R1 and R2 exchange Ethernet frames via an interface or VLAN subinterfaces.
- OSPF behaves the same as on Ethernet:
 - OSPF network type = multiaccess broadcast network
 - DR and BDR are elected
 - Routers form adjacencies with the DR and BDR only



© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-1.5

The Multiprotocol Label Switching (MPLS) backbone of the service provider is used to enable Layer 2 Ethernet connectivity between the customer routers R1 and R2, whether an EoMPLS or Metro Ethernet service is used.

R1 and R2 thus exchange Ethernet frames. Provider edge router (PE router) PE1 takes the Ethernet frames that are received from R1 on the link to PE1, encapsulates them into MPLS packets, and forwards them across the backbone to router PE2. PE2 de-encapsulates the MPLS packets and reproduces the Ethernet frames on the link toward R2. EoMPLS and Metro Ethernet typically do not participate in Spanning Tree Protocol (STP) and bridge protocol data unit (BPDU) exchanges, so EoMPLS and Metro Ethernet are transparent to the customer routers.

The Ethernet frames are transparently exchanged across the MPLS backbone. Keep in mind that customer routers can be connected either in a port-to-port fashion, in which PE routers take whatever Ethernet frame is received and forward the frames across the Layer 2 MPLS virtual private network (VPN) backbone, or in a VLAN subinterface fashion, in which frames for a particular VLAN—identified with subinterface in configuration—are encapsulated and sent across the Layer 2 MPLS VPN backbone.

When deploying OSPF over EoMPLS, there are no changes to the existing OSPF configuration from the customer perspective.

OSPF needs to be enabled, and network commands must include the interfaces that are required by the relevant OSPF area, to start the OSPF properly.

R1 and R2 form a neighbor relationship with each other over the Layer 2 MPLS VPN backbone. From an OSPF perspective, the Layer 2 MPLS VPN backbone, PE1, and PE2 are all invisible.


A neighbor relationship is established between R1 and R2 directly, and it behaves in the same way as on a regular Ethernet broadcast network.

Adjacency Behavior in MPLS Networks

This topic describes the adjacency behavior in an MPLS network link.

OSPF Adjacency Over MPLS VPN

- Customer routers run OSPF and exchange routing updates with the PE routers:
 - PE routers appear as another router in the customer network.
 - Service provider P routers are hidden from the customer.
 - Customer routers are unaware of MPLS VPN.
 - Customer and service provider must agree on OSPF parameters.
- Customer routers to PE connection can be of any type:
 - OSPF behaves per the connection type (point-to-point, broadcast, NBMA).



© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—1.6

With the Layer 3 MPLS VPN architecture, the ISP provides a peer-to-peer VPN architecture. In this architecture, PE routers participate in customer routing, guaranteeing optimum routing between customer sites. Therefore, the PE routers carry a separate set of routes for each customer, resulting in perfect isolation between the customers. The following applies to the Layer 3 MPLS VPN technology, even when running OSPF as a provider edge-customer edge (PE-CE) routing protocol:

- The customer routers should not be aware of MPLS VPN; they should run standard IP routing software.
- The provider routers (P routers) must not carry VPN routes for the MPLS VPN solution to be scalable.
- The PE routers must support MPLS VPN services and traditional Internet services.

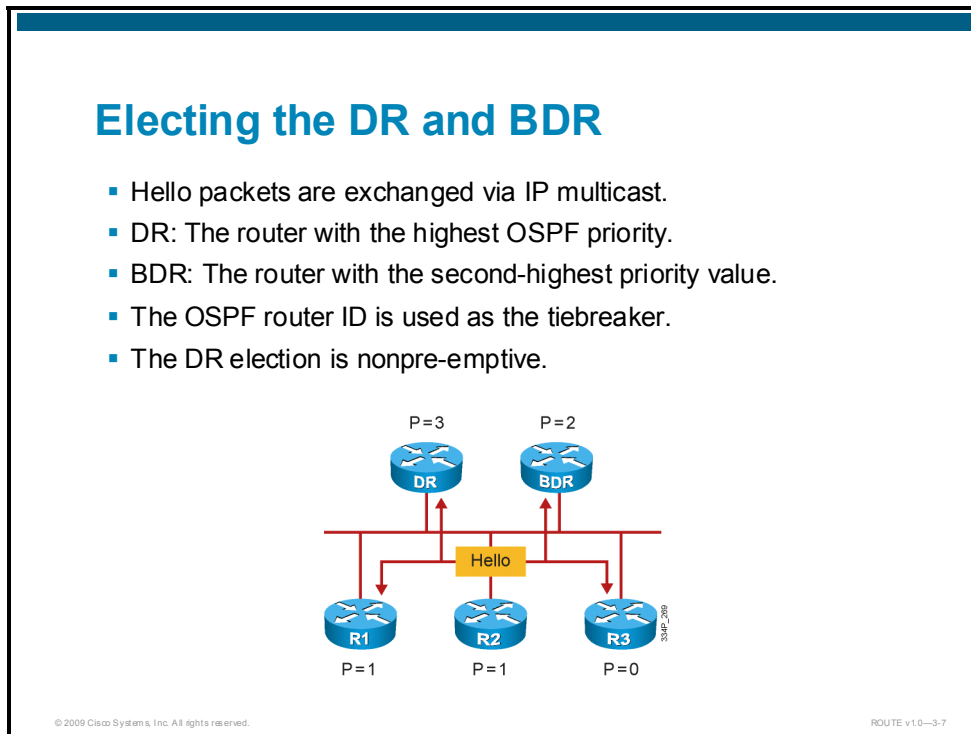
To OSPF, the Layer 3 MPLS VPN backbone looks like a standard corporate backbone, and it runs standard IP routing software. Routing updates are exchanged between the customer routers and the PE routers that appear as normal routers in the customer network. OSPF is enabled on proper interfaces using the **network** command. The standard design rules that are used for enterprise Layer 3 MPLS VPN backbones can be applied to the design of the customer network. The P routers are hidden from the customer view, and customer edge routers (CE routers) are unaware of MPLS VPN. Therefore, the internal topology of the Layer 3 MPLS backbone is totally transparent to the customer. The PE routers receive IP version 4 (IPv4) routing updates from the CE routers and install them in the appropriate virtual routing and forwarding (VRF) table. This part of the configuration and operation is the responsibility of a service provider.

The PE-CE can have any OSPF network type—point-to-point, broadcast, or even nonbroadcast multiaccess.

The only difference between a PE-CE design and a regular OSPF design is that the customer has to agree with the service provider about the OSPF parameters (autonomous system [AS] number, authentication password, and so on); usually, these parameters are governed by the service provider.

Selecting a DR and BDR

This topic describes how OSPF routers apply conditions to the OSPF priority values of other routers during the hello packet exchange process to elect a DR and BDR.



To elect a DR and BDR, the routers view the OSPF priority value of other routers during the hello packet exchange process and then use the following conditions to determine which router to select:

- The router with the highest priority value is the DR.
- The router with the second-highest priority value is the BDR.
- The default for the interface OSPF priority is 1. In the case of a tie, the router ID is used. The router with the highest router ID becomes the DR. The router with the second-highest router ID becomes the BDR.
- A router with a priority that is set to 0 cannot become the DR or BDR. A router that is not the DR or BDR is called a DROTHER.
- If a router with a higher priority value is added to the network, it does not pre-empt the DR and BDR. The only time that a DR or BDR changes is when one of them is out of service. If the DR is out of service, the BDR becomes the DR, and a new BDR is selected. If the BDR is out of service, a new BDR is elected.

To determine whether the DR is out of service, the BDR uses the wait timer. This timer is a reliability feature. If the BDR does not confirm that the DR is forwarding LSAs before the timer expires, the BDR assumes that the DR is out of service.

Note The highest IP address on an active interface is normally used as the router ID; however, you can override this selection by configuring an IP address on a loopback interface or using the **router-id** router configuration command.

In a multiaccess broadcast environment, each network segment has its own DR and BDR. A router that is connected to multiple multiaccess broadcast networks can be a DR on one segment and a regular router on another segment.

Note The DR operation is at the link level; a DR is selected for every multiaccess broadcast link in the OSPF network.

Setting the Priority for DR Election

```
DR (config-if) #
```

```
ip ospf priority 3
```

- This interface configuration command assigns the OSPF priority to an interface.
- Different interfaces on a router may be assigned different values.
- The default priority is 1. The range is from 0 to 255.
- “0” means the router cannot be the DR or BDR.
- A router that is not the DR or BDR is DROTHER.

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—3-8

Use the **ip ospf priority** command to designate which router interfaces on a multiaccess link are the DR and the BDR. The default priority is 1, and the range is from 0 to 255. The interface with the highest priority becomes the DR, and the interface with the second-highest priority becomes the BDR.

Interfaces that are set to a priority of 0 cannot be involved in the DR or BDR election process.

Note The priority of an interface takes effect only when the existing DR goes down. A DR does not relinquish its status just because a new interface reports a higher priority in its hello packet.

For more details about the **ip ospf priority** command, go to the Cisco IOS IP Routing: OSPF Command Reference via the following link:


http://www.cisco.com/en/US/docs/ios/iproute_ospf/command/reference/iro_book.html

OSPF over Different Frame Relay Implementations

This topic describes the adjacency behavior for a Frame Relay NBMA network.

NBMA Topology

- A single interface interconnects multiple sites.
- NBMA networks support multiple routers but have no broadcasting capabilities.
- Five modes of OSPF operation are available.



The diagram illustrates a Non-Broadcast Multiple Access (NBMA) network topology. It features four blue router icons arranged in a square, with red lines connecting each router to a central cloud. The cloud is labeled 'X.25 Frame Relay ATM'. The routers are arranged in two rows of two. The bottom-right router has the text 'SMP,271' next to it. The diagram is enclosed in a blue-bordered box.

© 2009 Cisco Systems, Inc. All rights reserved. ROUTE v1.0-3-9

When a single interface interconnects multiple sites over an NBMA network, the nonbroadcast nature of the network can lead to reachability issues. NBMA networks can support more than two routers but have no broadcast capability. For example, if the NBMA topology is not fully meshed, then a broadcast or multicast that is sent by one router will not reach all the other routers. Frame Relay, ATM, and X.25 are examples of NBMA networks.

To implement broadcasting or multicasting in an NBMA network, the router replicates the packets to be broadcast or multicast and sends them individually on each permanent virtual circuit (PVC) to all destinations. This process is CPU- and bandwidth-intensive.

The default OSPF hello and dead intervals on NBMA interfaces are 30 seconds and 120 seconds, respectively.

DR Election in NBMA Topology

- OSPF considers NBMA to be like other broadcast media.
- The DR and BDR need to have fully meshed connectivity with all other routers, but NBMA networks are not always fully meshed:
 - The DR and BDR each need a list of neighbors.
- OSPF neighbors are not automatically discovered by the router.

© 2009 Cisco Systems, Inc. All rights reserved.

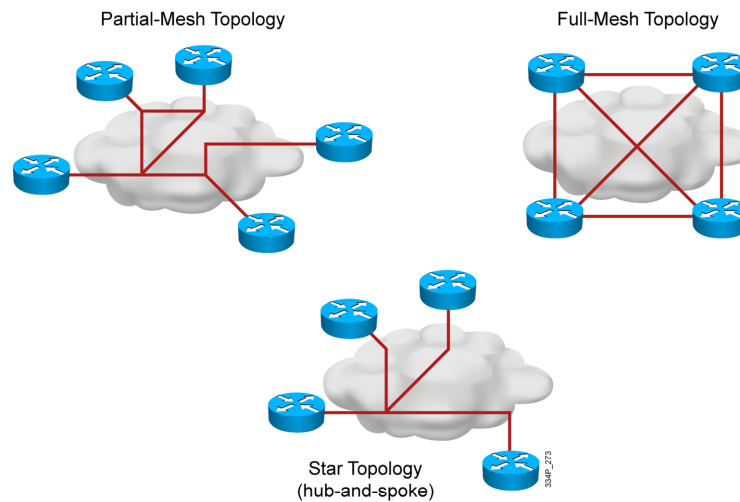
ROUTE v1.0-3-10

OSPF treats NBMA environments like any other broadcast media environment, such as Ethernet; however, NBMA clouds are usually built in hub-and-spoke topologies, using PVCs or switched virtual circuits (SVCs). A hub-and-spoke topology means that the NBMA network is only a partial mesh. In these cases, the physical topology does not provide multiaccess capability, on which OSPF relies.

The election of the DR becomes an issue in NBMA topologies, because the DR and BDR need to have complete physical connectivity with all routers in the NBMA network. The DR and BDR also need to have a list of all the other routers so that they can establish adjacencies.

OSPF cannot automatically build adjacencies with neighboring routers over NBMA interfaces.

Frame Relay Topologies



There are several OSPF configuration choices for a Frame Relay network, depending on the Frame Relay network topology.

With Frame Relay, remote sites interconnect in many ways. By default, interfaces that support Frame Relay are multipoint connection types. The following examples are types of Frame Relay topologies:

- **Star topology:** A star topology, also known as a hub-and-spoke configuration, is the most common Frame Relay network topology. In this topology, remote sites connect to a central site that generally provides a service or application. The star topology is the least expensive topology, because it requires the smallest number of PVCs. The central router provides a multipoint connection, because it typically uses a single interface to interconnect multiple PVCs.
- **Full-mesh topology:** In a full-mesh topology, all routers have virtual circuits to all other destinations. This method, although costly, provides direct connections from each site to all other sites and allows for redundancy. As the number of nodes in the full-mesh topology increases, the topology becomes increasingly expensive.

To figure out how many virtual circuits are needed to implement a fully meshed topology, use the formula $(n * (n - 1)) / 2$, where n is the number of nodes in the network.

- **Partial-mesh topology:** In a partial-mesh topology, not all sites have direct access to a central site. This method reduces the cost of implementing a full-mesh topology.

OSPF over NBMA Topology Modes of Operation

- There are five modes of OSPF operation.
- RFC 2328-compliant modes are as follows:
 - Nonbroadcast (NBMA)
 - Point-to-multipoint
- Additional modes from Cisco are as follows:
 - Point-to-multipoint nonbroadcast
 - Broadcast
 - Point-to-point

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-3-12

As described in RFC 2328, OSPF runs in one of the following two official modes in NBMA topologies:

- **Nonbroadcast:** The nonbroadcast (NBMA) mode simulates the operation of OSPF in broadcast networks. Neighbors must be manually configured, and DR and BDR election is required. This configuration is typically used with fully meshed networks.
- **Point-to-multipoint:** The point-to-multipoint mode treats the nonbroadcast network as a collection of point-to-point links. In this environment, the routers automatically identify their neighboring routers but do not elect a DR and BDR. This configuration is typically used with partially meshed networks.

The choice between NBMA and point-to-multipoint modes determines the way that the Hello protocol and flooding work over the nonbroadcast network. The main advantage of the point-to-multipoint mode is that it requires less manual configuration, and the main advantage of the nonbroadcast mode is that there is less overhead traffic.

Cisco has defined the following additional modes:

- Point-to-multipoint nonbroadcast
- Broadcast
- Point-to-point

The description of all the modes is as follows:

OSPF Network Type over NBMA Topology	Description
broadcast	Cisco extension: <ul style="list-style-type: none"> ■ Makes the WAN interface appear to be a LAN ■ Has one IP subnet ■ Uses multicast OSPF hello packets to automatically discover the neighbors ■ Elects the DR and BDR ■ Requires a full-mesh or a partial-mesh topology
non-broadcast	RFC-compliant mode: <ul style="list-style-type: none"> ■ Has one IP subnet ■ Requires neighbors to be manually configured ■ Elects the DR and BDR ■ Requires that the DR and BDR have full connectivity with all other routers ■ Is typically used in a full-mesh or a partial-mesh topology
point-to-multipoint	RFC-compliant mode: <ul style="list-style-type: none"> ■ Has one IP subnet ■ Uses multicast OSPF hello packets to automatically discover the neighbors ■ Does not require DR and BDR—router sends additional LSAs with more information about neighboring routers ■ Is typically used in partial-mesh or star topology
point-to-multipoint non-broadcast	Cisco extension: <ul style="list-style-type: none"> ■ Is used in place of RFC-compliant point-to-multipoint mode if multicast and broadcast are not enabled on the virtual circuits, because the router cannot dynamically discover its neighboring routers using hello multicast packets ■ Requires neighbors to be manually configured ■ Does not require DR and BDR election
point-to-point	Cisco extension: <ul style="list-style-type: none"> ■ Has a different IP subnet on each subinterface ■ Does not have DR or BDR election ■ Is used when only two routers need to form an adjacency on a pair of interfaces ■ Can be used with either LAN or WAN Interfaces

Broadcast mode is a workaround for statically listing all existing neighboring routers. The interface is set to broadcast and behaves as though the router connects to a LAN. DR and BDR election is still performed; therefore, take special care to ensure either a full-mesh topology or a static election of the DR based on the interface priority.


The other modes are examined in the following topics.

OSPF over Frame Relay NBMA

This topic describes how to configure an NBMA topology in an OSPF over Frame Relay network.

Nonbroadcast Mode (NBMA Mode)

- Treated as a broadcast network by OSPF (like a LAN).
- All serial ports are part of the same IP subnet.
- Frame Relay, X.25, and ATM networks default to nonbroadcast mode.
- Duplicates LSA updates.



© 2009 Cisco Systems, Inc. All rights reserved.

RD UTE v1.0-3-13

In nonbroadcast mode, OSPF emulates operation over a broadcast network. A DR and BDR are elected for the NBMA network, and the DR originates an LSA for the network. In this environment, the routers are usually fully meshed to facilitate the establishment of adjacencies among the routers. If the routers are not fully meshed, you should select the DR and BDR manually to ensure that the selected DR and BDR have complete connectivity to all other neighboring routers. Neighboring routers are statically defined to start the DR and BDR election process. When using nonbroadcast mode, all routers are on one IP subnet.

For flooding over a nonbroadcast interface, the link-state update (LSU) packet must be replicated for each PVC. The updates are sent to each of the neighboring routers that are defined in the neighbor table on the interface.

When there are few neighbors in the network, nonbroadcast mode is the most efficient way to run OSPF over NBMA networks, because it has less overhead than point-to-multipoint mode.


Frame Relay networks default to OSPF nonbroadcast mode.

Steps to Configure NBMA Mode

- Enable the OSPF routing process.
- Define the interfaces that OSPF will run on.

NBMA-specific configuration:

- Statically define a neighbor relationship.
- Define the OSPF network type.



© 2009 Cisco Systems, Inc. All rights reserved.

RO UTE v1.0—3-14

At the beginning, one or more OSPF routing processes must be enabled on the router, followed by configuration that defines which interfaces are involved in OSPF routing. Details will be described in the following lessons.

Once all the required information is defined and basic OSPF configuration is applied, an implementation plan showing the following tasks is required to configure the NBMA-specific configuration:

- Manual configuration of the neighbors
- Definition of the OSPF network type

Nonbroadcast Mode Operation

- Neighbors must be statically configured.
- The OSPF network type must be defined.

```
R1(config-router)#
```

```
neighbor 192.168.1.2 priority 0
```

- Use this command to statically define neighbor relationships in an NBMA network.

```
R1(config-if)#
```

```
ip ospf network non-broadcast
```

- This command defines the OSPF nonbroadcast network type.

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-3-15

To configure OSPF routers that are interconnecting to nonbroadcast networks, use the OSPF **neighbor** command in router configuration mode. Adjacent relationships must be statically defined in NBMA networks using the nonbroadcast mode.

One neighbor entry must be included in the Cisco IOS Software configuration for each known nonbroadcast network neighbor, and the neighbor address must be on the primary address of the interface.

Note You cannot use the OSPF **neighbor** command to specify an OSPF neighbor on nonbroadcast networks within an OSPF VPN routing instance.

The **neighbor 192.168.1.2 priority 0** command in the figure shows the configuration where the **priority 0** optional keyword is used. A number indicates the router priority value of the nonbroadcast neighbor that is associated with the specified IP address. The default value is 0. If this value is used, the specified neighbor cannot become a DR or BDR.

Note that this configuration is performed on one router for its neighbor. In other words, this example specifies on router R1 that the router R2 priority should be 0. Nevertheless, a priority value may already be set on R2. If this is the case, the priority that is defined locally on R2 will override the wished priority that is defined on R1 for R2, and the **priority** command will have no effect.

Additionally, an OSPF network type must be defined. To configure the OSPF network type to a type other than the default for a given medium, use the **ip ospf network** command in interface configuration mode. The **ip ospf network non-broadcast** command in the figure shows that the network type that is selected is NBMA.

It is very important that both adjacent routers agree on the media type. If one router considers the media as broadcast and the other as nonbroadcast, adjacency may not form because hello intervals are not the same on broadcast and nonbroadcast segments. On broadcast segments, hello messages are sent every 10 seconds. On nonbroadcast segments, hello intervals are sent every 30 seconds. The dead interval, which determines the interval during which at least one hello packet must be received from a neighbor before the router declares that neighbor down,


also depends on the hello interval (four times the hello interval by default). Neighboring routers need to have the same interval values to form an adjacency. You can manually adjust the intervals with the following interface commands:

- **ip ospf hello-interval** *seconds* adjusts the OSPF hello interval on this interface to the value expressed by the *seconds* parameter
- **ip ospf dead-interval** *seconds* adjusts the OSPF dead interval on this interface to the value expressed by the *seconds* parameter

For more details about the OSPF **neighbor** and **ip ospf network** commands, go to the Cisco IOS IP Routing: OSPF Command Reference via the following link:

http://www.cisco.com/en/US/docs/ios/iproute_ospf/command/reference/iro_book.html

NBMA Configuration Example



```
R1#
interface Serial0/0/0
 ip address 192.168.1.1 255.255.255.0
 encapsulation frame-relay
 ip ospf network non-broadcast
 <output omitted>

router ospf 100
 network 192.168.0.0 0.0.255.255 area 0
 neighbor 192.168.1.2 priority 0
 neighbor 192.168.1.3 priority 0
 network 172.16.0.0 0.0.255.255 area 0

R2#
interface Serial0/0/0
 ip address 192.168.1.2 255.255.255.0
 encapsulation frame-relay
 ip ospf network non-broadcast
 <output omitted>

router ospf 100
 network 192.168.0.0 0.0.255.255 area 0
```

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-3-16

Example: NBMA Configuration


The figure illustrates an example of statically defined adjacencies and the configuration of a nonbroadcast OSPF network type. All three routers are using the default nonbroadcast mode on their Frame Relay interfaces; therefore, each must manually configure its neighboring routers.

The priority should be set to 0 for routers R2 and R3, because a full-mesh topology does not exist. This configuration ensures that R1 becomes the DR, because only R1 has complete connectivity to the other two routers. No BDR will be elected in this case.

Note The default priority on the neighbor command is 1. Setting the OSPF priority to 0 at the interface level on R2 and R3 resulted in a priority of 0 and the routers not being elected as a DR or BDR.

In an NBMA network, neighbor statements are required only on the DR and BDR. In a hub-and-spoke topology, neighbor statements must be used on the hub, which must be configured to become the DR. Neighbor statements are not mandatory on the spoke routers. In a full-mesh NBMA topology, you may need neighbor statements on all routers unless you have statically configured the DR and BDR using the **priority** command.

The show ip ospf neighbor Command



```
R1# show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
192.168.1.3	0	FULL/DROTHER	00:01:57	192.168.1.3	Serial0/0/0
192.168.1.2	0	FULL/DROTHER	00:01:33	192.168.1.2	Serial0/0/0
172.16.1.1	1	FULL/BDR	00:00:34	172.16.1.1	FastEthernet0/0

To display OSPF neighbor information on a per-interface basis, you must use the **show ip ospf neighbor** command. This table describes the output of the **show ip ospf neighbor** command:

Field	Description
Neighbor ID	Neighbor router ID
Priority	Router priority of the neighbor
State	Neighbor state
Dead Time	Expected time before Cisco IOS Software will declare the neighbor dead
Address	IP address of the interface
Interface	Local interface to reach the neighbor

Example: show ip ospf neighbor Command

In the figure, R1 has a serial Frame Relay NBMA interface and a Fast Ethernet interface.

The serial 0/0/0 interface on this router has two neighbors; both show a state of FULL/DROTHER. DROTHER means that the neighboring router is not a DR or BDR (because R1 is the DR, and there is no BDR in this network).

The neighbor that is learned on Fast Ethernet 0/0 shows a state of FULL/BDR, which means that it has successfully exchanged LSDB information with the router issuing the **show** command, and that it is the BDR.

For more details about the **show ip ospf neighbor** command, go to the Cisco IOS IP Routing: OSPF Command Reference via the following link:

http://www.cisco.com/en/US/docs/ios/iproute_ospf/command/reference/iro_book.html

Using Subinterfaces in OSPF over Frame Relay

This topic describes how to configure a physical interface into multiple subinterfaces.

Using Subinterfaces

- Several logical subinterfaces can be created over all multiaccess WAN networks:
 - Point-to-point
 - Multipoint
- Each subinterface requires an IP subnet.
- Logical interfaces behave in exactly the same way as physical interfaces for routing purposes.
- Statistics and traffic-shaping behavior differ between interfaces and subinterfaces.

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-3-18

A physical interface can be split into multiple logical interfaces, called subinterfaces. Each subinterface is defined as a point-to-point or a point-to-multipoint interface. Subinterfaces were originally created to better handle issues that are caused by split horizon over an NBMA for distance vector-based routing protocols. Each subinterface requires an IP subnet.

Logical interfaces behave in the same way as physical interfaces for routing purposes. Statistics and traffic-shaping behavior differs between interfaces and subinterfaces.

During the configuration of subinterfaces, you must choose the **point-to-point** or **multipoint** keywords. The choice of modes affects the operation of OSPF.

The default OSPF mode on a point-to-point Frame Relay subinterface is the point-to-point mode; the default OSPF mode on a Frame Relay point-to-multipoint subinterface is the nonbroadcast mode. The default OSPF mode on a main Frame Relay interface is also the nonbroadcast mode.

The following topics will describe point-to-point and multipoint subinterfaces and their use in the OSPF NBMA mode.

Point-to-Point Subinterfaces

- Each PVC gets its own subinterface.
- PVCs are treated like point-to-point links.
- Each subinterface requires a subnet.
- OSPF point-to-point mode is the default:
 - DR and BDR are not used.
 - You do not need to configure neighbors.

R1 (config) #

```
interface serial 0/0/0.1 point-to-point
```

- This shows how to configure a point-to-point subinterface.

© 2009 Cisco Systems, Inc. All rights reserved.


RO UTE v1.0—3-19

When point-to-point subinterfaces are configured, each PVC gets its own subinterface, and PVCs are treated like point-to-point links.

A point-to-point subinterface has the properties of any physical point-to-point interface and requires its own subnet. There is no DR or BDR election process. Neighbor discovery is automatic, so neighbors do not need to be configured.

To configure the point-to-point interface, you must select the **point-to-point** keyword during subinterface configuration, as shown on the figure.

Point-to-Point Subinterface Example



R1#

```
interface Serial0/0/0
no ip address
encapsulation frame-relay
!
interface Serial0/0/0.1 point-to-point
ip address 192.168.1.1 255.255.255.0
frame-relay interface-dlci 121
interface Serial0/0/0.2 point-to-point
ip address 192.168.2.1 255.255.255.0
frame-relay interface-dlci 132
```

R2#

```
interface Serial0/0/0
no ip address
encapsulation frame-relay
!
interface Serial0/0/0.1 point-to-point
ip address 192.168.1.2 255.255.255.0
frame-relay interface-dlci 121
```

394P_283

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-1-20

Example: Point-to-Point Subinterface

Point-to-point mode is used when only two nodes exist; this mode is typically used only with point-to-point subinterfaces. Each point-to-point connection is one IP subnet. An adjacency forms over the point-to-point network with no DR or BDR election.

In the figure, the R1 serial 0/0/0 interface is configured with point-to-point subinterfaces. Although all three routers have only one physical serial port, R1 appears to have two logical ports. Each logical port (subinterface) has its own IP address and operates as a point-to-point OSPF network type. This type of configuration avoids the need for a DR or BDR and removes the requirement to define the neighbors statically.

Multipoint Subinterfaces

- Multiple PVCs are on a single subinterface.
- Each subinterface requires a subnet.
- OSPF nonbroadcast mode is the default:
 - The DR is used.
 - Neighbors need to be statically configured.

R1 (config) #

```
interface serial 0/0/0.1 multipoint
```

- This shows how to configure a multipoint subinterface.

© 2009 Cisco Systems, Inc. All rights reserved.


RO UTE v1.0—3-21

When multipoint subinterfaces are configured, there are multiple PVCs on a single subinterface. Each subinterface requires a subnet.

Multipoint Frame Relay subinterfaces default to OSPF nonbroadcast mode, which requires neighbors to be statically configured and performs a DR and BDR election.

To configure the multipoint interface, you must select the **multipoint** keyword during subinterface configuration, as shown in the figure.

Multipoint Subinterface Example



```
R1#
interface Serial0/0/0.1 point-to-point
ip address 192.168.1.1 255.255.255.0
<output omitted>
interface Serial0/0/0.2 multipoint
ip address 192.168.2.1 255.255.255.0
<output omitted>

router ospf 100
network 192.168.0.0 0.0.255.255 area 0
neighbor 192.168.2.2 priority 0
neighbor 192.168.2.3 priority 0

R2#
interface Serial0/0/0
ip address 192.168.1.2 255.255.255.0
<output omitted>

R3#
interface Serial0/0/0
ip address 192.168.2.2 255.255.255.0
ip ospf priority 3
<output omitted>
router ospf 100
network 192.168.0.0 0.0.255.255 area 0
```

© 2009 Cisco Systems, Inc. All rights reserved.

RO UTE v1.0-3-22

Example: Multipoint Subinterface

In this figure, R1 has one single interface, serial 0/0/0, which has been logically separated into two subinterfaces: one point-to-point (S0/0/0.1) and one point-to-multipoint (S0/0/0.2). The multipoint subinterface supports two other routers in a single subnet. Each subinterface requires a subnet. On a multipoint interface, neighbors are statically defined.

OSPF defaults to point-to-point mode on the point-to-point subinterface.


OSPF defaults to nonbroadcast mode on the point-to-multipoint interface.

Implementing OSPF over a Point-to-Point Frame Relay Network

This topic describes how to configure a point-to-point topology in an OSPF over Frame Relay network.

Point-to-Point Mode

- Leased-line emulation.
- Automatic configuration of adjacency.
- DR is not used.
- Only a single subnet is used.




© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-1-23

Networks in point-to-point mode are designed to emulate leased-line networks. This is a Cisco proprietary mode. Because they are treated as point-to-point links, the DR and BDR are not used, and the adjacency forms over the point-to-point network with no DR or BDR election.

Only a single subnet is used per point-to-point link.

Point-to-Point Configuration Example



```
R1#  
interface Serial0/0/0  
ip address 192.168.1.1 255.255.255.0  
encapsulation frame-relay  
ip ospf network point-to-point  
<output omitted>  
  
router ospf 100  
log-adjacency-changes  
network 172.16.0.0 0.0.255.255 area 0  
network 192.168.0.0 0.0.255.255 area 0
```

```
R2#  
interface Serial0/0/0  
ip address 192.168.1.2 255.255.255.0  
encapsulation frame-relay  
ip ospf network point-to-point
```

© 2009 Cisco Systems, Inc. All rights reserved.

RO UTE v1.0-3-24

Example: Point-to-Point Configuration

The figure shows partial configuration of R1 and R2 in point-to-point mode. This configuration does not require subinterfaces and uses only a single subnet.

In point-to-point mode, a DR or BDR is not used; therefore, DR and BDR election and priorities are not a concern.

For more details about the **ip ospf network** command, go to the Cisco IOS IP Routing: OSPF Command Reference via the following link:

http://www.cisco.com/en/US/docs/ios/iproute_ospf/command/reference/iro_book.html

Point-to-Point Verification Example

```
RI#show ip ospf interface s0/0/0
Serial0/0/0 is up, line protocol is up
 Internet Address 192.168.1.1/24, Area 0
 Process ID 100, Router ID 192.168.1.1, Network Type POINT_TO_POINT, Cost: 781
 Transmit Delay is 1 sec, State POINT_TO_POINT
 Timer intervals configured, Hello 30, Dead 120, Wait 120, Retransmit 5
  oob-resync timeout 120
  Hello due in 00:00:26
 Supports Link-local Signaling (LLS)
 Index 2/2, flood queue length 0
 Next 0x0(0)/0x0(0)
 Last flood scan length is 1, maximum is 1
 Last flood scan time is 0 msec, maximum is 4 msec
 Neighbor Count is 1, Adjacent neighbor count is 1
   Adjacent with neighbor 192.168.1.2
 Suppress hello for 0 neighbor(s)
```

The **show ip ospf interface** command displays key OSPF details for each interface.

The OSPF network type, area number, cost, and state of the interface are all displayed. The hello interval for a point-to-multipoint interface is 30 seconds, with a dead interval of 120 seconds.

The point-to-point and broadcast modes default to a 10-second hello timer. The hello and dead timers on the neighboring interfaces must match for the neighbors to form successful adjacencies.

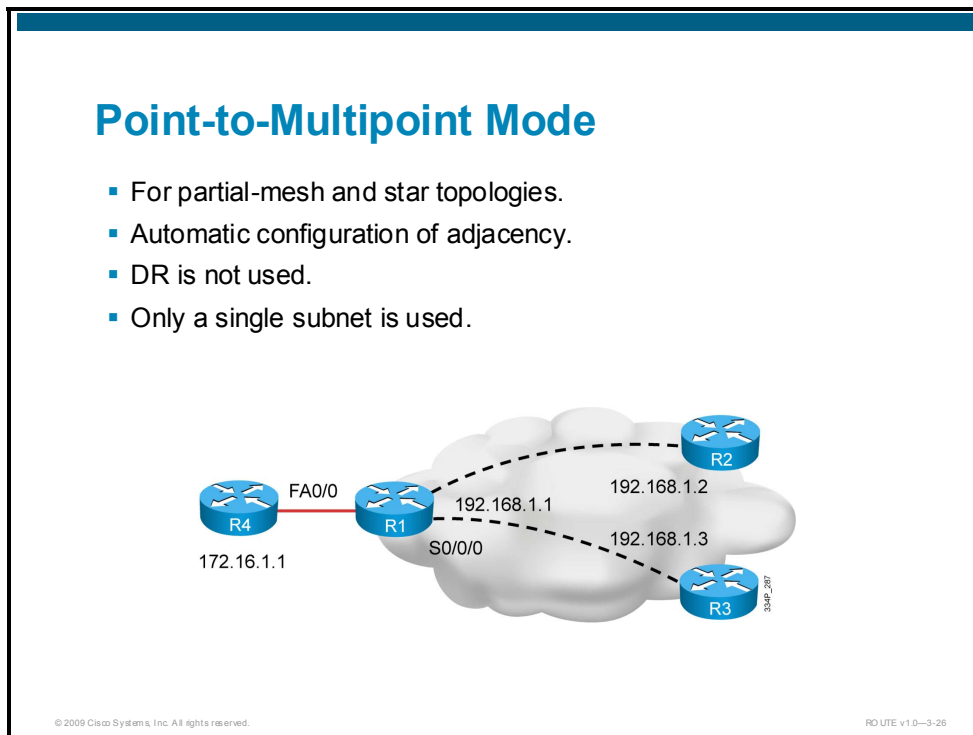
The listed adjacent neighboring routers are all dynamically learned. Manual configuration of neighboring routers is not necessary.

For more details about the **show ip ospf interface** command, go to the Cisco IOS IP Routing: OSPF Command Reference via the following link:

http://www.cisco.com/en/US/docs/ios/iproute_ospf/command/reference/iro_book.html

Implementing OSPF over a Point-to-Multipoint Frame Relay Network

This topic describes how to configure a point-to-multipoint topology and point-to-multipoint nonbroadcast topologies in an OSPF over Frame Relay network.




Networks in point-to-multipoint mode are designed to work with partial-mesh or star topologies. With RFC 2328-compliant point-to-multipoint mode, OSPF treats all router-to-router connections over the nonbroadcast network as if they were point-to-point links. In point-to-multipoint mode, DRs are not used, and a type 2 LSA (as described in the next lesson) is not flooded to adjacent routers. Instead, OSPF point-to-multipoint works by exchanging additional LSUs that are designed to automatically discover neighboring routers and add them to the neighbor table.

In large networks, using point-to-multipoint mode reduces the number of PVCs that are required for complete connectivity, because you are not required to have a full-mesh topology. In addition, not having a full-mesh topology reduces the number of neighbor entries in the neighbor table. Point-to-multipoint mode has the following properties:

- **Does not require a fully meshed network:** This environment allows routing to occur between two routers that are not directly connected but are connected through a router that has virtual circuits to each of the two routers. All three routers that are connected to the Frame Relay network in the figure can be configured for point-to-multipoint mode.
- **Does not require a static neighbor configuration:** In nonbroadcast mode, neighboring routers are statically defined to start the DR election process and allow the exchange of routing updates. Because the point-to-multipoint mode treats the network as a collection of point-to-point links, multicast hello packets discover neighboring routers dynamically. Statically configuring neighboring routers is not necessary.

- **Uses one IP subnet:** As in nonbroadcast mode, when you are using point-to-multipoint mode, all routers are on one IP subnet.
- **Duplicates LSA packets:** Also as in nonbroadcast mode, when flooding out a nonbroadcast interface in point-to-multipoint mode, the router must replicate the LSU. The LSU packet is sent to each of the neighboring routers of the interface, as defined in the neighbor table.

Point-to-Multipoint Configuration Example



R1#

```
interface Serial0/0/0
ip address 192.168.1.1 255.255.255.0
encapsulation frame-relay
ip ospf network point-to-multipoint
<output omitted>

router ospf 100
log-adjacency-changes
network 172.16.0.0 0.0.255.255 area 0
network 192.168.0.0 0.0.255.255 area 0
```

R2#

```
interface Serial0/0/0
ip address 192.168.1.2 255.255.255.0
encapsulation frame-relay
ip ospf network point-to-multipoint
<output omitted>

router ospf 100
network 192.168.0.0 0.0.255.255 area 0
```

300P_260

© 2009 Cisco Systems, Inc. All rights reserved.

RO UTE v1.0-3-27

Example: Point-to-Multipoint Configuration

The figure shows partial configurations of R1 and R2 in point-to-multipoint mode. This configuration does not require subinterfaces and uses only a single subnet. The R3 configuration would be very similar to the configuration that is displayed for R2, using R3-specific IP addressing information.

In point-to-multipoint mode, a DR or BDR is not used; therefore, DR and BDR election and priorities are not a concern.

Point-to-Multipoint Verification Example

```
R1#show ip ospf interface s0/0/0
Serial0/0/0 is up, line protocol is up
Internet Address 192.168.1.1/24, Area 0
Process ID 100, Router ID 192.168.1.1, Network Type POINT_TO_MULTIPPOINT, Cost: 781
Transmit Delay is 1 sec, State POINT_TO_MULTIPPOINT
Timer intervals configured, Hello 30, Dead 120, Wait 120, Retransmit 5
 oob-resync timeout 120
 Hello due in 00:00:26
Supports Link-local Signaling (LLS)
Index 2/2, Flood queue length 0
Next 0x0(0)/0x0(0)
Last flood scan length is 1, maximum is 1
Last flood scan time is 0 msec, maximum is 4 msec
Neighbor Count is 2, Adjacent neighbor count is 2
  Adjacent with neighbor 192.168.1.3
  Adjacent with neighbor 192.168.1.2
Suppress hello for 0 neighbor(s)
```

The **show ip ospf interface** command displays key OSPF details for each interface.

The OSPF network type, area number, cost, and state of the interface are all displayed. The hello interval for a point-to-multipoint interface is 30 seconds, with a dead interval of 120 seconds.

The point-to-multipoint and the nonbroadcast modes default to a 30-second hello timer, while the point-to-point and broadcast modes default to a 10-second hello timer. The hello and dead timers on the neighboring interfaces must match for the neighbors to form successful adjacencies.

The listed adjacent neighboring routers are all dynamically learned. Manual configuration of neighboring routers is not necessary.

For more details about the **show ip ospf interface** command, go to the Cisco IOS IP Routing: OSPF Command Reference via the following link:

http://www.cisco.com/en/US/docs/ios/iproute_ospf/command/reference/iro_book.html

Point-to-Multipoint Nonbroadcast

- Cisco extension to the RFC-compliant point-to-multipoint mode
- Must manually define neighbors—as with NBMA mode
- DR, BDR not used—as with point-to-multipoint mode
- Used in special cases where neighbors cannot be automatically discovered
 - Example: Virtual circuits without multicast and broadcast enabled

```
R1(config-if)#
```

```
ip ospf network point-to-multipoint non-broadcast
```

- Defines the OSPF network type

© 2009 Cisco Systems, Inc. All rights reserved.

RO UTE v1.0-3-29

Cisco defines additional modes for the OSPF neighbor relationship, including point-to-multipoint nonbroadcast. This mode is a Cisco extension of the RFC-compliant point-to-multipoint mode. You must statically define neighbors, and you can modify the cost of the link to the neighboring router to reflect the different bandwidths of each link. The RFC point-to-multipoint mode was developed to support underlying point-to-multipoint virtual circuits that support multicast and broadcast; therefore, this mode allows dynamic neighboring router discovery. For this reason, no DR or BDR is used.

This mode is used in special cases where neighbors cannot be automatically discovered. If, for example, multicast and broadcast are not enabled on the virtual circuits, the RFC-compliant point-to-multipoint mode cannot be used, because the router cannot dynamically discover its neighboring routers using the hello multicast packets; this Cisco mode should be used instead.

To define point-to-multipoint nonbroadcast mode, you must use the **ip ospf network point-to-multipoint non-broadcast** command in an interface configuration mode, as shown in the figure.

OSPF over NBMA Topology Summary

OSPF Mode	NBMA-Preferred Topology	Subnet Address	Hello Timer	Adjacency	RFC or Cisco
Broadcast	Full or partial mesh	Same	10 sec	Automatic, DR and BDR elected	Cisco
Nonbroadcast (NBMA)	Full or partial mesh	Same	30 sec	Manual configuration, DR and BDR elected	RFC
Point-to-multipoint	Partial mesh or star	Same	30 sec	Automatic, no DR or BDR	RFC
Point-to-multipoint nonbroadcast	Partial mesh or star	Same	30 sec	Manual configuration, no DR or BDR	Cisco
Point-to-point	Partial mesh or star, using subinterface	Different for each subinterface	10 sec	Automatic, no DR or BDR	Cisco

© 2009 Cisco Systems, Inc. All rights reserved.

RO UTE v1.0—3-30

Example: OSPF over NBMA Topology Summary

The figure provides a concise comparison of the various modes of operation for OSPF over NBMA topologies.

Summary

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

Summary

- OSPF defines three types of networks: point-to-point, broadcast, and NBMA.
- On point-to-point links, the adjacency is dynamic, uses multicast addresses, and has no DR or BDR.
- On broadcast links, the adjacency is dynamic and includes election of a DR and BDR. All updates are sent to the DR, which forwards the updates to all routers.
- OSPF over Metro Ethernet and EoMPLS requires no changes to the OSPF configuration from the customer perspective.
- OSPF over MPLS VPN requires the customer routers to run OSPF and exchange routing updates with the PE routers.
- The router with the highest OSPF priority is selected as the DR. The router with the second-highest priority value is selected as the BDR.

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-3-31

Summary (Cont.)

- The OSPF mode of operation on Frame Relay depends on the underlying Frame Relay network. OSPF mode options include nonbroadcast, broadcast, point-to-multipoint, point-to-multipoint nonbroadcast, and point-to-point.
- By default on NBMA links, adjacency requires the manual definition of neighbors for the DR and BDR, because OSPF will consider the network similar to broadcast media.
- A physical interface can be split into multiple logical interfaces called subinterfaces. Each subinterface requires an IP subnet.
- With point-to-point mode, leased line is emulated, the adjacency is automatically configured, and no DR is required.
- In point-to-multipoint mode, no DR or BDR is needed, and neighbors are automatically discovered. In point-to-multipoint nonbroadcast mode, no DR or BDR is needed, but neighbors must be statically configured.

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-3-32

Configuring and Verifying OSPF Routing

Overview

This lesson discusses the primary configuration commands for a single-area and multiarea Open Shortest Path First (OSPF) protocol design and how to configure the router OSPF and network commands. The network command for OSPF requires a special inverse mask that is defined in this lesson.

For verification, several OSPF **show** commands are described in this lesson, introducing the OSPF verification process together with a description of the link-state database (LSDB) and details about link-state advertisements (LSAs).

OSPF design limitations and solutions, such as virtual links, the **passive-interface** command, and changing the cost metric, are also described.

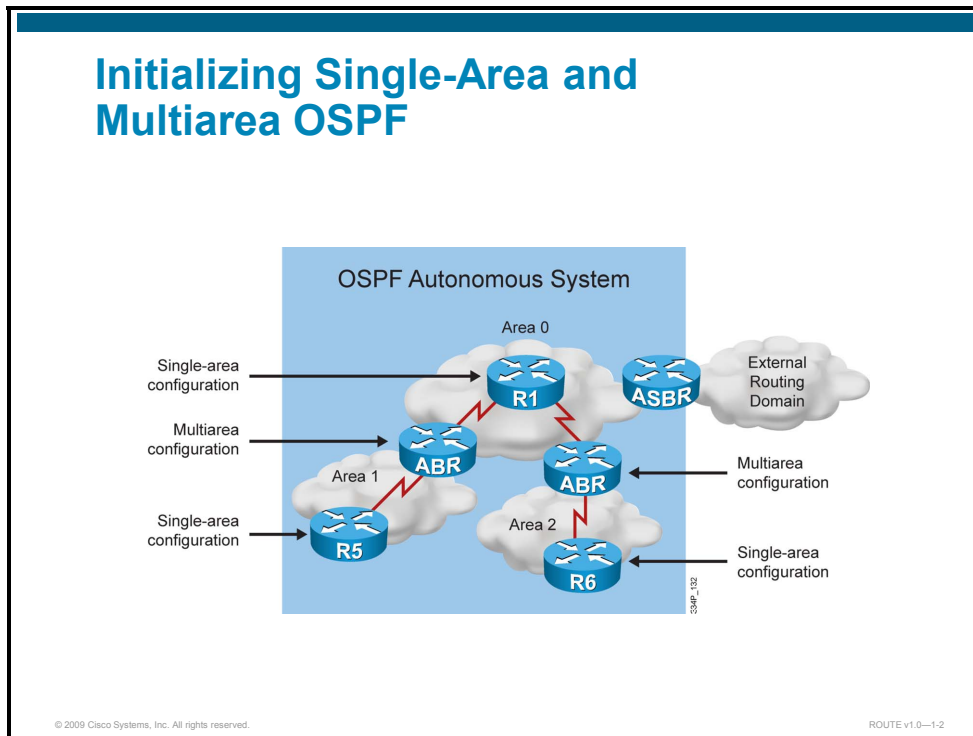
Objectives

Upon completing this lesson, you will be able to explain how to configure OSPF single-area and multiarea routing and how to configure several advanced features. This ability includes being able to meet these objectives:

- Initialize single-area and multiarea OSPF
- Activate OSPF within a router in single-area and multiarea
- Define the router ID
- Verify OSPF operations
- Identify LSA types within the LSDB
- Limit adjacencies in OSPF with the **passive-interface** command
- Define the design limitations of OSPF
- Define OSPF virtual links and solve the discontinuous area problem
- Change the cost metric

Initializing Single-Area and Multiarea OSPF

This topic describes the procedure to configure single-area and multiarea OSPF and the differences between the two.




Single-area design in OSPF puts all routers into a single OSPF area. This design results in many updates being processed on every router and in larger routing tables. The OSPF configuration follows a single-area design, in which all the routers are treated as being internal routers to the area and all the interfaces are members of that single area. The best design should use a single area as the backbone area.

Multiarea design is a better solution than single-area design. In a multiarea design, the network is segmented to limit the propagation of LSAs and to make routing tables smaller. There are two types of routers from the configuration point of view:

- **Routers with single-area configuration:** internal routers, backbone routers, and Autonomous System Boundary Routers (ASBRs) that are residing in one area
- **Routers with a multiarea configuration:** Area Border Routers (ABRs) and ASBRs that are residing in more than one area

Planning for OSPF

- Assess the requirements and options:
 - Contiguous IP addressing plan
 - Network topology with multiple areas
- Define different area types, ABRs, and ASBRs.
- Define summarization and redistribution points.
- Create an implementation plan.
- Configure the OSPF.



The type of OSPF routing protocol implementation that you should choose depends on your specific needs and topology. When you prepare to deploy OSPF routing in a network, you must first gather the existing state and requirements, and then consider different deployment options:

- The IP addressing plan governs how OSPF can be deployed and how well the OSPF deployment might scale. Thus, a detailed IP addressing plan, along with the IP subnetting information, must be created. A solid IP addressing plan should enable usage of OSPF area design and summarization, to more easily scale the network, as well as optimize OSPF behavior and propagation of LSAs.
- The network topology consists of links that connect the network equipment (routers, switches, and so on) and belong to different OSPF areas in a multiarea OSPF design. A detailed network topology plan should be presented to assess the OSPF scalability requirements and determine the different OSPF areas, ABRs, and ASBRs as well as summarization and redistribution points.
- An implementation plan must be created before configuring OSPF routing in the network.

Activating OSPF Within a Router in Single-Area and Multiarea

This topic describes the two-step process to configure basic single-area and multiarea OSPF.

Steps to Configure Basic OSPF

- Configure OSPF routing processes on every OSPF router:
 - Define one or more processes globally on the router.
 - Define the interfaces that OSPF will run on.

Or:

- Enable OSPF explicitly on an interface.

© 2009 Cisco Systems, Inc. All rights reserved. ROUTE v1.0—3-4


Once you have gathered all the required information, you must create an implementation plan showing the following tasks to perform basic OSPF configuration:

- Define one or more OSPF processes globally on the router.
- Define the interfaces that OSPF will run on.

The basic OSPF configuration can be applied directly on the router interface as well. In this case, there is just one step, which enables OSPF on an interface and assigns the interface to the desired OSPF area:

- Enable OSPF explicitly on an interface

Configuring OSPF for Multiple Areas



R1#

```
<output omitted>
interface Fast Ethernet0/0
 ip address 10.64.0.1 255.255.255.0

<output omitted>
router ospf 1
network 10.0.0.0 0.255.255.255 area 0
```

R2#

```
<output omitted>
interface Fast Ethernet0/0
 ip address 10.64.0.2 255.255.255.0

interface Serial 0/0/1
 ip address 10.2.1.2 255.255.255.0
 ip ospf 50 area 1

<output omitted>
router ospf 50
network 10.64.0.2 0.0.0.0 area 0
```

To configure the OSPF process, complete the following steps:

- Enable the OSPF process on the router.
- Identify which interfaces on the router are part of the OSPF process.

To configure an OSPF routing process, use the **router ospf** command in global configuration mode, as shown in the figure. The *process-id* parameter inside the **router ospf** command is an internally used identification parameter for the OSPF routing process. It is locally assigned and can be any positive integer. A unique value is assigned for each OSPF routing process in the router.

To define the interfaces on which OSPF runs and define the area IDs for those interfaces, use the **network area** command in router configuration mode, as shown in the figure. The **wildcard-mask** parameter in the **network area** command determines how to interpret the IP address. The mask has wildcard bits, in which 0 is a match and 1 indicates that the value is not significant. For example, 0.0.255.255 indicates a match in the first two octets.

Starting with Cisco IOS Release 12.3(11)T (and some specific versions of earlier releases), OSPF can be enabled directly on the interface using the **ip ospf area** command, which simplifies the configuration of unnumbered interfaces. Because the command is configured explicitly for the interface, it takes precedence over the **network area** command. The command is shown in the figure as well.

The figure shows the OSPF configuration for Fast Ethernet broadcast networks and serial point-to-point links. Router R1 is in Area 0, router R3 is in Area 1, and router R2 is the ABR between the two areas. R1 is configured for OSPF with a *process-id* value of 1, and a network statement assigns all interfaces that are defined in the 10.0.0.0 network to OSPF process 1. R2 is running OSPF process 50 and has a network statement for Area 0. The configuration for Area 1 in the figure uses the **ip ospf 50 area 1** interface command for serial point-to-point; alternatively, a separate **network** router configuration command could have been used.

Note The network statement and wildcard mask are not used for route summarization purposes. The network statement is used strictly to turn OSPF on for an interface or for multiple interfaces.

For more details about the **router ospf**, **network area** and **ip ospf area** commands, go to the Cisco IOS IP Routing: OSPF Command Reference via the following link:

http://www.cisco.com/en/US/docs/ios/iproute_ospf/command/reference/iro_book.html

Defining the Router ID

For an OSPF routing process to start successfully, it must be able to determine an OSPF router ID. This topic describes how to select the OSPF router ID using the **router-id** command.

OSPF Router ID

- The router is known to OSPF by the router ID number.
- This router ID is used in LSDBs to differentiate one router from the next.
- OSPF requires at least one active interface with an IP address.
- By default, the router ID is:
 - The highest IP address on an active interface at the moment of OSPF process startup.
 - If a loopback interface exists, the highest IP address on any active loopback interface. A loopback interface overrides the OSPF router ID.
- The OSPF **router-id** command can be used to override the default OSPF router ID selection process.
- Using a loopback interface or a **router-id** command is recommended for stability.

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-3-6

The OSPF database uses the OSPF router ID to describe each router in the network uniquely. Remember that every router keeps a complete topology database of all routers and links in an area (or network); therefore, each router should have a unique router ID.

The OSPF routing process chooses a router ID for itself when it starts up. The router ID is a unique IP address that can be assigned in the following ways:

- By default, the highest IP address of any active physical interface when OSPF starts is chosen as the router ID. The interface does not have to be part of the OSPF process, but it has to be up. There must be at least one IP interface that is up on the router for OSPF to use as a router ID. If no interface is up and available with an IP address when the OSPF process starts, the following error message occurs:

```
R1(config)#router ospf 1  
2w1d: %OSPF-4-NORTRID: OSPF process 1 cannot start.
```

- If there is a loopback interface, its address will always be preferred as the router ID over a physical interface address, because a loopback interface never goes down. If there is more than one loopback interface, the highest IP address on any active loopback interface becomes the router ID.
- Using the **router-id** command is the preferred procedure to set the router ID and is always used in preference to the other two procedures.

Once the OSPF router ID is set, it does not change, even if the interface that the router is using for the router ID goes down. The OSPF router ID changes only if the router reloads or if the OSPF routing process restarts.


Configuring Loopback Interfaces

R1 (config) #

```
interface loopback 0
ip address 172.16.17.5 255.255.255.255
```

OSPF is running and the new loopback takes effect in either of these two situations:

- When the router is reloaded
- When the OSPF process is removed and reconfigured



© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—3-7

To modify the OSPF router ID to a loopback address, define a loopback interface as follows:

```
R1 (config) #interface loopback 0
```

Configuring an IP address on a loopback interface overrides the highest IP address that is being used as the router ID. OSPF is more reliable if a loopback interface is configured, because the interface is always active and cannot fail, as opposed to a real interface. For this reason, you should use a loopback address on all key routers. If the loopback address is advertised with the **network** command, this address can be pinged for testing purposes. A private IP address can be used to save registered public IP addresses.

Note A loopback address requires a different subnet for each router, unless the host address itself is advertised. By default, OSPF advertises loopback addresses as /32 host routes.

If the OSPF process is already running, the router must be reloaded or the OSPF process must be removed and reconfigured before the new loopback address will take effect.

Setting OSPF Router ID

R1 (config-router) #

```
router-id 10.10.10.1
```

- This OSPF routing process configuration command changes the OSPF router ID.
- The 32-bit number in the IP address format is used.
- This must be configured before the OSPF process; otherwise, the OSPF process needs to be restarted or the router must be reloaded.

R1#

```
clear ip ospf process
```

- This is the command for a manual OSPF process restart.

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-3-8

Use the OSPF **router-id** command in router configuration mode to ensure that OSPF uses a specific router ID. The **ip-address** parameter can be any unique, arbitrary 32-bit value in an IP address dotted-decimal format.

After the **router-id** command is configured, use the **clear ip ospf process** command. This command restarts the OSPF routing process so that it will reselect the new IP address as its router ID.

Caution The **clear ip ospf process** command will temporarily disrupt an operational network.

Note Router IDs must be unique throughout the autonomous system (AS), no matter how they are configured.

For more details about the **router-id** and **clear ip ospf** commands, go to the Cisco IOS IP Routing: OSPF Command Reference via the following link:

http://www.cisco.com/en/US/docs/ios/iproute_ospf/command/reference/iro_book.html

OSPF Router ID Verification

```
R2#show ip ospf
Routing Process "ospf 50" with ID 10.64.0.2
<output omitted>

Number of areas in this router is 2. 2 normal 0 stub 0 nssa
Number of areas transit capable is 0
External flood list length 0
IETF NSF helper support enabled
Cisco NSF helper support enabled
Area BACKBONE(0)
  Number of interfaces in this area is 1
  Area has no authentication
  SPF algorithm last executed 00:01:17.896 ago
  SPF algorithm executed 4 times
<output omitted>
Area 1
  Number of interfaces in this area is 1
  Area has no authentication
  SPF algorithm last executed 00:00:46.668 ago
  SPF algorithm executed 3 times
<output omitted>
```

©2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-3-9

Use the **show ip ospf** command to verify the OSPF router ID. This command also displays OSPF timer settings and other statistics, including the number of times that the Shortest Path First (SPF) algorithm has been executed. This command also has optional parameters so that you can further specify the information that is to be displayed.

The figure shows partial output from this command on R2 in the last example; R2 is an ABR. The complete output is as follows:

```
R2#show ip ospf
Routing Process "ospf 50" with ID 10.64.0.2
Start time: 00:04:49.064, Time elapsed: 00:03:16.952
Supports only single TOS(TOS0) routes
Supports opaque LSA
Supports Link-local Signaling (LLS)
Supports area transit capability
It is an area border router
Router is not originating router-LSAs with maximum metric
Initial SPF schedule delay 5000 msec
Minimum hold time between two consecutive SPF's 10000 msec
Maximum wait time between two consecutive SPF's 10000 msec
Incremental-SPF disabled
Minimum LSA interval 5 secs
Minimum LSA arrival 1000 msec
LSA group pacing timer 240 secs
Interface flood pacing timer 33 msec
Retransmission pacing timer 66 msec
Number of external LSA 0. Checksum Sum 0x000000
Number of opaque AS LSA 0. Checksum Sum 0x000000
Number of DCbitless external and opaque AS LSA 0
Number of DoNotAge external and opaque AS LSA 0
```

```
Number of areas in this router is 2. 2 normal 0 stub 0 nssa
Number of areas transit capable is 0
External flood list length 0
IETF NSF helper support enabled
Cisco NSF helper support enabled
  Area BACKBONE(0)
    Number of interfaces in this area is 1
    Area has no authentication
    SPF algorithm last executed 00:01:17.896 ago
    SPF algorithm executed 4 times
    Area ranges are
    Number of LSA 4. Checksum Sum 0x031C95
    Number of opaque link LSA 0. Checksum Sum 0x000000
    Number of DCbitless LSA 0
    Number of indication LSA 0
    Number of DoNotAge LSA 0
    Flood list length 0
  Area 1
    Number of interfaces in this area is 1
    Area has no authentication
    SPF algorithm last executed 00:00:46.668 ago
    SPF algorithm executed 3 times
    Area ranges are
    Number of LSA 3. Checksum Sum 0x01198C
    Number of opaque link LSA 0. Checksum Sum 0x000000
    Number of DCbitless LSA 0
    Number of indication LSA 0
    Number of DoNotAge LSA 0
    Flood list length 0
```

For more details about the **show ip ospf** command, go to the Cisco IOS IP Routing: OSPF Command Reference via the following link:

http://www.cisco.com/en/US/docs/ios/iproute_ospf/command/reference/iro_book.html

Verify OSPF Operations

This topic describes how to configure a default route injection into OSPF.

Steps to Verify Basic OSPF

- Verify OSPF routing protocol.
- Verify OSPF interface information.
- Verify OSPF neighbors.
- Verify OSPF routes that are learned by the router in the IP routing table.
- Verify configured IP routing protocol processes.
- Verify OSPF LSDB.

R1#

```
<output omitted>
interface Fast Ethernet0/0
 ip address 10.64.0.1 255.255.255.0


<output omitted>
router ospf 1
network 10.0.0.0 0.255.255.255 area 0
```

R2#

```
<output omitted>
interface Fast Ethernet0/0
 ip address 10.64.0.2 255.255.255.0

interface Serial 0/0/1
 ip address 10.2.1.2 255.255.255.0
 ip ospf 50 area 1

<output omitted>
router ospf 50
network 10.64.0.2 0.0.0.0 area 0
```



© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-3-10

Once the basic OSPF is configured, you can use several verification commands to verify its operation. One of the first steps is to verify the presence of the OSPF routing process. If the OSPF process has been configured successfully, there must be an OSPF router inside the IP routing table. Additionally, you can verify the OSPF interfaces, routing process neighbors, and database.

Example: The show ip ospf Command

- This command displays the OSPF router ID, timers, and statistics.

```
RI#show ip ospf
Routing Process "ospf 1" with ID 10.64.0.1
Start time: 00:01:16.084, Time elapsed: 00:14:58.368
<output omitted>
Minimum hold time between two consecutive SPFs 10000 msec
Maximum wait time between two consecutive SPFs 10000 msec
Incremental-SPF disabled
Minimum LSA interval 5 secs
Minimum LSA arrival 1000 msec
LSA group pacing timer 240 secs
Interface flood pacing timer 33 msec
Retransmission pacing timer 66 msec
<output omitted>
Number of areas in this router is 1. 1 normal 0 stub 0 nssa
Number of areas transit capable is 0
<output omitted>
Area BACKBONE(0)
Number of interfaces in this area is 1
Area has no authentication
SPF algorithm last executed 00:07:26.520 ago
SPF algorithm executed 3 times
<output omitted>
```

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—3-11

Use the **show ip ospf** command to display general information about OSPF routing processes. As shown in the last topic, the **show ip ospf** command displays the OSPF router ID, OSPF timers, the number of times that the SPF algorithm has been executed, and LSA information.

For more details about the **show ip ospf** command, go to the Cisco IOS IP Routing: OSPF Command Reference via the following link:

http://www.cisco.com/en/US/docs/ios/iproute_ospf/command/reference/iro_book.html

Example: The show ip ospf interface Command

- This command displays the OSPF router ID, area ID, and adjacency information.

```
R2#show ip ospf interface FastEthernet 0/0
FastEthernet0/0 is up, line protocol is up
  Internet Address 10.64.0.2/24, Area 0
  Process ID 50, Router ID 10.64.0.2, Network Type BROADCAST, Cost: 1
  Transmit Delay is 1 sec, State BDR, Priority 1
  Designated Router (ID) 10.64.0.1, Interface address 10.64.0.1
  Backup Designated router (ID) 10.64.0.2, Interface address 10.64.0.2
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    oob-resync timeout 40
    Hello due in 00:00:05
  Supports Link-local Signaling (LLS)
  Cisco NSF helper support enabled
  IETF NSF helper support enabled
  Index 1/2, flood queue length 0
  Next 0x0(0)/0x0(0)
  Last flood scan length is 1, maximum is 1
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 10.64.0.1 (Designated Router)
  Suppress hello for 0 neighbor(s)
```

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—3-12

The **show ip ospf interface** command displays OSPF-related interface information. Use this command to verify that interfaces are configured in the intended areas. In addition, this command displays the timer intervals (including the hello interval) and shows the neighbor adjacencies.

The command output in the figure is from R2 from the last configuration example. The output details the OSPF status of the FastEthernet 0/0 interface. Use the output to verify that OSPF is running on this interface and which OSPF area that the interface is in.

This command also displays other information, such as the OSPF process ID, the OSPF router ID, the OSPF network type, the designated router (DR) and backup designated router (BDR), timers, and neighbor adjacency.

For more details about the **show ip ospf interface** command, go to the Cisco IOS IP Routing: OSPF Command Reference via the following link:

http://www.cisco.com/en/US/docs/ios/iproute_ospf/command/reference/iro_book.html

Example: The show ip ospf neighbor Command

- This command displays information about the OSPF neighbors, including the DR and BDR information.

```
R2#show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
10.64.0.1	1	FULL/DR	00:00:32	10.64.0.1	FastEthernet0/0
10.2.1.1	0	FULL/ -	00:00:37	10.2.1.1	Serial0/0/1

```
R2#show ip ospf neighbor detail
```

```
Neighbor 10.64.0.1, interface address 10.64.0.1
  In the area 0 via interface FastEthernet1/0
  Neighbor priority is 1, State is FULL, 6 state changes
  DR is 10.64.0.1 BDR is 10.64.0.2
<output omitted>
Neighbor 10.2.1.1, interface address 10.2.1.1
  In the area 1 via interface Serial2/0
  Neighbor priority is 0, State is FULL, 6 state changes
  DR is 0.0.0.0 BDR is 0.0.0.0
<output omitted>
```

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—3-13

One of the most important OSPF verification and troubleshooting commands is the **show ip ospf neighbor** command. OSPF does not send or receive updates without having full adjacencies between neighbors. The **show ip ospf neighbor** command displays a list of neighbors, including their OSPF router IDs, their OSPF priorities, their neighbor adjacency states (for example, INIT, exstart, or full), and their dead timers.

In the first output in the figure, R2 (from the last configuration example) has two neighbors. The first entry in the table represents the adjacency that is formed on the Fast Ethernet interface. A FULL state means that the LSDB has been exchanged successfully. The DR entry means that a router is the designated router.

The second line in the table represents the neighbor of R2 on the serial interface. DR and BDR are not used on point-to-point interfaces (indicated by a dash [-]).

The second output in the figure shows partial output of the **show ip ospf neighbor detail** command output, providing details of the neighbors of R2. The complete output is as follows:

```
R2#show ip ospf neighbor detail
Neighbor 10.64.0.1, interface address 10.64.0.1
  In the area 0 via interface FastEthernet1/0
  Neighbor priority is 1, State is FULL, 6 state changes
  DR is 10.64.0.1 BDR is 10.64.0.2
  Options is 0x52
  LLS Options is 0x1 (LR)
  Dead timer due in 00:00:30
  Neighbor is up for 00:18:09
  Index 1/1, retransmission queue length 0, number of
  retransmission 0
  First 0x0(0)/0x0(0) Next 0x0(0)/0x0(0)
  Last retransmission scan length is 0, maximum is 0
  Last retransmission scan time is 0 msec, maximum is 0 msec
Neighbor 10.2.1.1, interface address 10.2.1.1
```

```
In the area 1 via interface Serial2/0
Neighbor priority is 0, State is FULL, 6 state changes
DR is 0.0.0.0 BDR is 0.0.0.0
Options is 0x52
LLS Options is 0x1 (LR)
Dead timer due in 00:00:35
Neighbor is up for 00:16:14
Index 1/2, retransmission queue length 0, number of
retransmission 0
First 0x0(0)/0x0(0) Next 0x0(0)/0x0(0)
Last retransmission scan length is 0, maximum is 0
Last retransmission scan time is 0 msec, maximum is 0 msec
```

For more details about the **show ip ospf neighbor** and **show ip ospf neighbor detail** commands, go to the Cisco IOS IP Routing: OSPF Command Reference via the following link: http://www.cisco.com/en/US/docs/ios/iproute_ospf/command/reference/iro_book.html

Example: The show ip route ospf Command

- This command displays all OSPF routes that are learned by the router.

```
R1#show ip route ospf
10.0.0.0/24 is subnetted, 2 subnets
O IA 10.2.1.0 [110/65] via 10.64.0.2, 00:04:29, FastEthernet0/0
```

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—3-14

Use the **show ip route ospf** command to verify the OSPF routes in the IP routing table that are known to the router. This command is one of the best ways to determine connectivity between the local router and the rest of the internetwork. This command also has optional parameters so that you can further specify the information that is to be displayed, including the OSPF process ID.

The “O” code represents OSPF routes, and “IA” means “interarea.” In the figure, the 10.2.1.0 subnet is recognized on FastEthernet 0/0 via neighbor 10.64.0.2. The entry “[110/65]” in the routing table represents the administrative distance that is assigned to OSPF (110) and the total cost of the route to subnet 10.2.1.0 (cost of 65).

For more details about the **show ip route ospf** command, go to the Cisco IOS IP Routing: OSPF Command Reference via the following link:

http://www.cisco.com/en/US/docs/ios/iproute_ospf/command/reference/iro_book.html

Example: The show ip protocols Command

- This command verifies the configured IP routing protocol processes, parameters, and statistics.

```
R1#show ip protocols
Routing Protocol is "ospf 1"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Router ID 10.64.0.1
  Number of areas in this router is 1. 1 normal 0 stub 0 nssa
  Maximum path: 4
  Routing for Networks:
    10.0.0.0 0.255.255.255 area 0
  Reference bandwidth unit is 100 mbps
<output omitted>
```

©2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—3-15

Use the **show ip protocols** command to verify the OSPF routing protocol parameters about timers, filters, metrics, networks, and other information for the entire router.

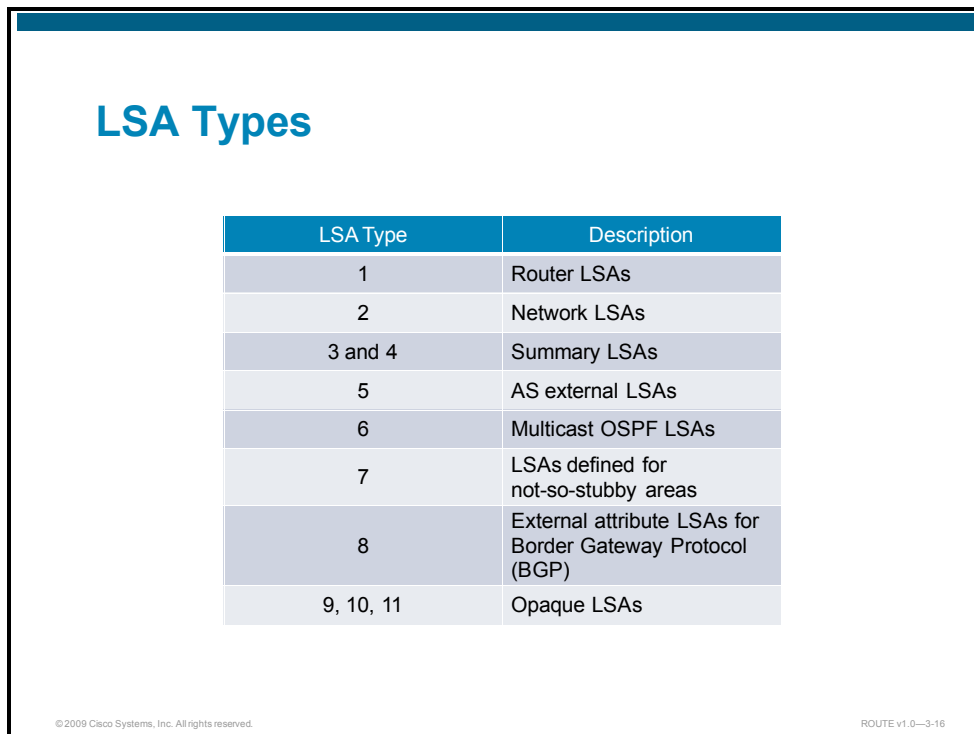
The command output in the figure shows that the OSPF routing protocol with process number 1 is configured. The router ID of the router is 10.64.0.1, and it belongs to Area 0.

For more details about the **show ip protocols** command, go to the Cisco IOS IP Routing: OSPF Command Reference via the following link:

http://www.cisco.com/en/US/docs/ios/iproute_ospf/command/reference/iro_book.html

Identifying LSA Types Within the LSDB

LSAs are the building blocks of the OSPF LSDB. Individually, they act as database records; in combination, they describe the entire topology of an OSPF network or area. This topic describes the LSAs that are defined by OSPF.



The image shows a slide titled "LSA Types" with a table listing various OSPF LSA types and their descriptions. The table has two columns: "LSA Type" and "Description".

LSA Type	Description
1	Router LSAs
2	Network LSAs
3 and 4	Summary LSAs
5	AS external LSAs
6	Multicast OSPF LSAs
7	LSAs defined for not-so-stubby areas
8	External attribute LSAs for Border Gateway Protocol (BGP)
9, 10, 11	Opaque LSAs

© 2009 Cisco Systems, Inc. All rights reserved. ROUTE v1.0-3-16

The following are descriptions of each type of LSA. LSA types 1 through 5 and 7 are explained in more detail in the following pages.

Type 1

Every router generates router link advertisements for each area to which it belongs. Router link advertisements describe the state of the router links to the area and are flooded only within that particular area. For all types of LSAs, there are 20-byte LSA headers. One of the fields of the LSA header is the link-state ID. The link-state ID of the type 1 LSA is the originating router ID.

Type 2

DRs generate network link advertisements for multiaccess networks that describe the set of routers that are attached to a particular multiaccess network. Network link advertisements are flooded in the area that contains the network. The link-state ID of the type 2 LSA is the IP interface address of the DR.

Types 3 and 4

ABRs generate summary link advertisements. Summary link advertisements describe the following interarea routes:

- Type 3 describes routes to networks and aggregates routes.
- Type 4 describes routes to ASBRs.

The link-state ID is the destination network number for type 3 LSAs and the router ID of the described ASBR for type 4 LSAs.

These LSAs are flooded throughout the backbone area to the other ABRs. The link entries are not flooded into totally stubby areas or not-so-stubby areas (NSSAs).

Type 5

ASBRs generate AS external link advertisements. External link advertisements describe routes to destinations that are external to the AS and are flooded everywhere with the exception of stub areas, totally stubby areas, and NSSAs. The link-state ID of the type 5 LSA is the external network number.

Type 6

Type 6 LSAs are specialized LSAs that are used in multicast OSPF applications.

Type 7

Type 7 LSAs are used in NSSAs for external routes.

Type 8


Type 8 LSAs are specialized LSAs that are used in internetworking OSPF and Border Gateway Protocol (BGP).

Types 9, 10, and 11

The opaque LSAs, types 9, 10, and 11, are designated for future upgrades to OSPF for application-specific purposes. For example, Cisco uses opaque LSAs for Multiprotocol Label Switching (MPLS) with OSPF. Standard LSDB flooding mechanisms are used for distribution of opaque LSAs. Each of the three types has a different flooding scope.

Type 1: Router LSA

- One router LSA for every router in an area (intra-area):
 - Includes a list of directly attached links
 - Links identified by the IP prefix and link type
- LSA identified by the router ID of the originating router
- Floods within its area only; does not cross an ABR



A router advertises a type 1 LSA that floods to all other routers in the area in which it originated. A type 1 LSA describes the collective states of the directly connected links (interfaces) of the router.

Each type 1 LSA is identified by the router ID.

Each router link is defined as one of four link types: type 1, 2, 3, or 4. The LSA includes a link ID field that identifies, by the network number and mask, the object to which this link connects.

Type 1 LSA Link Types

Depending on the type, the link ID has different meanings, as described in the table.


Link Type	Description	Link ID
1	Point-to-point connection to another router	Neighboring router ID
2	Connection to a transit network	IP address of DR
3	Connection to a stub network	IP network or subnet number
4	Virtual link	Neighboring router ID

A stub network is a dead-end link that has only one attached router.

In addition, the type 1 LSA describes whether the router is an ABR or ASBR.

Type 2: Network LSA

- One network LSA for each transit broadcast or NBMA network in an area (intra-area):
 - Includes a list of attached routers on the transit link
 - Includes a subnet mask of the link
- Advertised by the DR of the broadcast network
- Floods within its area only; does not cross an ABR



© 2009 Cisco Systems, Inc. All rights reserved.


ROUTE v1.0—3-18

A type 2 LSA is generated for every transit broadcast or nonbroadcast multiaccess (NBMA) network within an area. A transit network has at least two directly attached OSPF routers. A multiaccess network like Ethernet is an example of a transit network.

The DR of the network is responsible for advertising the network LSA. A type 2 network LSA lists each of the attached routers that make up the transit network, including the DR itself and the subnet mask that is used on the link. The type 2 LSA then floods to all routers within the transit network area. Type 2 LSAs never cross an area boundary. The link-state ID for a network LSA is the IP interface address of the DR that advertises it.

Type 3: Summary LSA

- Used to flood network information to areas outside the originating area (interarea)
 - Describes the network number and mask of the link
- Advertised by the ABR of the originating area; regenerated by all subsequent ABRs to flood throughout the AS
- Advertised for every subnet and not summarized, by default



© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—3-19

The ABR sends type 3 summary LSAs. Type 3 LSAs advertise any networks that are owned by an area to the rest of the areas in the OSPF AS, as shown in the figure.

The link-state ID is set to the network number; the mask is also advertised.

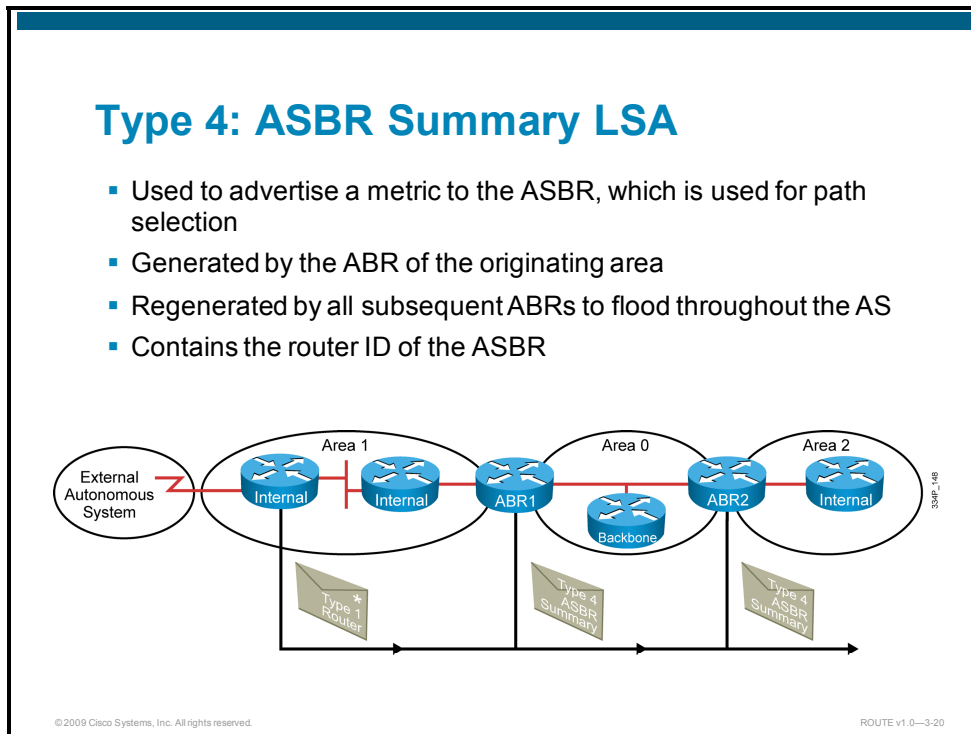
By default, OSPF does not automatically summarize groups of contiguous subnets, nor does it summarize a network to its classful boundary. The network operator, through configuration commands, must specify how the summarization will occur. By default, a type 3 LSA is advertised into the backbone area for every subnet that is defined in the originating area, which can cause significant flooding problems. Consequently, you should always consider using manual route summarization at the ABR.

Summary LSAs are flooded throughout a single area only, but are regenerated by ABRs to flood into other areas.

Note By default, summary LSAs do not contain summarized routes.

Type 4: ASBR Summary LSA

- Used to advertise a metric to the ASBR, which is used for path selection
- Generated by the ABR of the originating area
- Regenerated by all subsequent ABRs to flood throughout the AS
- Contains the router ID of the ASBR




A type 4 summary LSA is generated by an ABR only when an ASBR exists within an area. A type 4 LSA identifies the ASBR and provides a route to it. The link-state ID is set to the ASBR router ID. All traffic that is destined to an external AS requires routing table knowledge of the ASBR that originated the external routes.

Example: LSA Type 4—ASBR Summary LSA

In the figure, the ASBR sends a type 1 router LSA with a bit (known as the external bit [e bit]) that is set to identify itself as an ASBR. When the ABR (identified with the border bit [b bit] in the router LSA) receives this type 1 LSA, it builds a type 4 LSA and floods it to the backbone, Area 0. Subsequent ABRs regenerate a type 4 LSA to flood into their areas.

Type 5: External LSA

- Used to advertise networks from other autonomous systems
- Advertised and owned by the originating ASBR
- Flooded throughout the entire AS
- The advertising router ID (ASBR) is unchanged throughout the AS
- Type 4 LSA needed to find the ASBR
- By default, routes not summarized



© 2009 Cisco Systems, Inc. All rights reserved.


ROUTE v1.0—3-21

Type 5 external LSAs describe routes to networks outside the OSPF AS. Type 5 LSAs are originated by the ASBR and are flooded to the entire AS.

The link-state ID is the external network number. Because of the flooding scope and depending on the number of external networks, the default lack of route summarization can also be a major issue with external LSAs. Therefore, you should always attempt to summarize blocks of external network numbers at the ASBR to reduce flooding problems.

Type 7: NSSA External LSA

- Used to advertise networks from other autonomous systems injected into the NSSA
- Has the same format as a type 5 external LSA
- Advertised and owned by the originating ASBR
- Translated to type 5 LSA on first NSSA subsequent ABR
- By default, routes not summarized



Type 7 external LSAs describe routes to networks outside the OSPF AS. Redistribution from an external AS into an NSSA area creates this special type 7 LSA, which can only exist in an NSSA area. An NSSA ASBR generates this LSA, and an NSSA ABR translates it into a type 5 LSA, which gets propagated into the OSPF domain to all areas that can support type 5 LSAs.

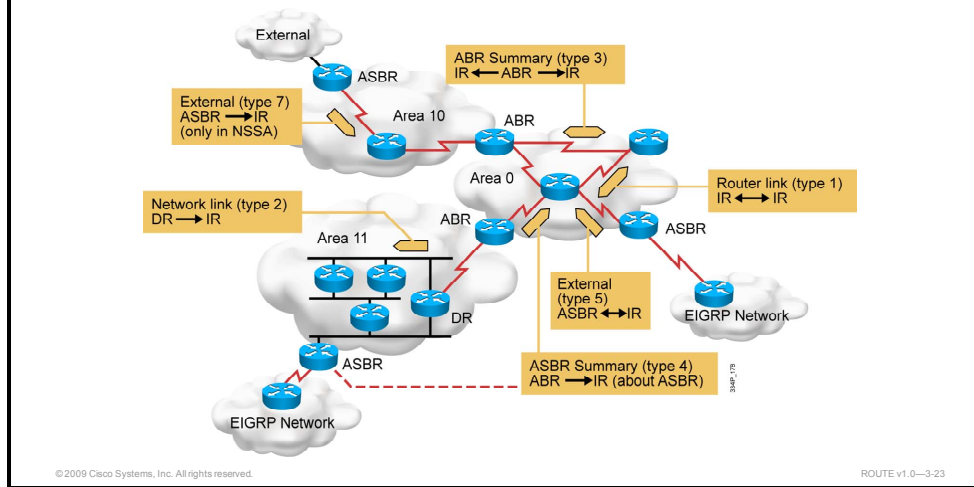
Routers that are operating NSSA areas set the n-bit to signify that they can support type 7 NSSAs. These option bits must be checked during neighbor establishment. They must match for an adjacency to form.

The link-state ID is the external network number. Because of the flooding scope and depending on the number of external networks, the default lack of route summarization can also be a major issue with external LSAs. Therefore, you should always attempt to summarize blocks of external network numbers at the ASBR to reduce flooding problems.

The advertising router is set to the router ID of the router that injected the external route into OSPF—a router inside this NSSA area. This address is also set as the “Forward Address” for this prefix—the address that is used to determine the path to take toward this external destination.

Example of Different LSAs

Note: Only one example of each LSA type exchange is demonstrated in this figure.



Example: Types of LSAs


LSAs are packets that OSPF uses to advertise changes in the condition of a specific link to other OSPF routers. There are various types of link-state packets that are used by OSPF, each of which is generated for a different purpose and flooded in the network.

The following are the different types of LSA packets that can be generated by the source router and entered into the destination router LSA database.

- **Type 1—Router link LSAs:** Router LSAs are generated by each router for each area that it is in. The link-state ID is the originating router ID.
- **Type 2—Intra-area network link LSAs:** Network LSAs are generated by DRs. The link-state ID is the IP interface address of the DR.
- **Type 3—Network summary LSAs for ABRs:** Summary LSAs are generated by ABRs. The link-state ID is the destination network number.
- **Type 4—ASBR summary LSAs for ASBRs:** Summary LSAs are generated by ABRs. The link-state ID is the router ID of the described ASBR.
- **Type 5—AS External LSAs:** Type 5 LSAs are generated by the ASBRs. The link-state ID is the external network number.
- **Type 7: NSSA External LSAs:** Type 7 LSAs are generated by ASBRs.

Note OSPF has more LSAs. Types 9, 10, and 11 (opaque LSAs) may be used for distributing application-specific information through an OSPF domain. Type 9 LSAs are not flooded beyond the local network or subnetwork. Type 10 LSAs are not flooded beyond the borders of their associated area. Type 11 LSAs are flooded throughout the AS. The flooding scope of type 11 LSAs is equivalent to that of AS-external (type 5) LSAs. Cisco MPLS Traffic Engineering (Cisco MPLS TE) functionality has been implemented with type 10 opaque LSAs. For more information about opaque LSAs, please see RFC 2370. Type 8 LSAs are specialized LSAs that are used in internetworking OSPF and BGP.

OSPF LSDB: Intra-Area Routing



R1#

```
interface Fast Ethernet0/0
 ip address 10.64.0.1 255.255.255.0
<output omitted>
router ospf 1
network 10.0.0.0 0.255.255.255 area 0
```

R2#

```
interface Fast Ethernet0/0
 ip address 10.64.0.2 255.255.255.0

interface Serial 0/0/1
 ip address 10.2.1.2 255.255.255.0
 ip ospf 50 area 1

router ospf 50
network 10.64.0.2 0.0.0.0 area 0
```

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—3-24

The figure shows the topology that will be used to describe the OSPF LSDB for intra-area routing. All routers are configured for OSPF routing; the **show ip ospf database** command is used to get information about an OSPF LSDB. The router link states that are presented in the next figure are type 1 and type 2 LSAs.

OSPF LSDB: Intra-Area Routing (Cont.)

```
R2#show ip ospf database
  OSPF Router with ID (10.64.0.2) (Process ID 50)
```

Router Link States (Area 0)					
Link ID	ADV Router	Age	Seq#	Checksum	Link count
10.64.0.1	10.64.0.1	305	0x80000002	0x00C93B	1
10.64.0.2	10.64.0.2	237	0x80000004	0x00C638	1

Net Link States (Area 0)				
Link ID	ADV Router	Age	Seq#	Checksum
10.64.0.1	10.64.0.1	305	0x80000001	0x008D7E

The figure illustrates the use of the **show ip ospf database** command to get information about an OSPF LSDB. The router link states are type 1 LSAs, and the net link states are type 2 LSAs.

The database columns are as follows:

- **Link ID:** This column identifies each LSA.
- **ADV router:** This column shows the address of the advertising router—the source router of the LSA.
- **Age:** This column shows the maximum age counter in seconds; the maximum configurable age counter is 1 hour, or 3600 seconds.
- **Seq#:** This column shows the sequence number of the LSA; the number begins at 0x80000001 and increases with each update of the LSA.
- **Checksum:** This column shows the checksum of the individual LSA, which can be used to ensure reliable receipt of that LSA.
- **Link count:** This column shows the total number of directly attached links, which is used only on router LSAs. The link count includes all point-to-point, transit, and stub links. Each point-to-point serial link counts as two; all other links count as one, including Ethernet links.

The output in the figure is partial output from an ABR—router R2. The complete command output from this router is as follows:

```
R2#sh ip ospf database

  OSPF Router with ID (10.64.0.2) (Process ID 50)

  Router Link States (Area 0)

  Link ID          ADV Router      Age           Seq#
  Checksum Link count

  10.64.0.1        10.64.0.1      305          0x80000002
  0x00C93B 1
  10.64.0.2        10.64.0.2      237          0x80000004
  0x00C638 1
```

10.64.0.1 0x00C73C 1	10.64.0.1	91	0x80000003
10.64.0.2 0x00C439 1	10.64.0.2	46	0x80000005

Net Link States (Area 0)

Link ID Checksum	ADV Router	Age	Seq#
10.64.0.1 0x008B7F	10.64.0.1	91	0x80000002

Summary Net Link States (Area 0)

Link ID Checksum	ADV Router	Age	Seq#
10.2.1.0 0x00FBA6	10.64.0.2	46	0x80000002

Router Link States (Area 1)

Link ID Checksum	Link count	ADV Router	Age	Seq#
10.2.1.1 0x001709 2		10.2.1.1	1969	0x80000002
10.64.0.2 0x005687 2		10.64.0.2	46	0x80000004

Summary Net Link States (Area 1)


Link ID Checksum	ADV Router	Age	Seq#
10.64.0.0 0x00A301	10.64.0.2	47	0x80000002

R2#

For more details about the **show ip ospf database** command, go to the Cisco IOS IP Routing: OSPF Command Reference via the following link:

http://www.cisco.com/en/US/docs/ios/iproute_ospf/command/reference/iro_book.html

OSPF LSDB: Interarea Routing



R1#

```
interface Loopback0
ip address 1.1.1.1 255.0.0.0
interface Ethernet2/0/0
ip address 4.0.0.1 255.0.0.0
interface Serial2/1/0
ip unnumbered Ethernet2/0/0
router ospf 1
network 4.0.0.0 0.255.255.255 area 0
```

R3#

```
interface Loopback0
ip address 3.3.3.3 255.0.0.0
interface ATM2/0.20 point-to-point
ip address 200.0.0.3 255.255.255.0
router ospf 2
network 200.0.0.0 0.255.255.255 area 1
```

R2#

```
interface Loopback0
ip address 2.2.2.2 255.0.0.0
interface Ethernet0/0/4
ip address 6.0.0.2 255.0.0.0
interface Serial0/1/0
ip unnumbered Ethernet0/0/4
interface ATM1/0.20 point-to-point
ip address 200.0.0.2 255.255.255.0
router ospf 2
network 6.0.0.0 0.255.255.255 area 0
network 200.0.0.0 0.255.255.255 area 1
```

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—3-26

The figure presents the topology that will be used to describe the OSPF LSDB for interarea routing. All routers are configured for OSPF routing; the **show ip ospf database** command is used to get information about an OSPF LSDB. The router link states that are presented in the next figure are type 1 and type 3 LSAs.

OSPF LSDB: Interarea Routing (Cont.)

```
R2#show ip ospf database
      OSPF Router with ID (2.2.2.2) (Process ID 2)
```

Router Link States (Area 0)					
Link ID	ADV Router	Age	Seq#	Checksum	Link count
1.1.1.1	1.1.1.1	697	0x80000040	0x5A21	2
2.2.2.2	2.2.2.2	696	0x80000045	0xEE82	2

Summary Net Link States (Area 0)					
Link ID	ADV Router	Age	Seq#	Checksum	
200.0.0.0	2.2.2.2	352	0x80000001	0x2546	

Router Link States (Area 1)					
Link ID	ADV Router	Age	Seq#	Checksum	Link count
2.2.2.2	2.2.2.2	351	0x8000000B	0xCA9D	2
3.3.3.3	3.3.3.3	354	0x80000006	0x71F7	2

Summary Net Link States (Area 1)					
Link ID	ADV Router	Age	Seq#	Checksum	
4.0.0.0	2.2.2.2	689	0x80000001	0xFFE6	
6.0.0.0	2.2.2.2	700	0x80000001	0x63C1	


© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-3-27

The figure illustrates the use of the **show ip ospf database** command to get information about an OSPF LSDB. The router link states are type 1 LSAs, and the summary net link states are type 3 LSAs. Because R2 is the ABR, it has the database for both areas to which it is connected, making it the best place to see the OSPF database. The output in the figure is the complete command output from this router.

To advertise routes from one area into another, the ABR creates summary links, which you can see using the **show ip ospf database summary** command.

OSPF LSDB: External Routes



R1#

```
interface Loopback0
ip address 1.1.1.1 255.0.0.0
interface Serial2/1/0
ip address 5.0.0.1 255.0.0.0
interface Ethernet2/0/0
ip address 4.0.0.1 255.0.0.0
router ospf 4
redistribute static metric 5 metric-type 1
network 5.0.0.0 0.255.255.255 area 1
network 4.0.0.0 0.255.255.255 area 1
ip route 9.0.0.0 255.0.0.0 4.0.0.2
```

R2#

```
interface Loopback0
ip address 2.2.2.2 255.0.0.0
interface Serial0/1/0
ip address 5.0.0.2 255.0.0.0
interface ATM1/0.20
ip address 6.0.0.2 255.0.0.0
router ospf 2
network 5.0.0.0 0.255.255.255 area 1
network 6.0.0.0 0.255.255.255 area 0
```

R3#

```
interface Loopback0
ip address 3.3.3.3 255.0.0.0
interface ATM2/0.20 point-to-point
ip address 6.0.0.3 255.0.0.0
router ospf 2
network 6.0.0.0 0.255.255.255 area 0
```

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—328

The figure shows the topology that will be used to describe the OSPF LSDB for external routes. All routers are configured for OSPF routing; the **show ip ospf database** command is used to get information about an OSPF LSDB. The router link states that are presented in the next figure are type 1, type 3, type 4, and type 5 LSAs.

OSPF LSDB: External Routes (Cont.)

```
R2#show ip ospf database
      OSPF Router with ID (2.2.2.2) (Process ID 2)
```

Router Link States (Area 0)					
Link ID	ADV Router	Age	Seq#	Checksum	Link count
2.2.2.2	2.2.2.2	93	0x80000020	0xCD0B	
3.3.3.3	3.3.3.3	1225	0x8000000D	0x9057	

Summary Net Link States (Area 0)				
Link ID	ADV Router	Age	Seq#	Checksum
4.0.0.0	2.2.2.2	73	0x80000001	0xFFE6
5.0.0.0	2.2.2.2	1651	0x80000006	0x8466

Summary ASB Link States (Area 0)				
Link ID	ADV Router	Age	Seq#	Checksum
1.1.1.1	2.2.2.2	74	0x80000001	0x935C

Type-5 AS External Link States					
Link ID	ADV Router	Age	Seq#	Checksum	Tag
9.0.0.0	1.1.1.1	135	0x80000001	0x3AE8	0

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-3-29

The figure illustrates the use of the **show ip ospf database** command to get information about an OSPF LSDB for external routes. The router link states are type 1 LSAs, and the summary net link states are type 3 LSAs. Additionally, the ASBR creates external LSAs (type 5) to advertise external routes into OSPF. External LSAs are flooded unaltered into all areas. However, the ASBR is not in Area 0. Routers in Area 0 do not know how to reach the ASBR. To advertise the reachability of an ASBR into other areas, the ABR creates ASBR summary LSAs (type 4).

The output in the figure is partial output from ABR R2. The complete command output from this router is as follows:

```
R2#show ip ospf database
```

```
      OSPF Router with ID (2.2.2.2) (Process ID 2)
```

Router Link States (Area 0)

Link ID	ADV Router	Age	Seq#	Checksum	Link count
2.2.2.2	2.2.2.2	93	0x80000020	0xCD0B	2
3.3.3.3	3.3.3.3	1225	0x8000000D	0x9057	2

Summary Net Link States (Area 0)

Link ID	ADV Router	Age	Seq#	Checksum
4.0.0.0	2.2.2.2	73	0x80000001	0xFFE6
5.0.0.0	2.2.2.2	1651	0x80000006	0x8466

Summary ASB Link States (Area 0)

Link ID	ADV Router	Age	Seq#	Checksum
1.1.1.1	2.2.2.2	74	0x80000001	0x935C

```
1.1.1.1      2.2.2.2      74      0x80000001      0x935C
```

Router Link States (Area 1)

Link ID count	ADV Router	Age	Seq#	Checksum	Link
1.1.1.1	1.1.1.1	89	0x80000011	0xFF59	3
2.2.2.2	2.2.2.2	88	0x80000033	0x2130	2


Summary Net Link States (Area 1)

Link ID	ADV Router	Age	Seq#	Checksum
6.0.0.0	2.2.2.2	94	0x8000001F	0xCC43

Type-5 AS External Link States

Link ID	ADV Router	Age	Seq#	Checksum	Tag
9.0.0.0	1.1.1.1	135	0x80000001	0x3AE8	0

OSPF LSDB: NSSA



```
R1#
interface Loopback0
ip address 1.1.1.1 255.0.0.0

interface Serial2/1/0
ip address 5.0.0.1 255.0.0.0

interface Ethernet2/0/0
ip address 4.0.0.1 255.0.0.0

router ospf 4
redistribute static metric 5 metric-type 1
network 5.0.0.0 0.255.255.255 area 1
network 4.0.0.0 0.255.255.255 area 1
area 1 nssa

ip route 9.0.0.0 255.0.0.0 4.0.0.2
```

```
R2#
interface Loopback0
ip address 2.2.2.2 255.0.0.0
interface Serial0/1/0
ip address 5.0.0.2 255.0.0.0
interface ATM1/0.20
ip address 6.0.0.2 255.0.0.0
router ospf 2
network 5.0.0.0 0.255.255.255 area 1
network 6.0.0.0 0.255.255.255 area 0
area 1 nssa
```

```
R3#
interface Loopback0
ip address 3.3.3.3 255.0.0.0
interface ATM2/0.20 point-to-point
ip address 6.0.0.3 255.0.0.0
router ospf 2
network 6.0.0.0 0.255.255.255 area 0
```

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-3-30

The figure presents the topology that will be used to describe the OSPF LSDB for NSSA routing. All routers are configured for OSPF routing; the **show ip ospf database** command is used to get information about an OSPF LSDB. The router link states that are presented in the next figure are type 1, type 3, type 5, and type 7 LSAs.

OSPF LSDB: NSSA (Cont.)

```
R2#show ip ospf database
  OSPF Router with ID (2.2.2.2) (Process ID 2)
```

Router Link States (Area 0)					
Link ID	ADV Router	Age	Seq#	Checksum	Link count
2.2.2.2	2.2.2.2	1235	0x8000001D	0xD9FF	
3.3.3.3	3.3.3.3	1100	0x8000000B	0x9455	

Summary Net Link States (Area 0)					
Link ID	ADV Router	Age	Seq#	Checksum	
4.0.0.0	2.2.2.2	1979	0x80000002	0xFDE7	
5.0.0.0	2.2.2.2	1483	0x80000004	0x8864	

Type-7 AS External Link States (Area 1)						
Link ID	ADV Router	Age	Seq#	Checksum	Tag	
9.0.0.0	1.1.1.1	334	0x80000005	0xD738	0	

Type-5 AS External Link States						
Link ID	ADV Router	Age	Seq#	Checksum	Tag	
9.0.0.0	2.2.2.2	1725	0x80000004	0x50C6	0	

The figure illustrates the use of the **show ip ospf database** command to get information about an OSPF LSDB for NSSA routing. The router link states are type 1 LSAs, and the summary net link states are type 3 LSAs; both are used to advertise routes from one area into another. To advertise external routes into an NSSA, the ASBR also creates NSSA external LSAs (type 7). The ABR converts type 7 LSAs into type 5 LSAs and propagates the type 5 LSAs into normal areas. The ASBR summary LSAs are not needed in this case, because the ABR originates the external LSA and is reachable within Area 0. You can compare this example with a scenario in which the NSSA is a normal area by looking at the previous database example.

The output in the figure is partial output from ABR R2. The complete command output from this router is as follows:

```
R2#show ip ospf database
```

```
  OSPF Router with ID (2.2.2.2) (Process ID 2)
```

Router Link States (Area 0)

Link ID	ADV Router	Age	Seq#	Checksum	Link count
2.2.2.2	2.2.2.2	1235	0x8000001D	0xD9FF	2
3.3.3.3	3.3.3.3	1100	0x8000000B	0x9455	2

Summary Net Link States (Area 0)

Link ID	ADV Router	Age	Seq#	Checksum
4.0.0.0	2.2.2.2	1979	0x80000002	0xFDE7
5.0.0.0	2.2.2.2	1483	0x80000004	0x8864

Router Link States (Area 1)

Link ID count	ADV Router	Age	Seq#	Checksum	Link
1.1.1.1	1.1.1.1	319	0x8000000C	0xAFA8	3
2.2.2.2	2.2.2.2	220	0x8000002F	0xD478	2

Summary Net Link States (Area 1)

Link ID	ADV Router	Age	Seq#	Checksum
6.0.0.0	2.2.2.2	1483	0x8000001C	0x7894

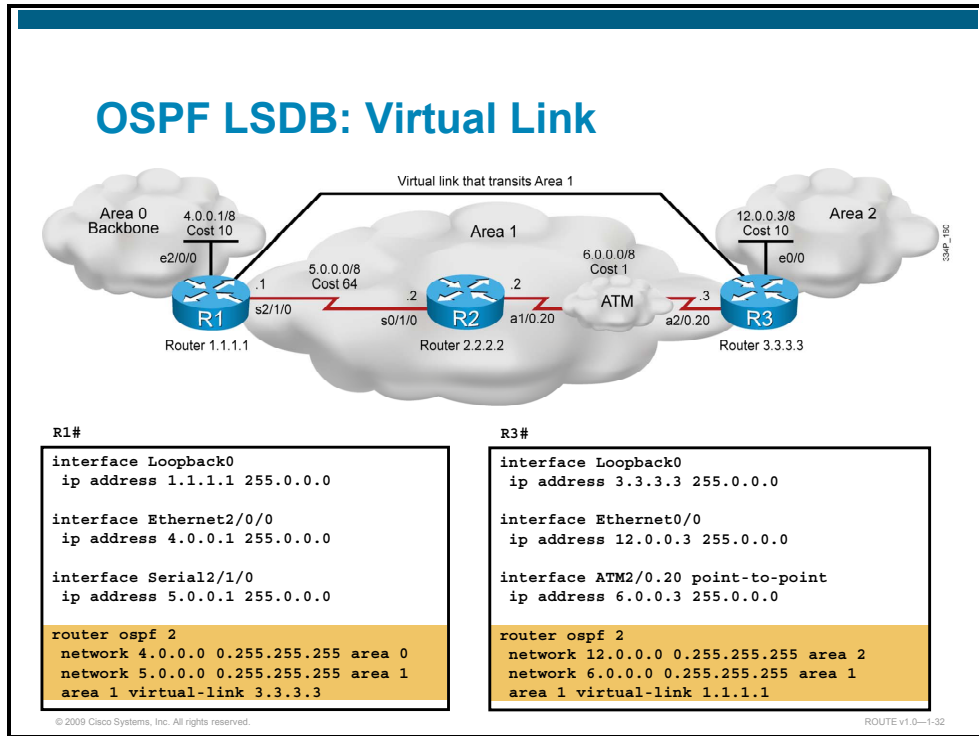
Type-7 AS External Link States (Area 1)

Link ID	ADV Router	Age	Seq#	Checksum	Tag
9.0.0.0	1.1.1.1	334	0x80000005	0xD738	0

Type-5 AS External Link States

Link ID	ADV Router	Age	Seq#	Checksum	Tag
9.0.0.0	2.2.2.2	1725	0x80000004	0x50C6	0

OSPF LSDB: Virtual Link



The figure presents the topology that will be used to describe the OSPF LSDB when virtual links are configured. All routers are configured for OSPF routing; the **show ip ospf database** command is used to get information about an OSPF LSDB. The router link states that are presented in the next figure are type 1 and type 3 LSAs.

OSPF LSDB: Virtual Link (Cont.)

```
R1#show ip ospf database
  OSPF Router with ID (1.1.1.1) (Process ID 2)

      Router Link States (Area 0)

Link ID  ADV Router  Age          Seq#          Checksum  Link
count
1.1.1.1  1.1.1.1    919          0x80000003   0xD5DF    2
3.3.3.3  3.3.3.3    5 (DNA)      0x80000002   0x3990    1

      Summary Net Link States (Area 0)

Link ID  ADV Router  Age          Seq#          Checksum
5.0.0.0  1.1.1.1    1945        0x80000002   0xAA48
5.0.0.0  3.3.3.3    9 (DNA)     0x80000001   0x7A70
6.0.0.0  1.1.1.1    1946        0x80000002   0xA749
6.0.0.0  3.3.3.3    9 (DNA)     0x80000001   0xEA3F
12.0.0.0 3.3.3.3    9 (DNA)     0x80000001   0xF624
```

©2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-3-33

The figure illustrates the use of the **show ip ospf database** command to get information about an OSPF LSDB when virtual links are configured. The router link states are type 1 LSAs, and the summary net link states are type 3 LSAs; both are used to advertise routes from one area into another.

Notice that LSAs that are learned through the virtual link have the DoNotAge (DNA) option. The virtual link is treated like a demand circuit.

The output in the figure is partial output from R1. The complete command output from this router is as follows:

```
R1#show ip ospf database
  OSPF Router with ID (1.1.1.1) (Process ID 2)

      Router Link States (Area 0)

Link ID  ADV Router  Age          Seq#          Checksum  Link count
1.1.1.1  1.1.1.1    919          0x80000003   0xD5DF    2
3.3.3.3  3.3.3.3    5 (DNA)      0x80000002   0x3990    1

      Summary Net Link States (Area 0)

Link ID  ADV Router  Age          Seq#          Checksum
5.0.0.0  1.1.1.1    1945        0x80000002   0xAA48
5.0.0.0  3.3.3.3    9 (DNA)     0x80000001   0x7A70
6.0.0.0  1.1.1.1    1946        0x80000002   0xA749
6.0.0.0  3.3.3.3    9 (DNA)     0x80000001   0xEA3F
12.0.0.0 3.3.3.3    9 (DNA)     0x80000001   0xF624

      Router Link States (Area 1)
```

Link ID	ADV Router	Age	Seq#	Checksum	Link count
1.1.1.1	1.1.1.1	1946	0x80000005	0xDDA6	2
2.2.2.2	2.2.2.2	10	0x80000009	0x64DD	4
3.3.3.3	3.3.3.3	930	0x80000006	0xA14C	2

Summary Net Link States (Area 1)

Link ID	ADV Router	Age	Seq#	Checksum
4.0.0.0	1.1.1.1	1947	0x80000002	0x9990
4.0.0.0	3.3.3.3	911	0x80000001	0xEBF5
12.0.0.0	1.1.1.1	913	0x80000001	0xBF22
12.0.0.0	3.3.3.3	931	0x80000001	0xF624

R3 considers itself an ABR, because it has a link to Area 0 (the virtual link). As a result, it generates a summary LSA for 12.0.0.0 into Area 1 and Area 0, which you can see when you issue the **show ip ospf database summary** command on R3:

R3#show ip ospf database summary 12.0.0.0

OSPF Router with ID (3.3.3.3) (Process ID 2)

Summary Net Link States (Area 0)

```

LS age: 1779
Options: (No TOS-capability, DC)
LS Type: Summary Links(Network)
Link State ID: 12.0.0.0 (summary Network Number)
Advertising Router: 3.3.3.3
LS Seq Number: 80000001
Checksum: 0xF624
Length: 28
Network Mask: /8
TOS: 0 Metric: 10

```

Summary Net Link States (Area 1)

```

LS age: 1766
Options: (No TOS-capability, DC)
LS Type: Summary Links(Network)
Link State ID: 12.0.0.0 (summary Network Number)
Advertising Router: 1.1.1.1
LS Seq Number: 80000001
Checksum: 0xBF22
Length: 28
Network Mask: /8
TOS: 0 Metric: 75

```

```

LS age: 1781
Options: (No TOS-capability, DC)

```

LS Type: Summary Links(Network)
Link State ID: 12.0.0.0 (summary Network Number)
Advertising Router: 3.3.3.3
LS Seq Number: 80000001
Checksum: 0xF624
Length: 28
Network Mask: /8
TOS: 0 Metric: 10

The show ip route Command

```
R1#show ip route
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

172.31.0.0/24 is subnetted, 2 subnets
O IA 172.31.2.0 [110/1563] via 10.1.1.1, 00:12:35, FastEthernet0/0
O IA 172.31.1.0 [110/782] via 10.1.1.1, 00:12:35, FastEthernet0/0
10.0.0.0/8 is variably subnetted, 6 subnets, 2 masks
C    10.200.200.13/32 is directly connected, Loopback0
C    10.1.3.0/24 is directly connected, Serial0/0/0
O    10.1.2.0/24 [110/782] via 10.1.3.4, 00:12:35, Serial0/0/0
C    10.1.1.0/24 is directly connected, FastEthernet0/0
O    10.1.0.0/24 [110/782] via 10.1.1.1, 00:12:37, FastEthernet0/0
O E2 10.254.0.0/24 [110/50] via 10.1.1.1, 00:12:37, FastEthernet0/0
```

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—3-34

The **show ip route** command example in the figure depicts both external type routes (O E2) and interarea (O IA) routes.

The last entry (O E2) is an external route (from the ASBR, via the ABR). The two numbers in brackets, [110/50], are the administrative distance and the total cost of the route to a specific destination network. In this case, the administrative distance is set to the default of 110 for all OSPF routes, and the total cost of the route has been calculated as 50.

Interpreting the Routing Table: Types of Routes

Router Designator		Description
O	OSPF intra-area (router LSA) and network LSA	<ul style="list-style-type: none">▪ Networks from within the area of the router▪ Advertised by means of router LSA and the network LSAs
O IA	OSPF interarea (summary LSA)	<ul style="list-style-type: none">▪ Networks from outside the area of the router, but within the OSPF AS▪ Advertised by means of summary LSAs
O E1	Type 1 external routes	<ul style="list-style-type: none">▪ Networks outside of the AS of the router▪ Advertised by means of external LSAs
O E2	Type 2 external routes	

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—3-35

The figure defines each of the routing table descriptors for OSPF. The router and network LSAs describe the details within an area. The routing table reflects this link-state information with a designation of “O,” meaning that the route is OSPF intra-area.

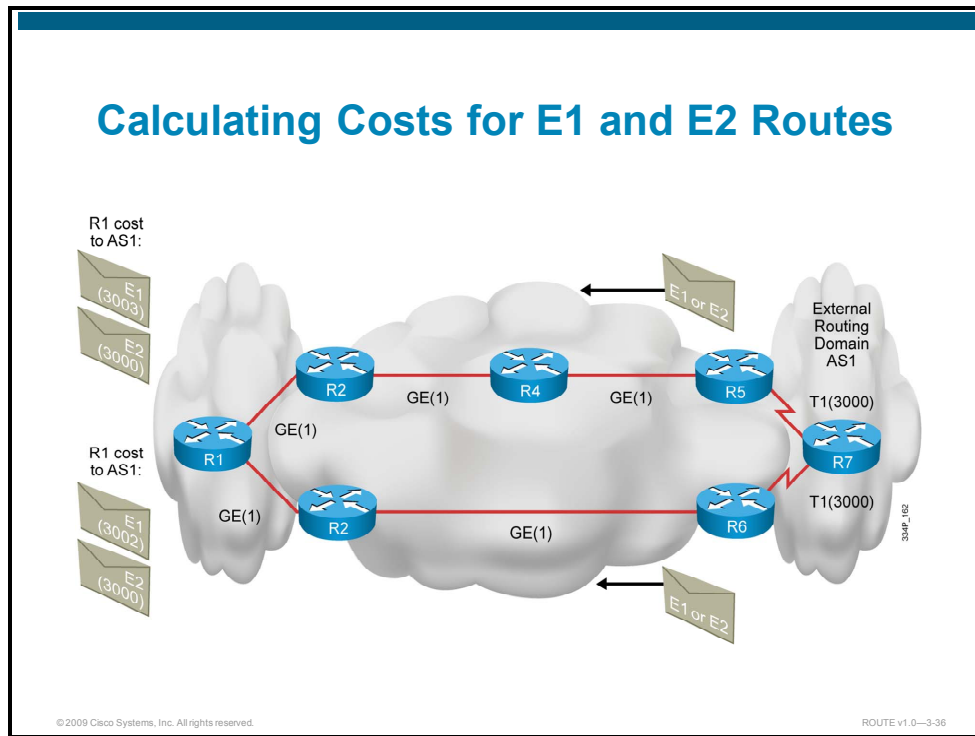
When an ABR receives summary LSAs, it adds them to its LSDB and regenerates them into the local area. When an ABR receives external LSAs, it adds them to its LSDB and floods them into the area. The internal routers then assimilate the information into their databases. Summary LSAs appear in the routing table as IA (interarea routes). External LSAs appear in the routing table marked as external type 1 (E1) or external type 2 (E2) routes.

The SPF algorithm is then run against the LSDB to build the SPF tree. The SPF tree is used to determine the best paths. The order in which the best paths are calculated is as follows:

1. All routers calculate the best paths to destinations within their areas (intra-area) and add these entries to the routing table. These are the type 1 and type 2 LSAs, which are noted in the routing table with a routing designator of “O.”
2. All routers calculate the best paths to the other areas within the internetwork. These best paths are the interarea route entries, or type 3 and type 4 LSAs, and are noted with a routing designator of O IA.
3. All routers (except those that are in a form of stub area) calculate the best paths to the external AS (type 5) destinations; these are noted with either an O E1 or an O E2 route designator, depending on the configuration.

At this point, a router can communicate with any network within or outside the OSPF AS.

Calculating Costs for E1 and E2 Routes



The cost of an external route varies, depending on the external type that is configured on the ASBR. The following external packet types can be configured:

- **E1:** Type O E1 external routes calculate the cost by adding the external cost to the internal cost of each link that the packet crosses. Use this type when there are multiple ASBRs that are advertising an external route to the same AS, to avoid suboptimal routing.
- **E2 (default):** The external cost of O E2 packet routes is always the external cost only. Use this type if only one ASBR is advertising an external route to the AS.

The figure shows the network diagram with two ASBRs (R5 and R6). Both ASBRs are sending external routes into the OSPF AS. External routes can be sent as E1 or E2. R1 receives the same external routes from R2 and R3. In the figure, the path from R1 to R6 is shorter than the path from R1 to R5. If external routes are received as E2 routes (default setting), the cost is the same regardless of the topology in the OSPF domain, which means that there will be suboptimal routing. If external routes are received as E1 routes, the cost is different, because the internal OSPF cost is added to the external cost. The routing is optimal, and the shortest path is selected to the destination.

OSPF LSDB Overload Protection

- Excessive LSAs that are generated by other routers can drain local router resources.
- This feature can limit the processing of nonself-generated LSAs for a defined OSPF process.
- Only a warning message can be sent, or neighbors can be ignored.

```
R1(config-router)#
```

```
max-lsa 12000
```

- Limit the number of nonself-generated LSAs.

©2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—3-37

If other routers are misconfigured—causing, for example, a large number of prefixes to be redistributed—large numbers of LSAs can be generated. These excessive LSAs can drain local CPU and memory resources. OSPF LSDB overload protection can be configured to protect against this issue with Cisco IOS Release 12.3(7)T and later (and some specific versions of earlier releases) by using the **max-lsa** command.

When this feature is enabled, the router keeps count of the number of received (nonself-generated) LSAs that it keeps in its LSDB. An error message is logged when this number reaches a configured threshold number, and a notification is sent when it exceeds the threshold number.

If the LSA count still exceeds the threshold after one minute, the OSPF process takes down all adjacencies and clears the OSPF database; this is called the “ignore” state. In the ignore state, no OSPF packets are sent or received by interfaces that belong to that OSPF process.

The OSPF process remains in the ignore state for the time that is defined by the **ignore-time** parameter. The **ignore-count** parameter defines the maximum number of times that the OSPF process can consecutively enter the ignore state before remaining permanently down and requiring manual intervention.

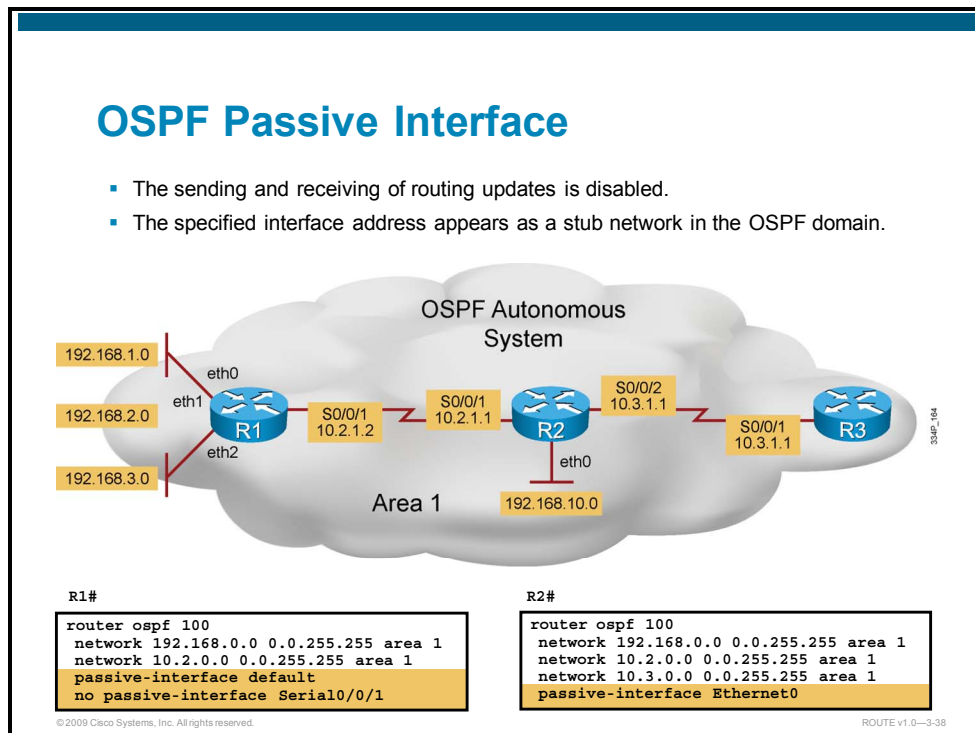
If the OSPF process remains normal for the time that is defined by the **reset-time** parameter, the ignore state counter is reset to 0.

For more details about the **max-lsa** command, go to the Cisco IOS IP Routing: OSPF Command Reference via the following link:

http://www.cisco.com/en/US/docs/ios/iproute_ospf/command/reference/iro_book.html

Limiting Adjacencies in OSPF with the passive-interface Command

This topic describes how routing advertisements can be controlled using the **passive-interface** command.



To prevent other routers on a local network from learning about routes dynamically, you can keep routing update messages from being sent through a router interface. This feature applies to all IP-based routing protocols except BGP. OSPF behaves somewhat differently. In OSPF, the interface address that you specify as passive appears as a stub network in the OSPF domain. OSPF routing information is neither sent nor received through the specified router interface.

To disable the sending and receiving of OSPF routing updates on an interface, use the **passive-interface** command in router configuration mode.

Within ISPs and large enterprise networks, many of the distribution routers have more than 200 interfaces. Therefore, a large number of LSAs can be flooded over the domain. The OSPF routing protocol can be configured on all interfaces, and the **passive-interface** command can be set on the interfaces where adjacency is not desired. In some networks, this means the coding of 100 or more passive interface statements. With the Default Passive Interface feature, this problem is solved by allowing all interfaces to be set as passive by default using a single **passive-interface default** command, then configuring individual interfaces where adjacencies are desired using the **no passive-interface** command.

In the figure, routers are configured for the OSPF routing protocol. The configuration of R1 and R2 is observed, where the passive interface feature is configured.

R1 has a set of interfaces that act as a stub network. LSAs are not received through these interfaces anyway, and there is no need for LSAs to be sent from there. The only interface that should participate in the OSPF process is interface S0/0/1. The **passive-interface default** command is used on R1, because it is easier to make all the interfaces passive and just enable one or more, where OSPF LSAs can be sent and received. In the figure, the **no passive-interface Serial0/0/1** command is used to enable the propagation of LSAs on interface S0/0/1.

R2 is slightly different. Only one interface is acting as a stub interface, where the propagation of LSAs should be stopped. The **passive-interface Ethernet0** command is stopping the propagation of LSAs through interface Ethernet0, and other interfaces are normally processing OSPF LSAs.

For more details about the **passive-interface** command, go to the Cisco IOS IP Routing: OSPF Command Reference via the following link:

http://www.cisco.com/en/US/docs/ios/iproute_ospf/command/reference/iro_book.html

Design Limitations of OSPF

This topic describes the effects of using a discontinuous backbone area, Area 0, or an area that does not connect to Area 0.

Design Limitations of OSPF

- If more than one area is configured, one of these areas has to be Area 0, the backbone area.
- All areas must be connected to Area 0.
- Area 0 must be contiguous.

© 2009 Cisco Systems, Inc. All rights reserved. ROUTE v1.0—3-39

OSPF has special restrictions when multiple areas are involved throughout the OSPF AS. If more than one area is configured, one of these areas has to be Area 0. This is called the backbone area. When designing networks, it is good practice to start from the core layer, which becomes Area 0, and then expand into other areas later.

The backbone has to be at the center of all other areas, and other areas have to be physically connected to the backbone. The main reason is that the OSPF expects all areas to inject routing information into the backbone, which distributes that information into other areas.

In the figure, it is evident that all areas are directly connected to the backbone. In the rare situations in which a new area is introduced that cannot have direct physical access to the backbone, a solution is required. Virtual links, which are discussed in the next section, must be configured in such cases.

Virtual Links as a Solution

- An extension to the backbone
- Carried by a nonbackbone area
- Cannot be created across a stub or NSSA area, or over unnumbered links
- Are used to:
 - Allow areas to connect to areas other than Area 0
 - Repair a discontinuous Area 0 (for example, if two companies merge and have separate backbone areas)

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—3-40

OSPF has a solution to extend the reach of the backbone across other areas. This solution is called the virtual link. It provides an extension to the OSPF backbone and allows a router to connect logically to the backbone, even though there is no direct physical link. Between the two routers that are involved in the creation of the virtual link, there is a nonbackbone area. The routers at each end become part of the backbone, and both act as ABRs.

The virtual link relies on intra-area routing, and its stability is dependent on the stability of the underlying area. Virtual links cannot run through more than one area or over stub areas; they can only run through normal nonbackbone areas. If a virtual link needs to be attached to the backbone across two nonbackbone areas, then two virtual links are required—one virtual link to serve each area.

Virtual links are used for two purposes:

- Linking an area that does not have a physical connection to the backbone
- Patching the backbone in case a discontinuity with Area 0 occurs


A good example of when virtual links might be required is when two companies merge that do not have a direct link between their backbone areas. A similar problem occurs when you add a nonbackbone area to an OSPF network, and the new area does not have a direct physical connection to the existing OSPF backbone area, Area 0.

OSPF Virtual Links and Solutions to Discontiguous Area Problems

This topic describes OSPF virtual links, which are used to address the discontiguous area issues, and how to configure these virtual links.

No Direct Physical Connection to Area 0

- Area 20 is added with no physical access to Area 0.
- A virtual link provides a logical path to the backbone area.
- The OSPF database treats the link between routers ABR1 and ABR2 as a direct link.




All areas in an OSPF AS must be physically connected to the backbone area (Area 0). When this is not possible, you can use a virtual link to connect to the backbone through a nonbackbone area. The area through which you configure the virtual link is known as a transit area and must have complete routing information.

In the figure, Area 20 is added to the existing OSPF topology. Because it has no direct physical link to the backbone area, Area 0, a virtual link across nonbackbone Area 10 is created. In this case, the virtual link provides a logical path between Area 20 and the backbone area.

The OSPF database treats the virtual link between ABR1 and ABR2 as a direct link. For greater stability, the loopback interface is used as a router ID and virtual links are created using these loopback addresses.

Discontiguous Area 0

- Two companies merge without a direct link between them.
- Virtual links are used to connect the discontiguous Areas 0.
- A logical link is built between ABR1 and ABR2.
- Virtual links are recommended for backup or temporary connections as well.



© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—3-42

The two-tiered area hierarchy of OSPF requires that all areas are directly connected to the backbone area, Area 0, and that Area 0 is contiguous.


A virtual link is a link that allows discontiguous areas to be connected, or that allows a disconnected area to be connected to Area 0, via a transit area. The OSPF virtual link feature should be used only in very specific cases, for temporary connections or backup after a failure. Virtual links should not be used as a primary backbone design feature.

Virtual links are part of the OSPF open standard and have been a part of Cisco IOS Software since Cisco IOS Release 10.0. In the figure, Area 0 is discontiguous, because the two companies have merged and there is no direct link between their backbone areas. A logical link (virtual link) is built between the two ABRs—routers ABR1 and ABR2. This virtual link is similar to a standard OSPF adjacency. However, in a virtual link, the routers do not have to be directly attached to neighboring routers.

The Hello protocol works over virtual links as it does over standard links, in 10-second intervals. However, LSA updates work differently on virtual links. An LSA usually refreshes every 30 minutes; LSAs that are learned through a virtual link have the DNA option set, so that the LSA does not age out. This DNA technique is required to prevent excessive flooding over the virtual link.

OSPF Virtual Link Configuration

- Configure a virtual link.
- The router ID of the remote router is used in the command.



ABR1#

```
router ospf 100
network 172.16.0.0 0.0.255.255 area 1
network 10.0.0.0 0.255.255.255 area 0
area 1 virtual-link 10.2.2.2
```

ABR2#

```
router ospf 100
network 172.16.0.0 0.0.255.255 area 1
network 10.0.0.0 0.255.255.255 area 0
area 1 virtual-link 10.1.1.1
```

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—3-43

Use the **area virtual-link** command in router configuration mode, along with any necessary optional parameters, to define an OSPF virtual link. The configuration must be done on both sides of the virtual link, and the **area virtual-link** command on each side must include the router ID of the far-end router. To find the router ID of the far-end router, use the **show ip ospf** command, **show ip ospf interface** command, or **show ip protocol** command on that remote router.

In the figure, two companies have merged, and because there is no direct physical link, Area 0 is discontinuous. A virtual link is used as a strategy to temporarily connect the Company 1 and Company 2 backbone areas. Area 1 is used as the transit area.

ABR1 builds a virtual link to ABR2, and ABR2 builds a virtual link to ABR1. Each router points to the router ID of the other router.

For more details about the **area virtual-link** command, go to the Cisco IOS IP Routing: OSPF Command Reference via the following link:

http://www.cisco.com/en/US/docs/ios/iproute_ospf/command/reference/iro_book.html

Virtual Link Verification

ABR1#

```
show ip ospf virtual-links
```

- Verify the configuration of the virtual link.

```
ABR1#show ip ospf virtual-links
```

```
Virtual Link OSPF_VL0 to router 10.2.2.2 is up
Run as demand circuit
DoNotAge LSA allowed.
Transit area 1, via interface Serial0/0/1, Cost of using 781
Transmit Delay is 1 sec, State POINT_TO_POINT,
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
Hello due in 00:00:07
Adjacency State FULL (Hello suppressed)
Index 1/2, retransmission queue length 0, number of retransmission 1
First 0x0(0)/0x0(0) Next 0x0(0)/0x0(0)
Last retransmission scan length is 1, maximum is 1
Last retransmission scan time is 0 msec, maximum is 0 msec
```

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-3-44

Use the **show ip ospf virtual-links** command to verify that the configured virtual link works properly. In the figure, it is evident that the virtual link toward the neighboring router with the router ID 10.2.2.2 is up, and Area 1 is used as a transit area. The output shows also that interface Serial0/0/1 is used to form the virtual link as well as several OSPF timers.

Other commands that are useful when troubleshooting virtual links are **show ip ospf neighbor**, **show ip ospf database**, and **debug ip ospf adj**. Routers that are across a virtual link become adjacent and exchange LSAs via the virtual link in a way that is similar to how they exchange LSAs over a physical link. Example output from the **show ip ospf neighbor** command follows:

```
ABR1#show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
10.2.2.2	0	FULL/ -	-	172.16.1.2	OSPF_VL0
10.2.2.2	0	FULL/ -	00:00:32	172.16.1.2	Serial0/0/1

For more details about the **show ip ospf virtual-links** and **show ip ospf neighbor** commands, go to the Cisco IOS IP Routing: OSPF Command Reference via the following link:

http://www.cisco.com/en/US/docs/ios/iproute_ospf/command/reference/iro_book.html

Virtual Link Verification in OSPF LSDB

ABR1#

```
show ip ospf database
```

- Verify the virtual link in the OSPF database.

```
ABR1#show ip ospf database
      OSPF Router with ID (10.1.1.1) (Process ID 100)

      Router Link States (Area 0)

Link ID  ADV Router  Age          Seq#          Checksum  Link
count
10.1.1.1 10.1.1.1    718         0x80000002   0x189A    2
10.2.2.2 10.2.2.2    4           (DNA)        0x80000001 0x2980    1

      Summary Net Link States (Area 0)
<output omitted>
```

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—3-45

The **show ip ospf database** command shows the virtual link as one of the router links. The LSAs that are learned through the virtual link have the DNA option.

Example output from the **show ip ospf database router 10.2.2.2** command shows the information from ABR2. The output of this show command for the virtual link LSA is as follows:

```
ABR1#show ip ospf database router 10.2.2.2
```

```
      OSPF Router with ID (10.1.1.1) (Process ID 1000)
```

```
      Router Link States (Area 0)
```

```
Routing Bit Set on this LSA
LS age: 1 (DoNotAge)
Options: (No TOS-capability, DC)
LS Type: Router Links
Link State ID: 10.2.2.2
Advertising Router: 10.2.2.2
LS Seq Number: 80000003
Checksum: 0x8380
Length: 48
Area Border Router
Number of Links: 2
```

Link connected to: a Virtual Link

```
(Link ID) Neighboring Router ID: 10.1.1.1
(Link Data) Router Interface address: 172.16.1.2
Number of TOS metrics: 0
TOS 0 Metrics: 781
```

Link connected to: a Transit Network

(Link ID) Designated Router address: 10.1.2.2
(Link Data) Router Interface address: 10.1.2.2
Number of TOS metrics: 0
TOS 0 Metrics: 1

Router Link States (Area 1)

Routing Bit Set on this LSA
LS age: 1688
Options: (No TOS-capability, DC)
LS Type: Router Links
Link State ID: 10.2.2.2
Advertising Router: 10.2.2.2
LS Seq Number: 80000008
Checksum: 0xCC81
Length: 48
Area Border Router
Virtual Link Endpoint
Number of Links: 2

Link connected to: another Router (point-to-point)
(Link ID) Neighboring Router ID: 10.1.1.1
(Link Data) Router Interface address: 172.16.1.2
Number of TOS metrics: 0
TOS 0 Metrics: 781

Link connected to: a Stub Network
(Link ID) Network/subnet number: 172.16.1.0
(Link Data) Network Mask: 255.255.255.0
Number of TOS metrics: 0
TOS 0 Metrics: 781

Changing the Cost Metric

This topic explains how to change the cost metric from the default values.

OSPF Cost

- The cost, or metric, is an indication of the overhead to send packets over an interface.
- OSPF cost is used as the route selection criteria.
- Dijkstra's algorithm determines the best path by adding all link costs along a path.
- OSPF cost is computed automatically:
 - Cost = $10^8 / \text{bandwidth (in b/s)}$
 - Bandwidth is specified on the interface with the **bandwidth** command.
- OSPF cost is recomputed after every bandwidth change.

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—3-46

The OSPF cost is an indication of the overhead to send packets over an interface. OSPF cost is computed automatically for each interface that is assigned into an OSPF process, using the following formula:

$$\text{Cost} = 10^8 / \text{bandwidth}$$

(The bandwidth is specified on the interface with the **bandwidth** interface command.)

The cost value is a 16-bit positive number between 1 and 65,535, where a lower value is a more desirable metric. For example, a 64-kb/s link gets a metric of 1562, while a T1 link gets a metric of 64. Cost is applied on all router link paths, and route decisions are made on the total cost of a path. The metric is only relevant on an outbound path; route decisions are not made for inbound traffic. The OSPF cost is recomputed after every bandwidth change, and Dijkstra's algorithm determines the best path by adding all link costs along a path.

On high-bandwidth links (155 Mb/s and more), automatic cost assignment no longer works (it would result in all costs being equal to 1). On these links, OSPF costs must be set manually on each interface.

Changing the Default OSPF Cost

R1 (config-if) #


```
ip ospf cost 10
```

- Changes the OSPF cost on the specified interface to 10

R2 (config-router) #

```
auto-cost reference-bandwidth 10000
```

- Changes the reference bandwidth that is used to compute the default OSPF costs from 100 to 10000



In general, the OSPF cost in Cisco routers is calculated using the following formula:

$$(100 \text{ Mb/s}) / (\text{bandwidth in Mb/s})$$

However, the cost is calculated based on a maximum bandwidth of 100 Mb/s, which is a cost of 1. If you have faster interfaces, you may want to recalibrate the cost of 1 to a higher bandwidth.

When you are using the bandwidth of the interface to determine OSPF cost, always remember to use the **bandwidth** interface command to define the bandwidth of the interface (in kb/s) accurately. The **bandwidth** command is the reference for calculating the cost.

To override the automatically calculated default cost, manually define the cost using the **ip ospf cost** interface command on a per-interface basis. The cost value is an integer from 1 to 65,535. The lower the number, the better and more preferred the link. In the figure, the **ip ospf cost 10** command manually defines the cost of the Gigabit Ethernet interface.

Also in the figure, interfaces that are faster than 100 Mb/s are being used (Gigabit Ethernet and Fast Ethernet). The automatic cost assignment is returning the value 1. For all faster links, a new reference value must be configured. Use the **auto-cost** command in router configuration mode on all routers in the network to ensure accurate route calculations.

The **reference-bandwidth** value in the **auto-cost** command is a reference bandwidth in megabits per second; it ranges from 1 to 4,294,967, with a default value of 100. In the figure, the **reference-bandwidth** value in R2 is set to 10,000. The cost of the Fast Ethernet link is changed to 100, while the Gigabit Ethernet link cost is changed to 10. Therefore, the link costs are differentiated.

For more details about the **ip ospf cost** and **auto-cost** commands, go to the Cisco IOS IP Routing: OSPF Command Reference via the following link:

http://www.cisco.com/en/US/docs/ios/iproute_ospf/command/reference/iro_book.html

Summary

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

Summary

- There are four router types: internal routers, backbone routers, ABRs, and ASBRs.
- The configuration of OSPF is a two-step process:
 - Define one or more OSPF processes in the router.
 - Define the interfaces that OSPF will run on.
- OSPF selects a router ID at startup time:
 - Define the router ID with the **router-id** command.
 - If you do not define the router ID, and there is a loopback interface, the highest IP address of the loopback interface is used.
 - If you do not define the router ID, and there is no loopback interface, the highest IP address of all active interfaces is used.

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—3-48

Summary (Cont.)

- Use the **show ip protocols**, **show ip route ospf**, **show ip ospf interface**, **show ip ospf**, **show ip ospf neighbor**, and **show ip ospf database** commands to verify OSPF operation.
- There are 11 OSPF LSA types. The following are the most commonly used: type 1 router, type 2 network, type 3 and 4 summary, type 5 external, and type 7 external.
- To prevent other routers on a local network from learning about routes dynamically, you can keep routing update messages from being sent through a router interface by using the **passive-interface** command.

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—3-49

Summary (Cont.)

- All OSPF areas must be connected to a backbone area, Area 0, which must be contiguous.
- A virtual link allows discontinuous Areas 0 to be connected, or a disconnected area to be connected to Area 0 via a transit area. Virtual links should be used only for temporary connections or backup after a failure, not as a primary backbone design feature.
- The OSPF cost defaults to $(100 \text{ Mb/s}) / (\text{bandwidth in megabits per second})$. The cost can be changed on a per-interface basis, and so can the reference bandwidth (100 Mb/s).

Lab 3-1 Debrief

Overview

In Lab 3-1, you configured and verified OSPF to improve routing performance. First, you used OSPF to exchange routing information to achieve IP connectivity over the LAN interfaces. You then established reachability and configured OSPF on the WAN segment. The OSPF interface network type should reflect the Frame Relay WAN segment representation. You then configured OSPF over Frame Relay using the point-to-point and multipoint OSPF network types.

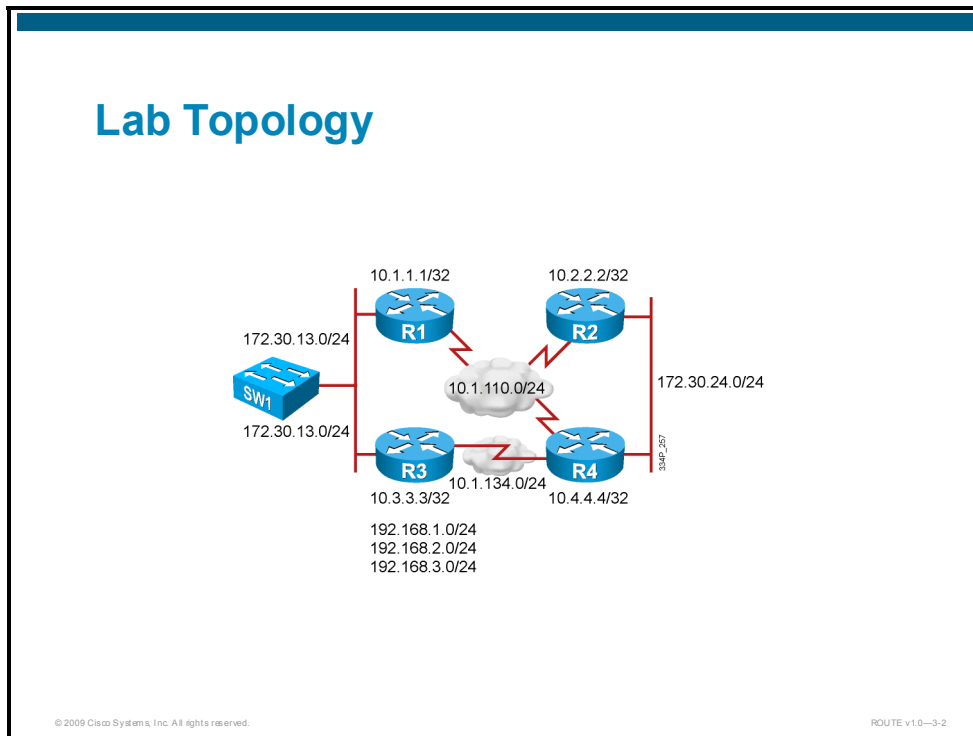
Objectives

Upon completing this lesson, you will be able to implement and verify OSPF to improve routing performance. This ability includes being able to meet these objectives:

- Identify the implementation, verification, and troubleshooting tasks for OSPF operations over LAN and WAN interfaces with different Frame Relay WAN segment representations
- Present a sample solution and identify possible alternative solutions

Lab Overview and Verification

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



The figure presents the physical lab topology that is used for Lab 3-1: “Configure and Verify OSPF to Improve Routing Performance.” The topology uses four pod routers.

Based on the topology, you have to identify the required parameters and configure an OSPF routing protocol to establish Layer 3 reachability in the network.

Lab Review: What Did You Accomplish?

- **Task 1:** Configure OSPF over LAN interfaces.
 - Which steps did you take to configure OSPF and advertise all the specific IP subnets used in the network?
 - How can you change the configuration to prevent any network added to the router from being advertised?
 - How does this change the default OSPF DR and BDR selection on the LAN segment?
- **Task 2:** Configure OSPF over Frame Relay using the point-to-point OSPF network type.
 - What must you change on R3 to be able to add a WAN segment to the OSPF in the same area as the LAN segment?
 - What must you change on R4 to be able to add a WAN segment to the OSPF in the same area as the LAN segment and where only the WAN segment subnet is advertised?

Lab Review: What Did You Accomplish? (Cont.)

- **Task 3:** Configure OSPF over Frame Relay using the point-to-multipoint OSPF network type.
 - How does the configuration reflect the Frame Relay multipoint network type when you include the WAN segment on R1 into OSPF and enable OSPF adjacency over the serial interface toward R2 and R4?
 - How do you change the configuration to also advertise, from the major network, any networks that are added later? This is the OSPF configuration on R2 and R4, where the WAN segment toward R1 is included in the OSPF.

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0–3-4

In the first task, you configured OSPF over LAN interfaces. You configured OSPF and advertised all the specific IP subnets that are used in the network. You then prevented any network that is added to the router from being advertised. To change the default OSPF designated router (DR) and backup designated router (BDR) selection on the LAN segment, you manipulated the OSPF router priority.

In the second task, you configured OSPF over Frame Relay using the point-to-point OSPF network type. You added the WAN segment to the OSPF in the same area as the LAN segment. You then configured OSPF to advertise the WAN segment subnet only.

In the third task, you configured OSPF over Frame Relay using the point-to-multipoint OSPF network type. You configured OSPF on the WAN segments on routers R2 and R4 toward R1, reflecting the Frame Relay multipoint network type. You configured OSPF to advertise any networks from the major network that are added later.

Verification

- Did you have enough information to create an implementation plan?
- Are the proper adjacencies established on the LAN segment?
- For how long are the adjacencies up?
- Do you see the proper OSPF routes in the IP routing table when comparing the IP routing table to the OSPF database?
- Which router is the DR on the LAN segment between R1 and R3?
- Are the proper adjacencies established on the WAN segment between R3 and R4?
- For how long are the adjacencies up?

Verification (Cont.)

- After enabling OSPF on the WAN segment, do you see the proper OSPF routes in the IP routing table when comparing the IP routing table to the OSPF database?
- Are the proper adjacencies established on the WAN segment between R1, R2, and R4 in your pod?
- For how long are the adjacencies up?
- After enabling OSPF on the WAN segment on R2 and R4, do you see the proper OSPF routes in the IP routing table when comparing the IP routing table to the OSPF database?

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0--3-6

A common approach to verifying the implementation process for a routing protocol is to answer the following questions:

- Did you have enough information to create an implementation plan?
- Are the proper adjacencies established on the LAN segment?
- For how long are the adjacencies up?
- Do you see the proper OSPF routes in the IP routing table when comparing the IP routing table to the OSPF database?
- Which router is the DR on the LAN segment between R1 and R3?
- Are the proper adjacencies established on the WAN segment between R3 and R4?
- For how long are the adjacencies up?
- After enabling OSPF on the WAN segment, do you see the proper OSPF routes in the IP routing table when comparing the IP routing table to the OSPF database?
- Are the proper adjacencies established on the WAN segment between R1, R2, and R4 in your pod?
- For how long are the adjacencies up?
- After enabling OSPF on the WAN segment on R2 and R4, do you see the proper OSPF routes in the IP routing table when comparing the IP routing table to the OSPF database?

Checkpoints

- Configure OSPF on the LAN segment.
- Advertise only the specific subnets used in the network.
- Check for the DR on the LAN segment.
- Configure OSPF on the WAN segment between R3 and R4.
- Advertise only the specific subnets used in the network.
- Check for the correct OSPF network type on the WAN segments.
- Configure OSPF on the WAN segments between R1, R2, and R4.
- Check for the proper entries in the IP routing table.

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v10-3-7

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


- Configure OSPF on the LAN segment.
- Advertise only the specific subnets that are used in the network.
- Check for the DR on the LAN segment.
- Configure OSPF on the WAN segment between R3 and R4.
- Advertise only the specific subnets that are used in the network.
- Check for the correct OSPF network type on the WAN segments.
- Configure OSPF on the WAN segments between R1, R2, and R4.
- Check for the proper entries in the IP routing table.

Sample Solution and Alternatives

This topic describes a sample solution and possible alternatives.

Sample Solution

- Configure OSPF on the LAN segment, and advertise only the specific subnets used in the network.
- Configure OSPF on the WAN segment between R3 and R4, and advertise only the specific subnets used in the network.
- Configure OSPF on the WAN segments between R1, R2, and R4.




A sample solution includes the implementation details and the details for each task of the implementation plan. Different solutions are possible; the figure shows some of the details of a successful configuration.

The proper implementation of OSPF configuration and verification to improve routing performance includes the following details:

- Configure OSPF on the LAN segment, and advertise only the specific subnets that are used in the network.
- Configure OSPF on the WAN segment between R3 and R4, and advertise only the specific subnets that are used in the network.
- Configure OSPF on the WAN segments between R1, R2, and R4.

Alternative Solutions

- Use several routing protocols to provide reachability in the network.
- Changing the routing protocol is not a realistic solution, although you can modify the existing solution.
- Fulfill the same requirements by reconfiguring to different Frame Relay types.



© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-3-9

You can use several routing protocols to provide reachability in the network. Changing the routing protocol is not a realistic solution, because it is not a common practice for fine-tuning the existing protocol. You can also fulfill the same requirements by reconfiguring to different Frame Relay types.

Q and A

1. Why is routing protocol selection important?
2. How is a DR and BDR selection performed?
3. How is OSPF enabled on a LAN segment?
4. Is an OSPF network type important when enabling OSPF on the WAN segment?

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-3-10

1. A routing protocol, with its metric and administrative distance, exchanges routing updates and populates its IP routing table, which is used for destination-based forwarding. Different routing protocols process routing updates in different ways.
2. LAN links elect one router as the DR and another as the BDR to ensure that all the routers on the same LAN have identical databases. The OSPF priority value is used to elect a DR and BDR. The router with the highest priority value is the DR. In a tie, the router with the highest router ID becomes the DR. The highest IP address on an active interface is normally used as the router ID, which can be manipulated by configuring an IP address on a loopback interface or using the **router-id** router configuration command. Once elected, the DR passes its database to any new routers that come up.
3. OSPF on a LAN segment is configured using the **network** command in OSPF router configuration mode with the network for that LAN segment. The router priority value can be added to the particular interface to manipulate DR and BDR election.
4. It is important to specify the OSPF network type on the WAN segment, because the OSPF network type defines how the OSPF adjacencies will be established, as well as how OSPF routing packets will be propagated.

Summary

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

Summary

- Configure OSPF on the LAN segment to advertise only specific subnets and to observe DR selection.
- Configure OSPF on the WAN segment to advertise only specific subnets.
- Configure OSPF on the WAN segment to advertise any additional networks from the major network, if added later.

Lab 3-2 Debrief

Overview

In Lab 3-2, you implemented and verified OSPF multiarea routing. First, you defined the backbone area, where the selected routers are configured for the OSPF backbone area. You then configured the remaining routers for nonbackbone normal areas, because there is no need to filter any subnets. You then tuned the OSPF configuration by manipulating the cost and router ID. You could perform additional optimization by suppressing routing updates.

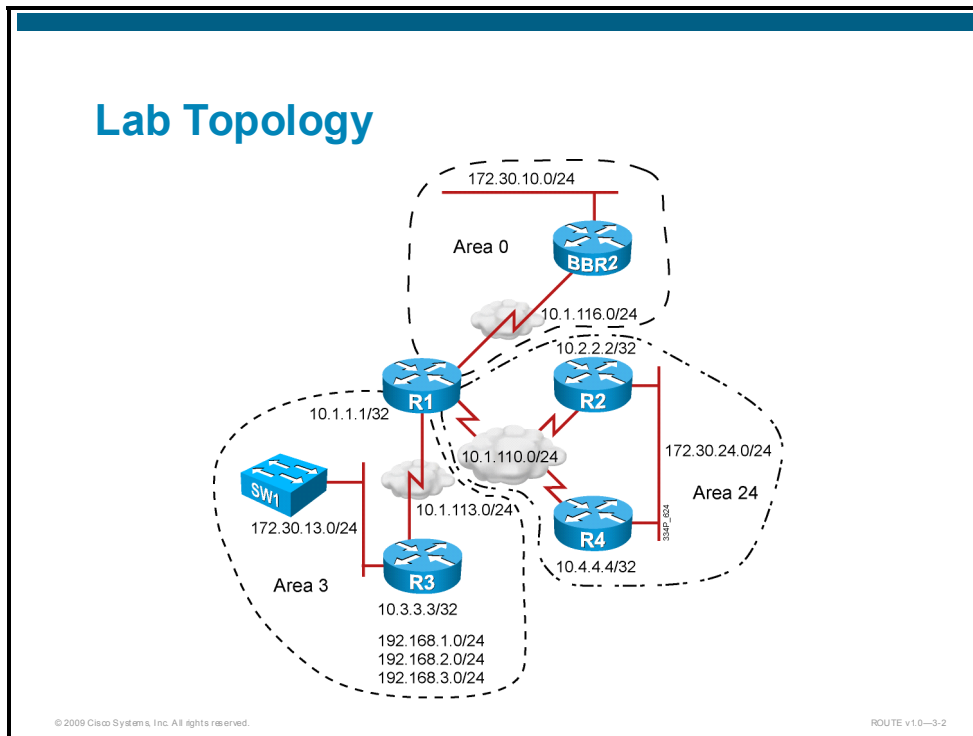
Objectives

Upon completing this lesson, you will be able to implement and verify OSPF multiarea routing. This ability includes being able to meet these objectives:

- Identify the implementation, verification, and troubleshooting tasks for OSPF operations in multiarea over LAN and WAN interfaces with different Frame Relay WAN segment representations
- Present a sample solution and identify possible alternative solutions

Lab Overview and Verification

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



The figure presents the physical lab topology that is used for Lab 3-2: “Implement and Verify OSPF Multiarea Routing.” The topology uses four pod routers and one backbone router.

Based on this topology, you have to divide the network into more OSPF areas (backbone and nonbackbone) and apply the correct configuration to each router.

Lab Review: What Did You Accomplish?

- **Task 1:** Configure an OSPF backbone area.
 - Which steps did you take to configure OSPF on a router belonging to the backbone area?
- **Task 2:** Configure OSPF nonbackbone areas.
 - Which steps did you take to configure OSPF on routers belonging to different nonbackbone areas?
- **Task 3:** Tune OSPF operation.
 - How can the default cost calculation be changed?
 - How can the router ID be changed?
 - How can you preserve CPU cycles on R3 by eliminating the unnecessary OSPF traffic on a LAN segment?

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-3-3

In the first task, you configured the OSPF backbone area, which means configuring routers in the backbone area. This task implied selecting the routers that would have one or several interfaces in the backbone area, and configuring the correct area type on each interface.

In the second task, you configured the OSPF nonbackbone areas and configured the routers in the nonbackbone areas. Some routers represent a nonbackbone area where they have all or just a few interfaces in a nonbackbone area.

In the third task, you tuned OSPF operation. You had to change the default cost calculation to manipulate the path for the packets. You had to change the router ID to manipulate the DR and BDR selection. Finally, you had to configure some interfaces as passive to preserve the CPU cycles on the router by eliminating unnecessary OSPF traffic.

Verification

- Did you have enough information to create an implementation plan?
- Is an adjacency in Area 0 between R1 and BBR2 established?
- Is the IP routing table populated with the correct OSPF routes?
- Is an adjacency established between routers in the nonbackbone areas?
- Is the IP routing table populated with the correct OSPF routes?
- How is a cost calculation change done?
- What is the router ID of R1?
- Is the IP routing table populated with the correct OSPF routes?
- Did R3 stop trying to set up an OSPF adjacency via the LAN segment?

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-3-4

A common approach to verifying the implementation process for a routing protocol is to answer the following questions:

- Did you have enough information to create an implementation plan?
- Is an adjacency in Area 0 between routers R1 and BBR2 established?
- Is the IP routing table populated with the correct OSPF routes?
- Is an adjacency established between routers in the nonbackbone areas?
- Is the IP routing table populated with the correct OSPF routes?
- How is a cost calculation change done?
- What is the router ID of R1?
- Is the IP routing table populated with the correct OSPF routes?
- Did R3 stop trying to set up an OSPF adjacency via the LAN segment?

Checkpoints

- Configure OSPF in Area 0.
- Check for adjacencies in Area 0 between R1 and BBR2.
- Check the IP routing table for the proper OSPF routes.
- Configure OSPF in nonbackbone areas.
- Check if an adjacency is established between the routers of nonbackbone areas.
- Check the IP routing table for proper OSPF routes.
- Change the cost calculation.
- Manipulate the OSPF router ID of R1.
- Check the IP routing table for the proper OSPF routes.
- Check if R3 has stopped trying to set up an OSPF adjacency via the LAN segment.

© 2009 Cisco Systems, Inc. All rights reserved.

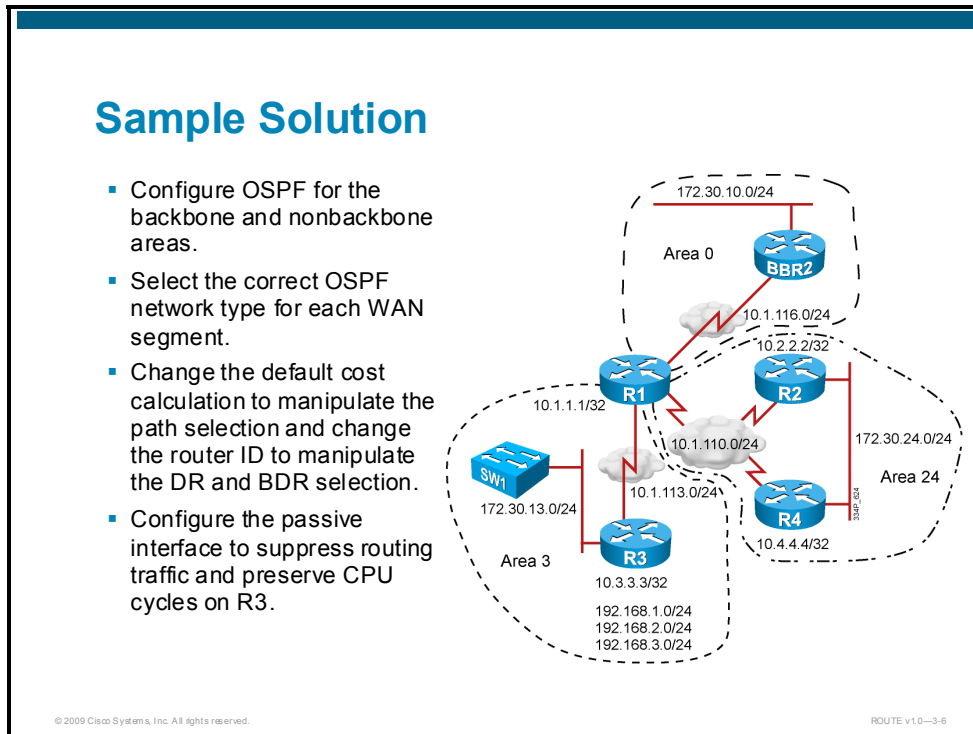
ROUTE v10-3-5

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

- Configure OSPF in Area 0.
- Check for adjacencies in Area 0 between R1 and BBR2.
- Check the IP routing table for the proper OSPF routes.
- Configure OSPF in nonbackbone areas.
- Check if an adjacency is established between the routers of nonbackbone areas.
- Check the IP routing table for the proper OSPF routes.
- Change the cost calculation.
- Manipulate the OSPF router ID of R1.
- Check the IP routing table for the proper OSPF routes.
- Check if R3 has stopped trying to set up an OSPF adjacency via the LAN segment.

Sample Solution and Alternatives

This topic describes a sample solution and possible alternatives.




A sample solution includes the implementation details and the details for each task of the implementation plan. Different solutions are possible; the figure shows a few details of a successful configuration.

The proper implementation and verification of OSPF multiarea routing includes the following details:

- Configure OSPF for the backbone and nonbackbone areas.
- Select the correct OSPF network type for each WAN segment.
- Change the default cost calculation to manipulate the path selection and change the router ID to manipulate the DR and BDR selection.
- Configure the passive interface to suppress routing traffic and preserve CPU cycles on R3.

Alternative Solutions

- Design different backbone and nonbackbone areas, for which nonbackbone areas can be nonstandard to reduce the number of routing updates.
- Configure a different IP address on the loopback interface to manipulate the router ID.
- Because changing the routing protocol is not a realistic solution, you can configure static and default routes.



© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v10-3-7

To create a similar configuration, you can design different backbone and nonbackbone areas, for which the nonbackbone areas can be nonstandard to reduce the number of routing updates.

You can configure a different IP address on the loopback interface to manipulate the router ID.

Because it is not realistic to change the routing protocol, you can configure static and default routes to provide reachability while still preserving OSPF as a routing protocol in the network.

Q and A

1. What is the purpose of backbone and nonbackbone areas?
2. How can a default cost calculation be changed?
3. Why is the router ID important?
4. How does a passive interface work in an OSPF routing protocol?

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—3-8

1. OSPF segments the network in different areas to optimize the routing in the whole network. All nonbackbone areas must be connected to the backbone area. Routing information is exchanged between the areas in a controlled way. Summarization takes place on the ABRs, limiting the number of updates so that the small changes locally in the area are not triggering the topology change recalculation in other areas.
2. You can change the default cost calculation by changing the cost per neighbor in OSPF routing process configuration mode.
3. OSPF adds the router ID to all OSPF updates and is visible in the OSPF topology database. The router ID inside the update packet defines the source of the update packets and must be unique in the network.
4. When a passive interface is enabled, the sending and receiving of routing updates is disabled. The specified interface address appears as a stub network in the OSPF domain.

Summary

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

Summary

- Configure OSPF for a backbone area.
- Configure OSPF for a nonbackbone area.
- Tune OSPF operation by changing how default cost calculation is performed. You can change the cost calculation by changing the router ID and configuring a passive interface. Doing this preserves CPU cycles by eliminating unnecessary OSPF traffic on the LAN segment.

Configuring and Verifying OSPF Route Summarization

Overview

Scalability, improved CPU and memory utilization, and the ability to mix small routers with large routers are all benefits of using proper route summarization techniques. A key feature of the Open Shortest Path First (OSPF) protocol is the ability to summarize routes at area and autonomous system (AS) boundaries.

Route summarization is important, because it reduces the amount of OSPF link-state advertisement (LSA) flooding and the sizes of link-state databases (LSDBs) and routing tables, which also reduces memory and CPU utilization on the routers. The OSPF network can scale to very large sizes, in part because of route summarization.

Default routes reduce routing table size, and also reduce memory and CPU utilization. OSPF injects a default route unconditionally or based on the presence of the default route inside the routing table.

This lesson defines the different types of route summarization and describes the configuration commands for each type. It also describes the benefits of default routes and how to configure them.

Objectives

Upon completing this lesson, you will be able to describe the procedure for configuring OSPF route summarization for interarea and external routes. This ability includes being able to meet these objectives:

- Describe OSPF route summarization
- Implement OSPF route summarization
- Identify the benefits of a default route in OSPF
- Use a default route in OSPF

OSPF Route Summarization

This topic describes the functions of interarea route summarization and external route summarization.

Summarization

- Networks are normally translated into type 3 LSAs in other areas.
- Route summarization is the consolidation of advertised addresses:
 - On ABR, summarize type 3 LSAs.
 - On ASBR, summarize type 5 LSAs.
- A good addressing plan is required.
- A drawback is the possibility of suboptimal routing.

© 2009 Cisco Systems, Inc. All rights reserved. ROUTE v1.0-3-2

Route summarization is a key to scalability in OSPF. Route summarization helps solve two major problems: large routing tables and frequent LSA flooding throughout the AS. Every time that a route disappears in one area, routers in other areas also get involved in shortest-path calculation. To reduce the size of the area database, you can configure summarization on an area boundary or AS boundary.

Normally, type 1 and type 2 LSAs are generated inside each area and translated into type 3 LSAs in other areas. With route summarization, the Area Border Routers (ABRs) or Autonomous System Boundary Routers (ASBRs) consolidate multiple routes into a single advertisement. ABRs summarize type 3 LSAs, and ASBRs summarize type 5 LSAs. Instead of advertising many specific prefixes, advertise only one summary prefix.


If the OSPF design includes many ABRs or ASBRs, suboptimal routing is possible. This is one of the drawbacks of summarization.

Route summarization requires a good addressing plan—an assignment of subnets and addresses that is based on the OSPF area structure and lends itself to aggregation at the OSPF area borders.

Note Recall that summary LSAs (type 3) and external LSAs (type 5) do not, by default, contain summarized routes.

Benefits of Route Summarization

- Minimizes the number of routing table entries
- Localizes the impact of a topology change
- Reduces LSA flooding and saves CPU resources



© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-3-3

Route summarization directly affects the amount of bandwidth, CPU power, and memory resources that the OSPF routing process consumes.

Without route summarization, every specific-link LSA is propagated into the OSPF backbone and beyond, causing unnecessary network traffic and router overhead.


With route summarization, only the summarized routes are propagated into the backbone (Area 0). Summarization prevents every router from having to rerun the Shortest Path First (SPF) algorithm, increases the stability of the network, and reduces unnecessary LSA flooding. Also, if a network link fails, the topology change is not propagated into the backbone (and other areas by way of the backbone). Specific-link LSA flooding outside the area does not occur.

Implementing OSPF Route Summarization

This topic describes how to configure route summarization in OSPF.

Interarea Route Summarization

- A summary route will be generated if at least one subnet within the area falls in the summary address range.
- A summarized route metric will be equal to the lowest cost of all subnets within the summary address range.
- Only for the summary routes of connected areas:
 - The ABR creates a route to Null0 to avoid loops.

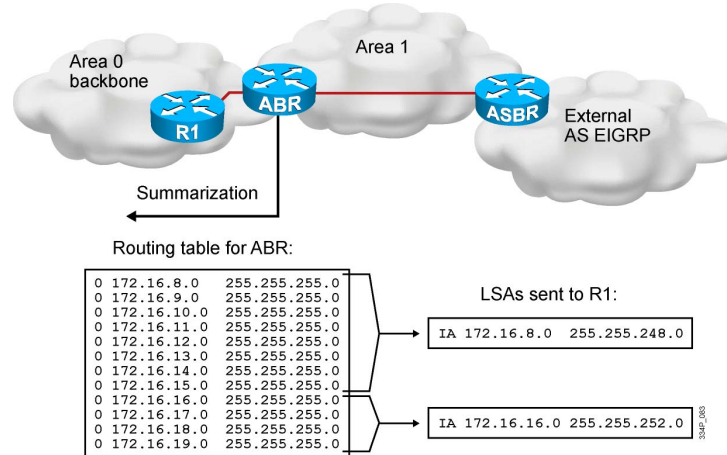


The summarization of internal routes can be done only by ABRs (router R2 in the figure). Without summarization, all the prefixes from an area are passed into the backbone as type 3 interarea routes. When summarization is enabled, the ABR intercepts this process and instead injects a single type 3 LSA, which describes the summary route into the backbone. Multiple routes inside the area are summarized.

A summary route is generated if at least one subnet within the area falls in the summary address range and the summarized route metric is equal to the lowest cost of all the subnets within the summary address range. Interarea summarization can only be done for the intra-area routes of connected areas, and the ABR creates a route to Null0 to avoid loops in the absence of more specific routes.

Using Route Summarization

- Interarea summary link carries a mask.
- One or more entries can represent several subnets.



OSPF is a classless routing protocol, which means that it carries subnet mask information along with route information. Therefore, OSPF supports multiple subnet masks for the same major network, which is known as variable-length subnet masking (VLSM).

OSPF also supports discontinuous subnets, because the subnet masks are part of the LSDB. Some older distance vector protocols do not support VLSM or discontinuous subnets.


For example, if a major network crosses the boundaries of an OSPF and an older distance vector routing protocol domain, then VLSM information that is redistributed into the older routing distance vector protocol is lost, and the static routes may have to be configured manually there.

Network numbers in areas should be assigned contiguously to ensure that these addresses can be summarized into a minimal number of summary addresses.

In the figure, the list of 12 networks in the routing table of the ABR can be summarized into two summary address advertisements. The block of addresses from 172.16.8.0 through 172.16.15.0/24 can be summarized using 172.16.8.0/21, and the block from 172.16.16.0 through 172.16.19.0/24 can be summarized using 172.16.16.0/22.

Configuring Interarea Route Summarization

- Configures type 3 summarization on ABRs



OSPF does not perform automatic summarization on major network boundaries. To consolidate and summarize routes at an area boundary, use the **area range** command in router configuration mode. The ABR will summarize routes for a specific area before injecting them into a different area via the backbone as type 3 summary LSAs.


Before summarization is possible, OSPF must be configured with the contiguous IP addressing implementation. Without proper addressing implementation, effective summarization cannot be done.

Cisco IOS Software creates a summary route to the Null0 interface when manual summarization is configured, to prevent routing loops. For example, if the summarizing router receives a packet to an unknown subnet that is part of the summarized range, the packet matches the summary route based on the longest match. The packet is forwarded to the Null0 interface (in other words, it is dropped), which prevents the router from forwarding the packet to a default route and possibly creating a routing loop.

For more details about the **area range** command, go to the Cisco IOS IP Routing: OSPF Command Reference via the following link:

http://www.cisco.com/en/US/docs/ios/iproute_ospf/command/reference/iro_book.html

Route Summarization Configuration Example at the ABR




The figure shows that route summarization can occur in both directions: from a nonbackbone area to Area 0 and from Area 0 to a nonbackbone area. In the example, the ABR configuration specifies the following summarization:

- **area 0 range 172.16.96.0 255.255.224.0:** Identifies Area 0 as the area containing the range of networks to be summarized into Area 1. Router ABR summarizes the range of subnets from 172.16.96.0 to 172.16.127.0 into one range: 172.16.96.0 255.255.224.0.
- **area 1 range 172.16.64.0 255.255.224.0:** Identifies Area 1 as the area containing the range of networks to be summarized into Area 0. Router ABR summarizes the range of subnets from 172.16.64.0 to 172.16.95.0 into one range: 172.16.64.0 255.255.224.0.

Note Depending on your network topology, you may not want to summarize Area 0 networks into other areas. For example, if you have more than one ABR between a nonbackbone area and the backbone area, sending a summary LSA (type 3) with the explicit network information into an area ensures that the shortest path to destinations outside the area is selected. If you summarize the addresses, suboptimal path selection may occur.

External Route Summarization

- Summarization can be used for external routes:
 - On an AS boundary for type 5 LSAs (redistributed routes)
 - On an NSSA ABR for type 5 translated from type 7
- A summary route to Null0 will be created for each summary range.




© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-3-8

Summarization can also be used for external routes. Each route that is redistributed into OSPF from other protocols is advertised individually with an external LSA. To reduce the size of the OSPF LSDB, you can configure a summary for external routes. Summarization of external routes can be done on the ASBR for type 5 LSAs (redistributed routes) before injecting them into the OSPF domain. Summarization of external routes can be done by the ABR in not-so-stubby areas (NSSAs) as well, where the ABR router creates type 5 summary routes from type 7 external routes. Without summarization, all the redistributed external prefixes from external autonomous systems are passed into the OSPF area. A summary route to Null0 is created automatically for each summary range.

Configuring External Route Summarization

- Configures type 5 summarization of redistributed routes on ASBRs



To create aggregate addresses for OSPF at an AS boundary, use the **summary-address** command in router configuration mode. The ASBR will summarize external routes before injecting them into the OSPF domain as type 5 external LSAs.


You must implement contiguous IP addressing before you can achieve optimal summarization.

Sometimes, only type 3 or type 5 summarization is required, but in large OSPF networks, both configurations can be applied.

For more details about the **summary-address** command, go to the Cisco IOS IP Routing: OSPF Command Reference via the following link:

http://www.cisco.com/en/US/docs/ios/iproute_ospf/command/reference/iro_book.html

Route Summarization Configuration Example at ASBR



```
ASBR (config)# router ospf 100
ASBR (config-router)# network 172.16.64.1 0.0.0.0 area 1
ASBR (config-router)# summary-address 172.16.32.0 255.255.224.0
```

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-3-10

The figure depicts route summarization on the ASBR. On the right side, an external AS that is running Enhanced Interior Gateway Routing Protocol (EIGRP) has its routes redistributed into OSPF.

Because of the contiguous subnet block in the external EIGRP network, it is possible to summarize the 32 different subnets into one summarized route. Instead of 32 external type 5 LSAs flooding into the OSPF network, there is only one.

Note EIGRP routes must also be redistributed into OSPF in this example; redistribution is covered in the “Implement an IPv4-Based Redistribution Solution” module.

Benefits of a Default Route in OSPF

Occasionally, you will be required to configure OSPF to advertise a default route into the route AS. This topic describes the benefits of a default route in OSPF.

Default Routes in OSPF

- A default route is injected into OSPF as an external type 5 LSA.
- Default route distribution is not on by default.
- Benefits of default routes include:
 - A smaller routing table
 - Fewer resources used in the router

The diagram shows an OSPF Domain containing two routers, R1 and R2. R1 is connected to Internet Service Provider A, and R2 is connected to Internet Service Provider B. R1 has a default route 0.0.0.0 with a cost of 20, and R2 has a default route 0.0.0.0 with a cost of 30. The diagram also shows the physical connections between the routers and the service providers.

© 2009 Cisco Systems, Inc. All rights reserved. R0 UTE v1.0-3-11

To be able to perform routing from an OSPF AS toward external networks or toward the Internet, you must either know all the destination networks or create a default route. To learn all the destinations, you can either create a number of static routes or configure redistribution. But the most scalable and optimized way is through the use of a default route.

The benefits of a default route include the following:

- Smaller routing table
- Fewer resources and less CPU power needed; no need to recalculate the SPF algorithm if one or more networks fail

The figure shows how OSPF injects a default route into a standard area (the different types of areas are described in the next lesson). Any OSPF router can originate the default routes that are injected into a standard area, but the OSPF routers by default do not generate a default route into the OSPF domain. For OSPF to generate a default route, you must run a configuration command.

A default route shows up in the OSPF database as a type 5 external LSA. Here is an example of how it looks in the database:

```
Type-5 AS External Link States
Link ID ADV Router  Age  Seq#           Checksum      Tag
0.0.0.0 198.1.1.1    601  0x80000001    0xD0D8       0
```

Using a Default Route in OSPF

This topic describes how to configure a default route injection into OSPF.

Configuring OSPF Default Route

- The first command allows the ASBR to originate a type 5 default route if it has the gateway of last resort.
- The second command allows the ASBR to originate a type 5 default route even if there is no gateway of last resort (optional).
- Use the route map to define a dependency on any condition inside the route map (optional).

```
ASBR (config-router) #
default-information originate
default-information originate always
```

© 2009 Cisco Systems, Inc. All rights reserved. RD UTE v1.0-3-12

To generate a default external route into an OSPF routing domain, use the **default-information originate** router configuration command, as shown in the figure. There are two ways to advertise a default route into a standard area:

- Advertise 0.0.0.0/0 into the OSPF domain when the advertising router already has a default route. Use the **default-information originate** command to allow the ASBR to originate a type 5 default route inside the OSPF AS.

You can also use different keywords in the configuration command to configure dependency on IP routing table entries:

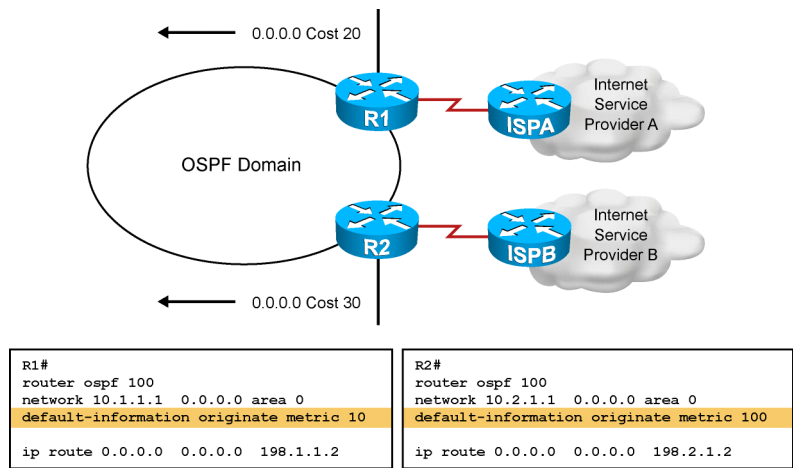
- Advertise 0.0.0.0/0, regardless of whether the advertising router already has a default route. The second method can be accomplished by adding the keyword **always** to the **default-information originate** command, as shown in the figure.

Whenever you use the **redistribute** or the **default-information** router configuration command to redistribute routes into an OSPF routing domain, the Cisco IOS Software automatically becomes an ASBR. However, an ASBR by default does not generate a default route into the OSPF routing domain. The software must still have a default route for itself before it generates one, except when you have specified the **always** keyword. You can also use a route map to define dependency on any condition inside the route map.

For more details about the **default-information originate** command, go to the Cisco IOS IP Routing: OSPF Command Reference via the following link:

http://www.cisco.com/en/US/docs/ios/iproute_ospf/command/reference/iro_book.html

Default Route Configuration Example



In the figure, an OSPF network is multi-homed to dual ISPs. Because the number of prefixes in the ISP routing table is high, it makes no sense to redistribute them into an OSPF AS. The default route is sent into the OSPF AS to provide connectivity to the external destinations. In the customer design, ISP A is preferred, and ISP B is used as a backup. To define the priority, the optional **metric** parameter has been used to establish a preference for the default route to ISP A. In the example, the default route advertised by R1 has a metric of 10. It is preferred over the default route that is advertised by R2, which has a metric of 100.

Because the **default-information originate** command is used without the optional parameter **always**, the default route must exist inside the IP routing table to advertise the default route into an OSPF domain. Both routers in the figure have a default route that is statically defined.

For more details about the **default-information originate** command, go to the Cisco IOS IP Routing: OSPF Command Reference via the following link:

http://www.cisco.com/en/US/docs/ios/iproute_ospf/command/reference/iro_book.html

Summary

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

Summary

- Route summarization improves CPU utilization, reduces LSA flooding, and reduces routing table sizes.
- The **area range** command is used to summarize at the ABR. The **summary-address** command is used to summarize at the ASBR.
- Default routes can be used in OSPF to prevent the need for a specific route to each destination network. The benefits include a much smaller routing table and an LSDB with complete reachability.
- OSPF uses the **default-information originate** command to inject a default route.

Lab 3-3 Debrief

Overview

In Lab 3-3, you configured and verified OSPF route summarization for interarea and external routes. After you performed an initial examination of the OSPF routing information and an IP routing table, it became evident that optimization was required. You had to optimize an OSPF link-state database—and, consequently, the IP routing table—to make them smaller with summarization. Once you summarized internal routes, you could do additional work for external routes as well.

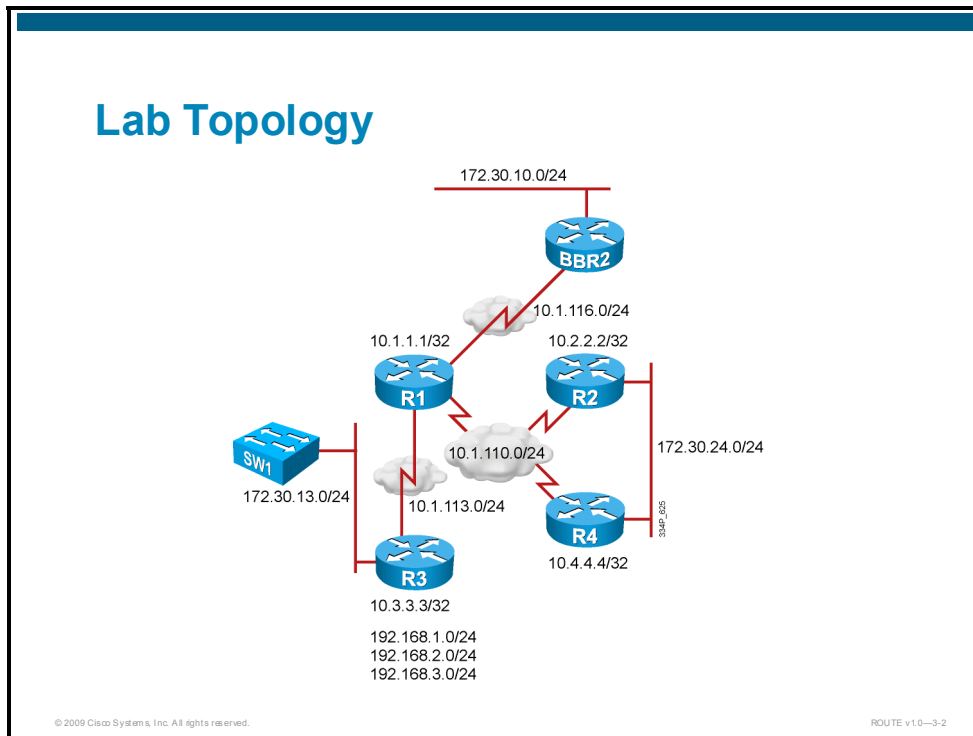
Objectives

Upon completing this lesson, you will be able to configure and verify OSPF route summarization for interarea and external routes. This ability includes being able to meet these objectives:

- Identify the implementation, verification, and troubleshooting tasks for OSPF route summarization for interarea and external routes
- Present a sample solution and identify possible alternative solutions

Lab Overview and Verification

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



The figure presents the physical the lab topology that is used for Lab 3-3: “Configure and Verify OSPF Route Summarization for Interarea and External Routes.” The topology uses four pod routers and one backbone router.

Based on this topology and the OSPF configuration, several internal and external OSPF routes exist in the IP routing tables in all the routers in different areas. The configuration will be optimized by using summarization.

Lab Review: What Did You Accomplish?

- **Task 1:** Examine OSPF routing information.
 - How can you verify the operation of an OSPF routing protocol?
 - What can you see by observing the OSPF neighbors, OSPF database, OSPF interfaces, and IP routing table?
- **Task 2:** Summarize OSPF internal routes.
 - How are internal routes defined?
 - How is summarization performed for internal routes?
- **Task 3:** Summarize OSPF external routes.
 - Where do you apply the summarization of external routes?
 - How many sources of external routes exist in the lab?

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-3-3

In the first task, you examined the operation of an OSPF routing protocol by checking the information from the OSPF neighbors, OSPF database, and OSPF interfaces. You also examined the IP routing table. Your pod was preconfigured with the OSPF configuration, and the proper configuration steps were needed to optimize the size of the OSPF database and IP routing table.

In the second task, summarization of the OSPF internal routes was required. An OSPF can differentiate between internal and external (redistributed) routes. When applying internal route summarization, the summarization command is applied from the direction (area) where the routes are received.

In the third task, OSPF summarization of OSPF external routes was required. Summarization of external routes should be done as close to the source as possible. Preferably, the router that is redistributing the routes should summarize them as well. There is only one source of external routes. Keep in mind that the summarization command is different from the command for the internal routes, because external routes have no concept of areas.

Verification

- Did you have enough information to create an implementation plan?
- Were you able to define the OSPF topology and the content of an IP routing table?
- Does the summary route exist on the routers in Area 3 and Area 24?
- Does the OSPF link-state database contain relevant information?
- Is the routing information for the external networks 192.168.x.0/24 presented as summary information in the OSPF link-state database and, therefore, in the IP routing tables on relevant routers?
- Do all the routers have IP connectivity to the pod external networks?

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-3-4

To verify the configuration of OSPF, answer the following questions:

- Did you have enough information to create an implementation plan?
- Were you able to define the OSPF topology and the content of an IP routing table?
- Does the summary route exist on the routers in Area 3 and Area 24?
- Does the OSPF link-state database contain relevant information?
- Is the routing information for the external networks 192.168.x.0/24 presented as summary information in the OSPF link-state database and, therefore, in the IP routing tables on relevant routers?
- Do all the routers have IP connectivity to the pod external networks?

Checkpoints

- Examine the IP routing information exchanged by routers configured with OSPF.
- Define the internal and external routes.
- Configure the summarization of internal routes.
- Check if the adjacencies are up.
- Check if the IP routing tables and OSPF databases reflect the summarization.
- Configure the summarization of external routes.
- Check if the adjacencies are up.
- Check if the IP routing tables and OSPF databases reflect the summarization.

© 2009 Cisco Systems, Inc. All rights reserved.

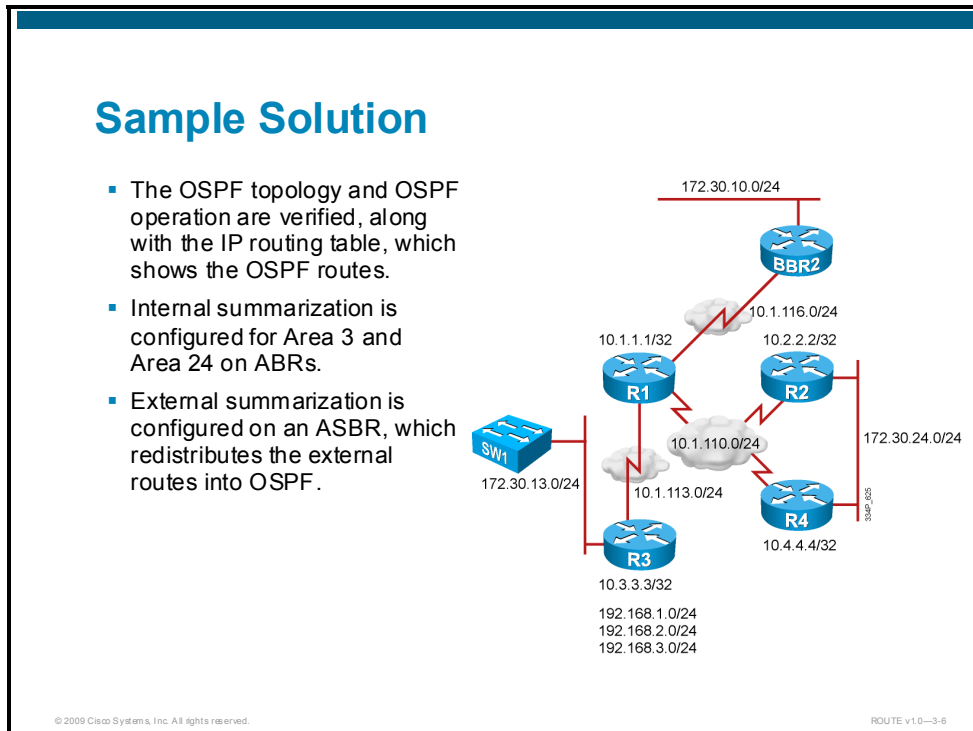
ROUTE v10-3-5

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

- Examine the IP routing information that is exchanged by routers that are configured with OSPF.
- Define internal and external routes.
- Configure the summarization of internal routes.
- Check if the adjacencies are up.
- Check if the IP routing tables and OSPF databases reflect the summarization.
- Configure the summarization of external routes.
- Check if the adjacencies are up.
- Check if the IP routing tables and OSPF databases reflect the summarization.

Sample Solution and Alternatives

This topic describes a sample solution and possible alternatives.




A sample solution includes the implementation details and the details for each task of the implementation plan. Different solutions are possible; the figure shows a few details of a successful configuration.

The proper implementation of route redistribution between multiple IP routing protocols includes the following details:

- The OSPF topology and OSPF operation are verified, along with the IP routing table, which shows the OSPF routes.
- Internal summarization summarizes the internal OSPF routes that are sent into Area 3 and Area 24. ABRs are configured for summarization.
- External summarization is configured on an ASBR, which redistributes the external routes into OSPF.

Alternative Solutions

- Change the area types to optimize the amount of OSPF information entering a specific area.
- Changing the routing protocol is not a realistic solution, but static and default routes can be configured together with route filtering.



© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v10-3-7

OSPF supports different area types, which prevent the insertion of all OSPF routing updates. By selecting different OSPF area types, you can filter network announcements on the ABR or ASBR routers and ensure that only a summary or default route is entered.

To provide reachability in the network, several routing protocols can be used. Changing the routing protocol is not a realistic solution, because network system administrators will not redesign and reconfigure the entire network for another routing protocol. Sometimes, reconfiguration is not even possible, because not all devices may support the desired routing protocol. Instead, static and default routes can be configured together with route filtering. Filtering prevents routing updates from being exchanged, and the static and default routes provide reachability. Static routes are not a scalable solution.

Q and A

1. How can you verify the OSPF routes in an IP routing table?
2. How can you verify the OSPF topology and operation?
3. Where do you perform summarization into Area 3 and Area 24?
4. Where do you perform summarization of external routes?
5. Is there a difference between summarization of internal and external routes?
6. How can you verify summarization results?

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0--3-8

1. You can verify the IP routing table by observing the contents of the IP routing table for the OSPF routes.
2. You can verify the operation of OSPF by checking the information of OSPF neighbors, the OSPF database, OSPF interfaces, and the IP routing table.
3. You should configure summarization of internal routes on the ABRs.
4. You should configure summarization of external routes as close as possible to the source of redistribution, preferably on the ASBRs.
5. You use different configuration commands to configure external and internal summarization.
6. You can verify summarization results by observing the IP routing table and OSPF database.

Summary

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

Summary

- IP routing information that is exchanged by routers configured with OSPF was examined.
- You can optimize the OSPF link-state database and, consequently, the IP routing table by using summarization to make the database smaller.
- OSPF routing operation can be further optimized by summarizing OSPF external routes.

Configuring and Verifying OSPF Special Area Types

Overview

The Open Shortest Path First (OSPF) protocol defines several special-case area types, including stub areas, totally stubby areas, and not-so-stubby areas (NSSAs).

The purpose of all three types of stub areas is to inject default routes into an area so that external and summary link-state advertisements (LSAs) are not flooded. Stub areas are designed to reduce the amount of flooding, the link-state database (LSDB) size, and the routing table size in routers within the area.

Network designers should always consider using stub area techniques when building networks. Stub area techniques improve performance in OSPF networks and allow the network to scale to significantly larger sizes. This lesson discusses OSPF area types and how to configure them.

Objectives

Upon completing this lesson, you will be able to implement and verify different OSPF area types, including the stub area, totally stubby area, NSSA, and totally NSSA. This ability includes being able to meet these objectives:

- Determine OSPF area types
- Define and implement a stub area
- Define and implement a totally stubby area
- Interpret routing tables for OSPF area types
- Define and implement NSSA and totally NSSA in OSPF
- Verify all OSPF area types

OSPF Area Types

This topic describes the OSPF area types.

OSPF Area Types and Structure

- OSPF is based on a two-level hierarchical area structure.
- Each area has its own topology database.
- Area types:
 - Backbone area: Connects all other areas
 - Normal area: Contains all internal and external routing information
 - Stub area: Contains internal and area routing information, but not external routing information
 - Totally stubby area: Contains area routing information only; Cisco proprietary
 - NSSA: Contains area and external routing information

© 2009 Cisco Systems, Inc. All rights reserved.

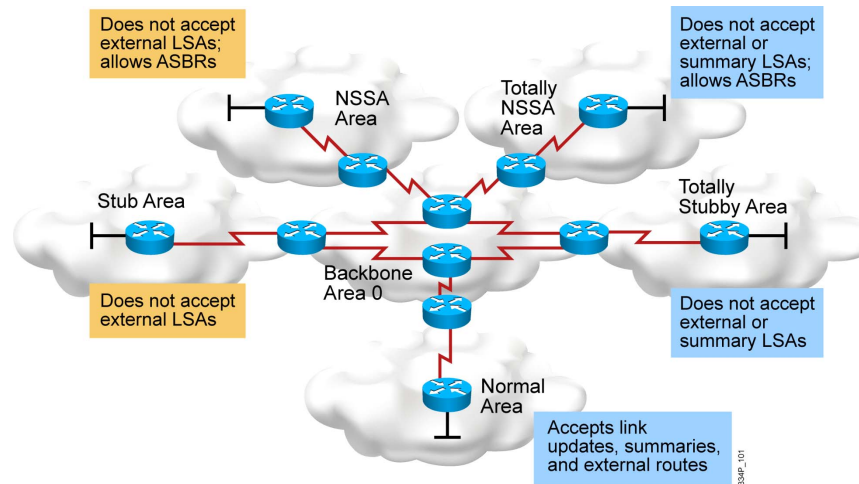
ROUTE v1.0—9-2

OSPF is based on a two-level hierarchical area structure. The hierarchy defines backbone and nonbackbone areas. Each area has its own topology, which is invisible from outside the area. A router that belongs to several areas (called an Area Border Router, or ABR) has several topology databases. All areas have to be connected to a backbone area or linked to it with a virtual link. The backbone area has to be contiguous. A nonbackbone area can be discontinuous.

Area types are as follows:

- **Backbone area:** Connects all other areas
- **Normal area:** Contains all internal and external routing information
- **Stub area:** Contains internal and area routing information, but not external routing information
- **Totally stubby area:** Contains area routing information only; Cisco proprietary
- **NSSA:** Contains area and external routing information

Types of Areas

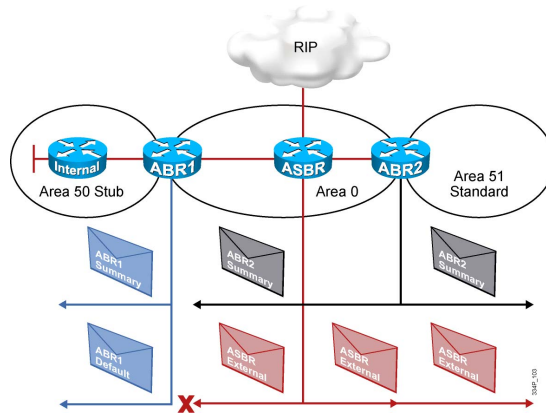


The characteristics that are assigned to an area control the type of route information that it receives. The possible area types are as follows:

- **Normal area:** This default area accepts link updates, route summaries, and external routes.
- **Backbone area (transit area):** The backbone area is the central entity to which all other areas connect. The backbone area is labeled Area 0. All other areas connect to this area to exchange and route information. The OSPF backbone includes all the properties of a standard OSPF area.
- **Stub area:** This area does not accept information about routes that are external to the autonomous system (AS), such as routes from non-OSPF sources. If routers need to route to networks outside the AS, they use a default route, which is noted as 0.0.0.0. Stub areas cannot contain Autonomous System Boundary Routers (ASBRs), except that ABRs may also be ASBRs.
- **Totally stubby area:** This area does not accept external AS routes or summary routes from other areas that are internal to the AS. If the router needs to send a packet to a network that is external to the area, it sends the packet using a default route. Totally stubby areas cannot contain ASBRs, except that ABRs may also be ASBRs.
- **NSSA:** NSSA is an addendum to the OSPF RFC. This area defines a special type 7 LSA. An NSSA offers benefits that are similar to those of a stub area or totally stubby area. However, NSSAs allow ASBRs, which are prohibited in a stub area.

OSPF Router and LSA Types

- ABR is generating summary LSAs.
- ASBR is generating external LSAs.
- Summary and external LSAs can be blocked; default route is sent instead.



The OSPF defines three main router types:

- **Internal router:** This type of router resides inside any area and is a member of one area only.
- **ABR:** This type of router is a member of more than one area. As a border router, it has the ability to control routing traffic from one area to another. Different types of LSAs are exchanged between areas. ABRs can transmit these LSAs, or block them and send default routes instead.
- **ASBR:** This type of router is used to insert external routing information from another non-OSPF AS. ASBRs generate external LSAs, which can be blocked by ABRs.

OSPF routers are able to generate and send the following six types of LSAs: Router link (type 1 LSA), network link (type 2 LSA), network summary (type 3 LSA), ASBR summary (type 4 LSA), external (type 5 LSA), and NSSA external (type 7 LSA).

Stub and Totally Stubby Area Rules

An area can be stub or totally stubby if:

- There is one ABR or more.
- All routers that are members of the stub area are configured as stub routers.
- There is no ASBR in the area.
- The area is not an Area 0.
- No virtual links go through the area.

© 2009 Cisco Systems, Inc. All rights reserved.

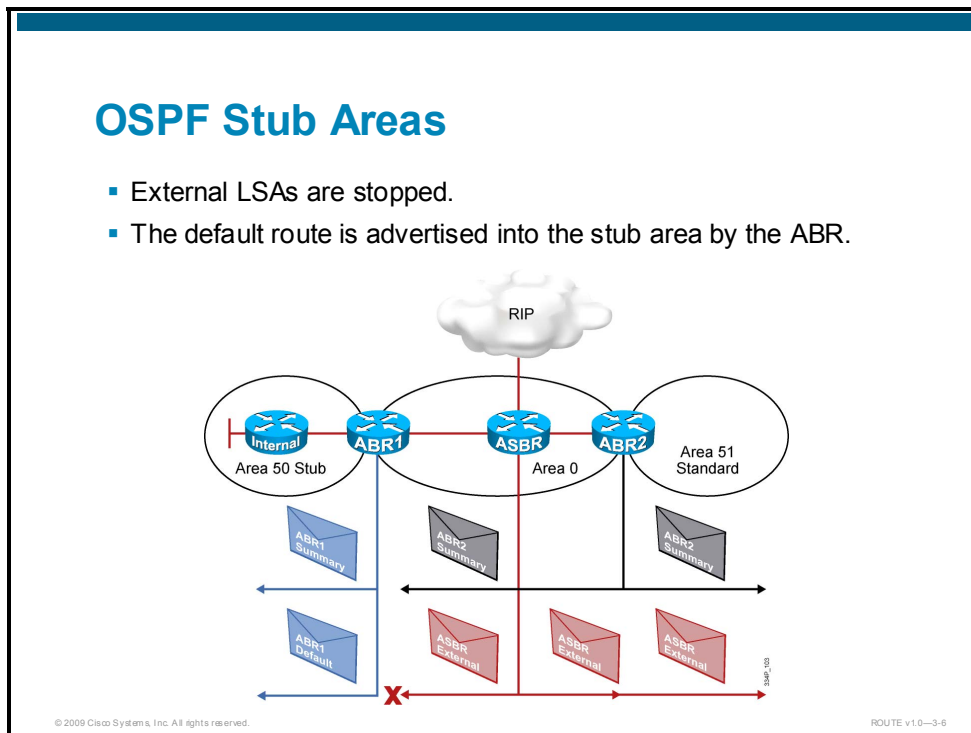
ROUTE v10-3-5

Stub and totally stubby areas do not carry any external routes (type 5 LSAs). An area can be qualified as a stub or totally stubby area if it has the following characteristics:

- There is either a single exit point from that area, or, if there are multiple exits, one or more ABRs inject a default into the stub area and suboptimal routing paths are acceptable. Routing to other areas or autonomous systems could take a suboptimal path to reach the destination by exiting the area at a point that is farther from the destination than other exit points.
- All OSPF routers inside the stub area, including ABRs and internal routers, must be configured as stub routers before they can become neighbors and exchange routing information.
- There is no ASBR inside the area.
- The area is not the backbone area, Area 0.
- The area is not needed as a transit area for virtual links. Recall that a virtual link is a link that allows an area to connect to the backbone via a transit area. Virtual links are generally used for temporary connections or backup after a failure; they should not be considered as part of the primary OSPF design.

Defining and Implementing a Stub Area

This topic defines OSPF stub areas and describes how to configure them.



Configuring a stub area reduces the size of the LSDB inside the area, resulting in reduced memory requirements for routers in that area. External network LSAs (type 5), such as those that are redistributed from other routing protocols into OSPF, are not permitted to flood into a stub area.

Routing from these areas to the outside is based on a default route (0.0.0.0). If a packet is addressed to a network that is not in the routing table of an internal router, the router automatically forwards the packet to the ABR, which sends a 0.0.0.0 LSA. Forwarding the packet to the ABR allows routers within the stub to reduce the size of their routing tables, because a single default route replaces many external routes.

A stub area is typically created when a hub-and-spoke topology is used, with each spoke being a stub area, such as a branch office. In this case, the branch office does not need to know about every network at the headquarters site, because it can use a default route to reach the networks.

Configuring Stub Area

R2 (config-router) #


```
area 2 stub
```

- This command turns on stub area networking.
- Configure all routers in the stub area as stub routers.

ABR (config-router) #

```
area 2 stub  
area 2 default-cost 10
```

- This command defines the cost of a default route sent into the stub area (default is 1); defining the cost is optional.



Once you have gathered all the required information for an implementation plan, you must complete the following tasks to configure OSPF stub areas:

- Configure all routers in the stub area as stub routers.
- Configure the cost for the default route on an ABR (optional).

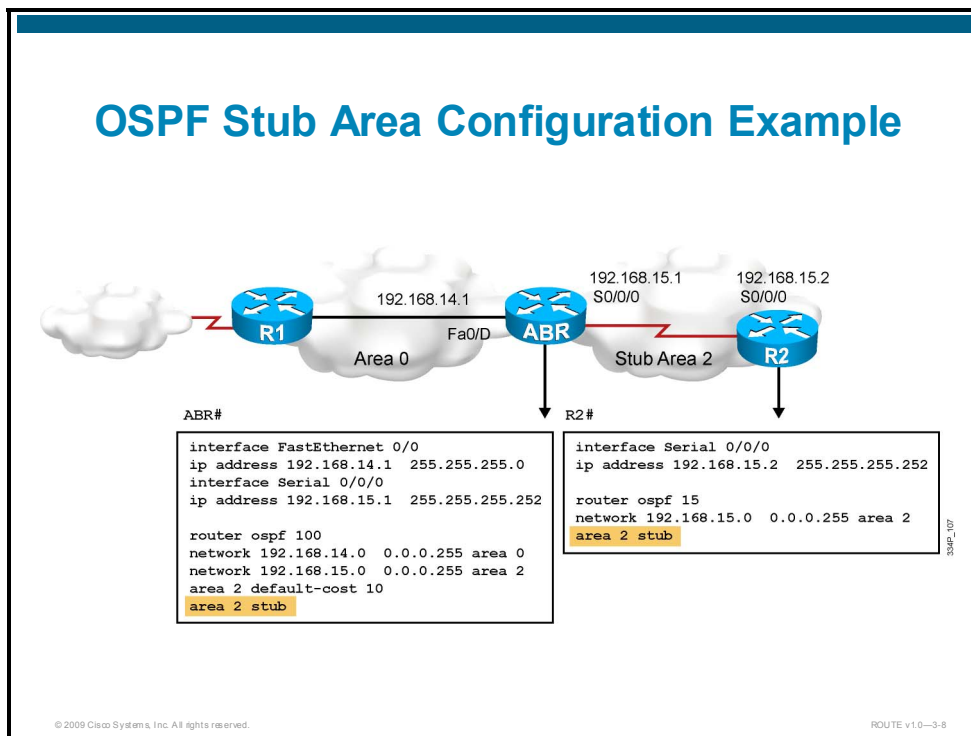
To configure an area as a stub, all routers inside the area must be configured as stub routers. The **area stub** router configuration mode command is used to define an area as a stub area.

By default, the ABR will advertise a default route with a cost of 1. You can change the cost of the default route by using the **area default-cost** command. You only use this command on ABRs that are attached to stub areas or NSSAs. The **default-cost** option provides the metric for the summary default route that is generated by the ABR into the stub area.

For more details about the **area stub** and **area default-cost** commands, go to the Cisco IOS IP Routing: OSPF Command Reference via the following link:

http://www.cisco.com/en/US/docs/ios/iproute_ospf/command/reference/iro_book.html

OSPF Stub Area Configuration Example



Area 2 in the figure is defined as the stub area. No routes from the external AS are forwarded into the stub area.

The last line in each configuration (**area 2 stub**) defines the stub area. Router ABR automatically advertises 0.0.0.0 (the default route) with a default cost metric of 1 into the stub area.

Each router in the stub area must be configured with the **area stub** command. The hello packet that is exchanged between OSPF routers contains a stub area flag that must match on neighboring routers. The **area 2 stub** command must be enabled on all routers in the stub so that they all have the stub flag set; they can then become neighbors and exchange routing information.

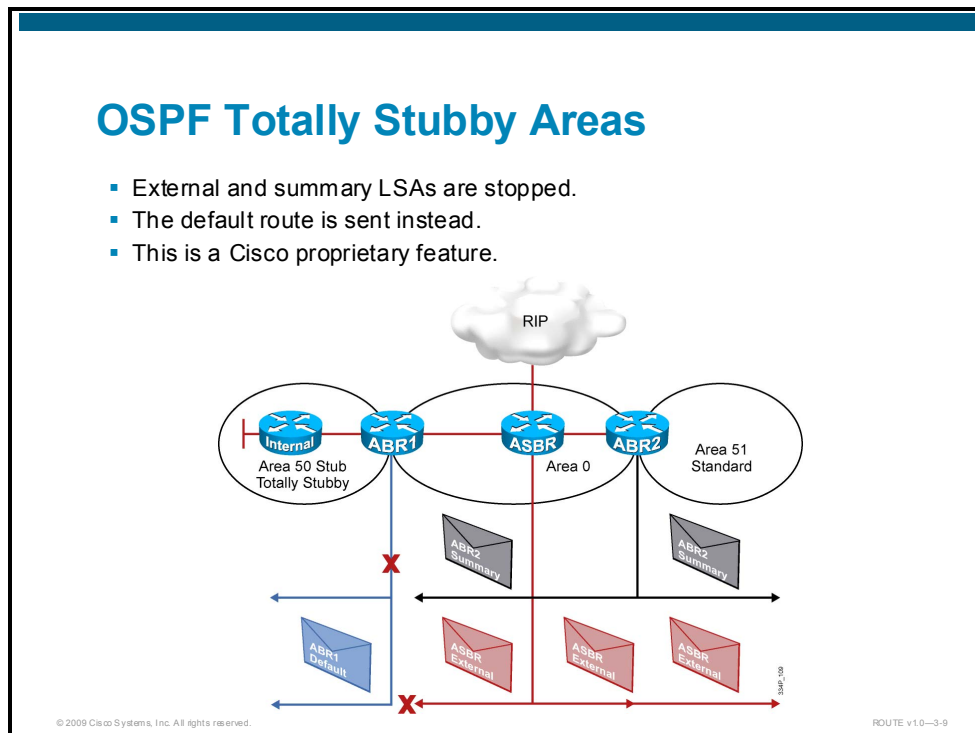
The routes that appear in the routing table of R2 are as follows:

- Intra-area routes, which are designated with an “O” in the routing table.
- The default route and interarea routes, which are both designated with an “IA” in the routing table.
- The default route is also denoted with an asterisk (O *IA).

Configuration of router ABR includes the **area 2 default-cost 10** command, which changes the default value of the cost for default routes from 1 to 10.

Defining and Implementing a Totally Stubby Area

This topic defines OSPF totally stubby areas and describes how to configure them.



The totally stubby area technique is a Cisco proprietary enhancement that further reduces the number of routes in the routing table. A totally stubby area is a stub area that blocks external type 5 LSAs as well as summary type 3 and type 4 LSAs (interarea routes) from entering the area.

Because it blocks these routes, a totally stubby area recognizes only intra-area routes and the default route of 0.0.0.0. ABRs inject the default summary link 0.0.0.0 into the totally stubby area. Each router picks the closest ABR as a gateway to everything outside the area.

Totally stubby areas minimize routing information further than stub areas and increase the stability and scalability of OSPF internetworks. Using totally stubby areas is typically a better solution than using stub areas, as long as the ABR is a Cisco router.

Configuring Totally Stubby Area

R2 (config-router)#


```
area 2 stub
```

- This command turns on stub area networking.
- Configure all routers in the stub area as stub routers.

ABR (config-router) #

```
area 2 stub no-summary  
area 2 default-cost 10
```

- First command defines the totally stubby area on the ABR router.
- Second command defines the cost of a default route sent into the totally stubby area (default is 1); defining the cost is optional.



© 2009 Cisco Systems, Inc. All rights reserved.

RO UTE v1.0-3-10

Once you have gathered all the required information for an implementation plan, you must complete the following tasks to configure OSPF totally stubby areas:

- Configure all routers in stub areas as stub routers.
- Configure the ABR as a totally stubby router.
- Configure the cost of the default route on the ABR, if desired.

To configure an area as totally stubby, you must configure all routers inside the area as stub routers. The **area stub** router configuration mode command defines an area as a stub area. Additionally, the ABR must be configured for totally stubby functionality. Run the **area stub** command with the **no-summary** keyword on the ABR to configure it as totally stubby.

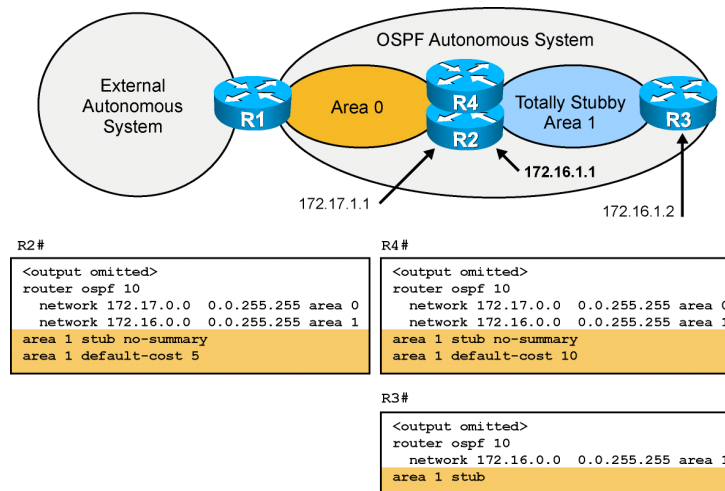
By default, the ABR will advertise a default route with a cost of 1. You can change the cost of the default route by using the **area default-cost** command.

For more details about the **area stub** and **area default-cost** commands, go to the Cisco IOS IP Routing: OSPF Command Reference via the following link:

http://www.cisco.com/en/US/docs/ios/iproute_ospf/command/reference/iro_book.html

Totally Stubby Configuration Example

- R2 is the preferred ABR.



Example: Totally Stubby Configuration

The figure shows an example of a totally stubby area configuration. All routes that are advertised into Area 1 (from Area 0 and the external AS) default to 0.0.0.0. The default route cost is set to 5 on router R2 and to 10 on router R4.

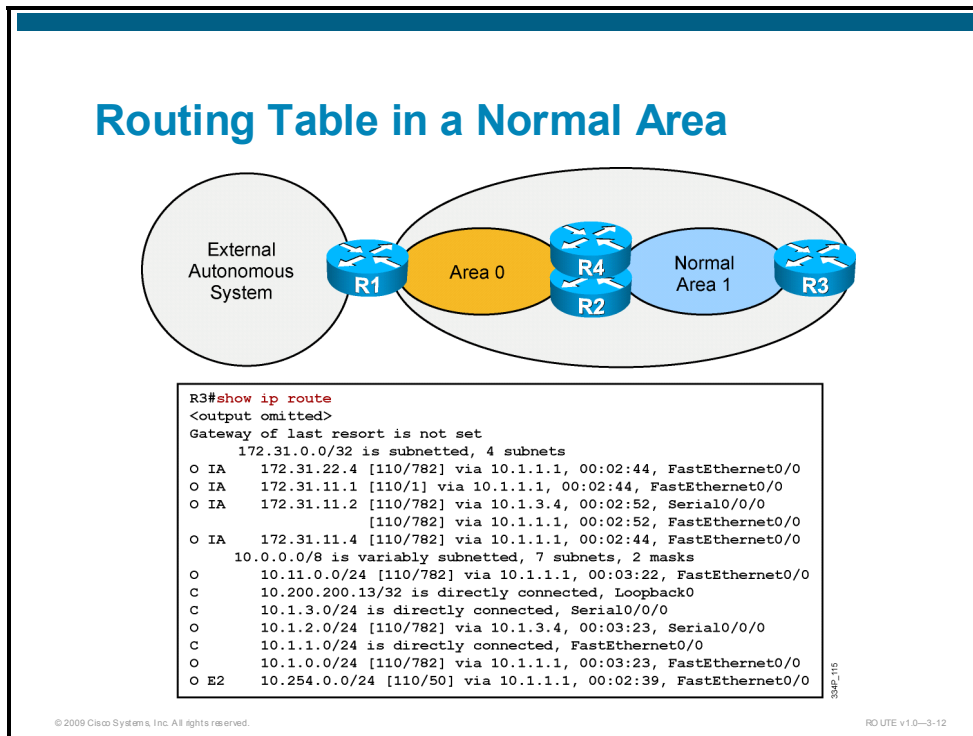
Both default routes are advertised into Area 1. However, the default route from R2 is advertised with a lower cost. This makes routes through R2 preferable even if the internal cost from R3 to R4 is the same as the internal cost from R3 to R2.

Notice that R3 requires the **area 1 stub** command, yet the **no-summary** extension is not required. Only ABRs use the **no-summary** extension to keep summary LSAs from being propagated into the stub area.

Note Remember that all routers in a stub or totally stubby area must be configured as stub routers. An OSPF adjacency will not form between stub and nonstub routers.

Interpreting Routing Tables for OSPF Area Types


This topic interprets information that is shown on routing tables for stub areas and totally stubby areas.



The figure shows how the routing table might look for an OSPF router in a standard area without a stub or totally stubby configuration. Intra-area, interarea, and external routes are all maintained in a standard area.

Routing Table in a Stub Area

- Use the **area 1 stub** command to configure a stub area.



```
R3#show ip route
<output omitted>
Gateway of last resort is 10.1.1.1 to network 0.0.0.0
 172.31.0.0/32 is subnetted, 4 subnets
 O IA 172.31.22.4 [110/782] via 10.1.1.1, 00:01:49, FastEthernet0/0
 O IA 172.31.11.1 [110/1] via 10.1.1.1, 00:01:49, FastEthernet0/0
 O IA 172.31.11.2 [110/782] via 10.1.3.4, 00:01:49, Serial0/0/0
      [110/782] via 10.1.1.1, 00:01:49, FastEthernet0/0
 O IA 172.31.11.4 [110/782] via 10.1.1.1, 00:01:49, FastEthernet0/0
 10.0.0.0/8 is variably subnetted, 6 subnets, 2 masks
 O 10.11.0.0/24 [110/782] via 10.1.1.1, 00:01:50, FastEthernet0/0
 C 10.200.200.13/32 is directly connected, Loopback0
 C 10.1.3.0/24 is directly connected, Serial0/0/0
 O 10.1.2.0/24 [110/782] via 10.1.3.4, 00:01:50, Serial0/0/0
 C 10.1.1.0/24 is directly connected, FastEthernet0/0
 O 10.1.0.0/24 [110/782] via 10.1.1.1, 00:01:50, FastEthernet0/0
 O*IA 0.0.0.0/0 [110/2] via 10.1.1.1, 00:01:51, FastEthernet0/0
```


© 2009 Cisco Systems, Inc. All rights reserved.

RO UTE v1.0—3-13

The figure shows how the same routing table looks if the area is configured as a stub area. Intra-area and interarea routes are all maintained. In a stub area, ABRs block external LSAs, and external routes are not visible in the routing table but accessible via the intra-area default route. The ABR inserts an intra-area default route instead of forwarding external LSAs into the stub area.

Routing Table in a Stub Area with Summarization

- Use the **area 1 stub** and **area 1 range** commands.



```
R3#show ip route
<output omitted>
Gateway of last resort is 10.1.1.1 to network 0.0.0.0
  172.31.0.0/16 is variably subnetted, 2 subnets, 2 masks
  O IA 172.31.22.4/32 [110/782] via 10.1.1.1, 00:13:08, FastEthernet0/0
  O IA 172.31.11.0/24 [110/1] via 10.1.1.1, 00:02:39, FastEthernet0/0
  10.0.0.0/8 is variably subnetted, 6 subnets, 2 masks
  O 10.11.0.0/24 [110/782] via 10.1.1.1, 00:13:08, FastEthernet0/0
  C 10.200.200.13/32 is directly connected, Loopback0
  C 10.1.3.0/24 is directly connected, Serial0/0/0
  O 10.1.2.0/24 [110/782] via 10.1.3.4, 00:13:09, Serial0/0/0
  C 10.1.1.0/24 is directly connected, FastEthernet0/0
  O 10.1.0.0/24 [110/782] via 10.1.1.1, 00:13:09, FastEthernet0/0
  O*IA 0.0.0.0/0 [110/2] via 10.1.1.1, 00:13:09, FastEthernet0/0
```


© 2009 Cisco Systems, Inc. All rights reserved.

RO UTE v1.0-3-14

The figure shows how the same routing table looks if stub area functionality is used and summarization is performed on each ABR. External LSAs are blocked by the ABRs, which also perform summarization. The intra-area and summarized interarea routes are all maintained in the routing tables of stub area routers. External routes are not visible in the routing tables for each stub area, but are accessible via the intra-area default routes for that stub area.

Routing Table in a Totally Stubby Area

- Use the **area 1 stub** command on all internal routers.
- Use the **area 1 stub no-summary** command on the ABRs.



```
R3#show ip route
<output omitted>

Gateway of last resort is 10.1.1.1 to network 0.0.0.0
 10.0.0.0/8 is variably subnetted, 6 subnets, 2 masks
 o 10.11.0.0/24 [110/782] via 10.1.1.1, 00:16:53, FastEthernet0/0
 c 10.200.200.13/32 is directly connected, Loopback0
 c 10.1.3.0/24 is directly connected, Serial0/0/0
 o 10.1.2.0/24 [110/782] via 10.1.3.4, 00:16:53, Serial0/0/0
 c 10.1.1.0/24 is directly connected, FastEthernet0/0
 o 10.1.0.0/24 [110/782] via 10.1.1.1, 00:16:53, FastEthernet0/0
 o*IA 0.0.0.0/0 [110/2] via 10.1.1.1, 00:00:48, FastEthernet0/0
```

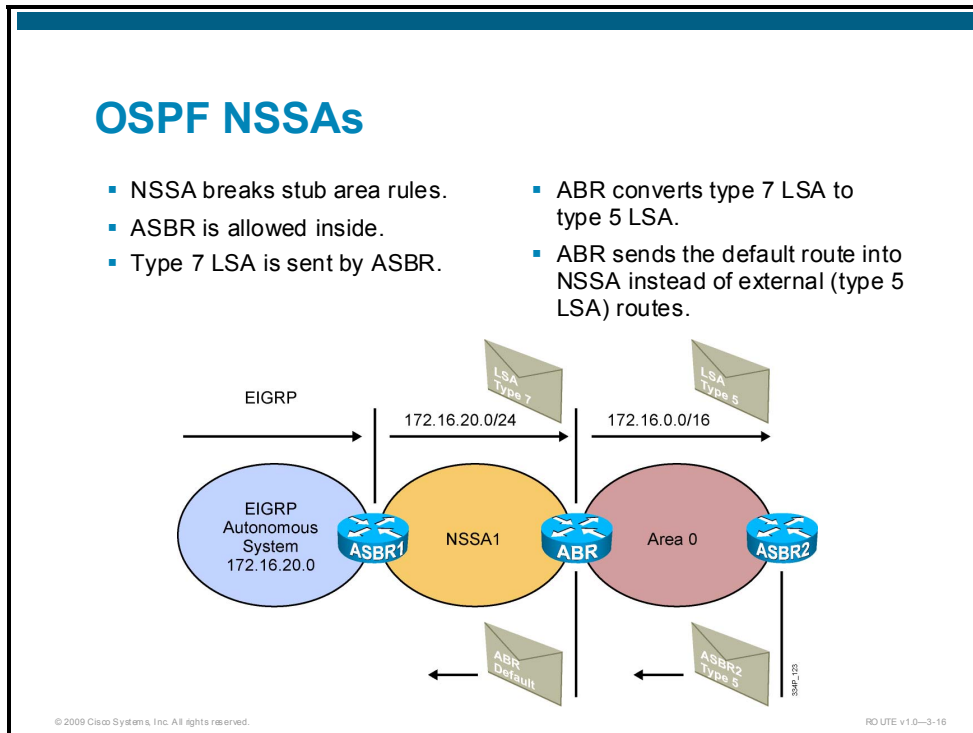
© 2009 Cisco Systems, Inc. All rights reserved.

RO UTE v1.0—3-15

The figure shows how the same routing table looks if the area is configured as a totally stubby area. Notice that routers in the totally stubby area have the smallest routing tables of all the configurations that you have seen so far. Intra-area routes are maintained. Interarea and external routes are not visible in the routing tables for each stub area, but are accessible via the intra-area default routes for that stub area. ABRs block interarea and external LSAs and insert the default routes instead.

Defining and Implementing an NSSA and Totally NSSA in OSPF

This topic defines OSPF NSSAs and describes how to configure them.



The OSPF NSSA feature is described by RFC 3101 and was first introduced in Cisco IOS Release 11.2. It is a nonproprietary extension of the existing stub area feature that allows the injection of external routes in a limited fashion into the stub area.

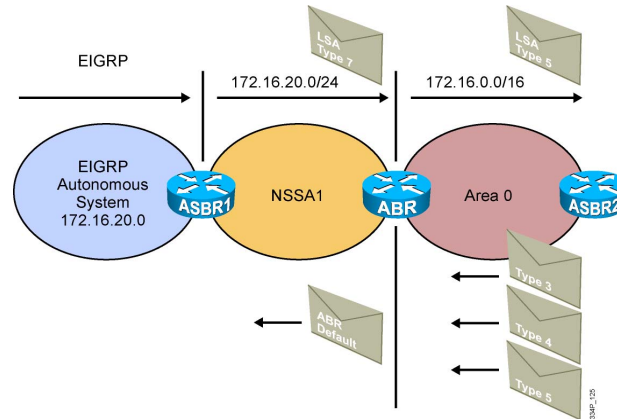
Redistribution into an NSSA creates a special type of LSA known as a type 7 LSA, which can exist only in an NSSA. An NSSA ASBR (router ASBR1 in the figure) generates this LSA, and an NSSA ABR translates it into a type 5 LSA, which gets propagated into the OSPF domain. Type 7 LSAs have a propagate (P) bit in the LSA header to prevent propagation loops between the NSSA and the backbone area. The NSSA retains the other stub area features; the ABR sends a default route into the NSSA instead of external routes from other ASBRs (for example, from ASBR2 in the figure).

The type 7 LSA is described in the routing table as an O N2 or O N1 (N means NSSA). N1 means that the metric is calculated like external type 1 (E1); N2 means that the metric is calculated like external type 2 (E2). The default is O N2.

Note The E1 metric adds external and internal costs together to reflect the whole cost to the destination. The E2 metric takes only the external cost, which is reflected in the OSPF cost.

OSPF Totally NSSA Areas

- ABR is blocking type 3, 4, and 5 LSAs.
- ABR is sending the default route into the NSSA instead.
- This is a Cisco proprietary feature.



The OSPF totally NSSA feature is an extension to the NSSA feature like the totally stubby feature is an extension to the stub area feature. It is a Cisco proprietary feature that blocks type 3, 4, and 5 LSAs. A single default route replaces both inbound-external (type 5) LSAs and summary (type 3 and 4) LSAs in the totally NSSA area.

The ABRs for the totally NSSA area must be configured to prevent the flooding of summary routes for other areas into the NSSA area. Only ABRs control the propagation of type 3 LSAs from the backbone. If an ABR is configured on any other routers in the area, it will have no effect at all.

Configuring Totally NSSA Area

R2 (config-router) #


```
area 2 nssa
```

- This command turns on NSSA area networking.
- Set on all routers in the NSSA area.

ABR (config-router) #

```
area 2 nssa no-summary  
area 2 default-cost 10
```

- The first command defines the totally NSSA area on ABRs.
- The second command defines the cost of a default route sent into the NSSA area (default is 1).



© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-3-18

Once all the required information is defined, you must complete the following tasks to configure OSPF totally NSSAs:

- Configure all routers in the NSSA area as NSSA routers.
- Configure the ABR as totally not-so-stubby (for configuration of the totally NSSA feature only).
- Configure the cost for the default route on the ABR (optional).

Note The default cost of the default route is 1.

To configure an area as an NSSA, you must configure all routers inside the area for NSSA functionality. The **area nssa** router configuration mode command is used to define each router in the NSSA area as not-so-stubby. Each ABR must be configured for NSSA, too.

Note Remember that all routers in the NSSA must be configured with the **area nssa** command in order to form adjacencies.

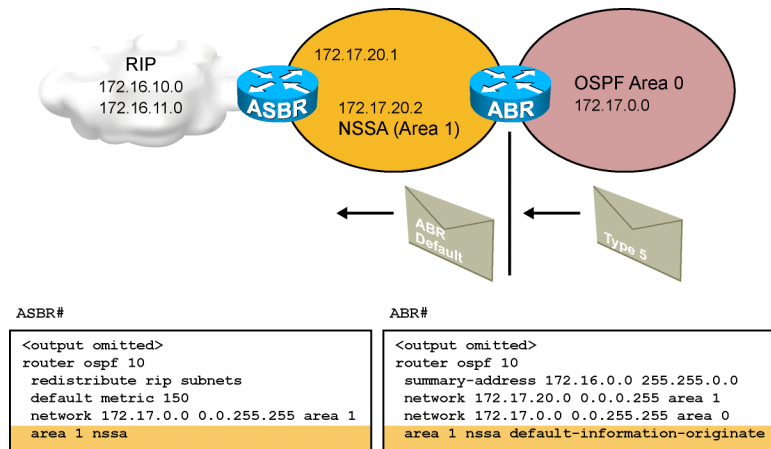
Totally NSSA functionality requires one more step; you must configure each ABR for totally NSSA functionality. The **area nssa** command with the **no-summary** keyword at the end is used to define the ABR as totally not-so-stubby.

By default, the ABR will advertise a default route with a cost of 1. You can change the cost of the default route by using the **area default-cost** command.

For more details about the **area nssa** and **area default-cost** commands, go to the Cisco IOS IP Routing: OSPF Command Reference via the following link:

http://www.cisco.com/en/US/docs/ios/iproute_ospf/command/reference/iro_book.html

NSSA Configuration Example




In the figure, router ASBR is redistributing Routing Information Protocol (RIP) routes into Area 1, the NSSA. Router ABR is the NSSA ABR. This router converts type 7 LSA into type 5 for advertisement into the backbone, Area 0. Router ABR is also configured to summarize the type 5 LSAs that originate from the RIP network; the 172.16.0.0 subnets will be summarized to 172.16.0.0/16 and advertised into Area 0.

To cause router ABR (the NSSA ABR) to generate an O *N2 default route (O *N2 0.0.0.0/0) into the NSSA, use the **area nssa** command with the **default-information-originate** option on R2.

For more details about the **area nssa** command, go to the Cisco IOS IP Routing: OSPF Command Reference via the following link:

http://www.cisco.com/en/US/docs/ios/iproute_ospf/command/reference/iro_book.html

Totally NSSA Configuration Example



ASBR#

```
<output omitted>
router ospf 10
 redistribute rip subnets
 default metric 150
 network 172.17.0.0 0.0.255.255 area 1
 area 1 nssa
```

ABR#

```
<output omitted>
router ospf 10
 summary-address 172.16.0.0 255.255.0.0
 network 172.17.20.0 0.0.0.255 area 1
 network 172.17.0.0 0.0.255.255 area 0
 area 1 nssa no-summary
```

© 2009 Cisco Systems, Inc. All rights reserved.

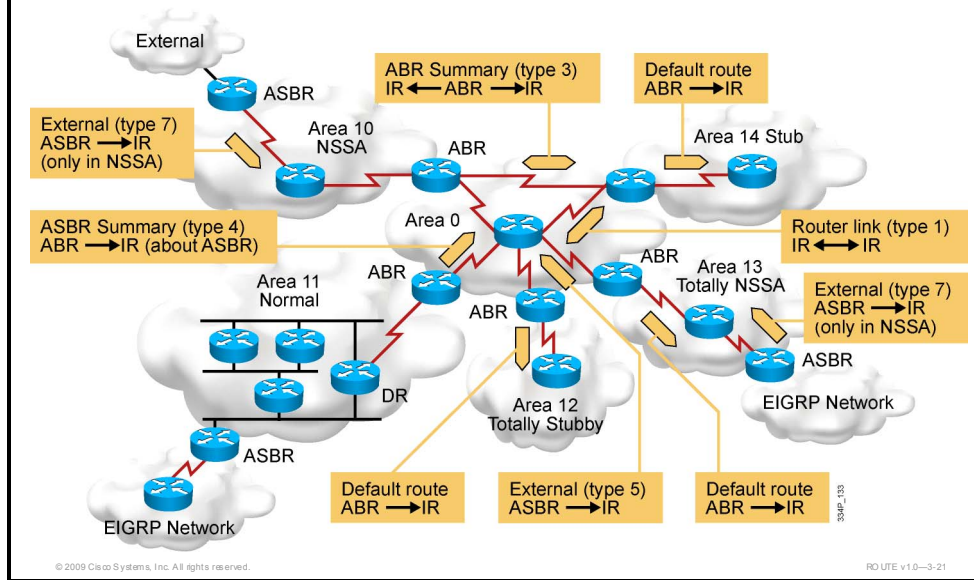
RO UTE v1.0-3-20

In the figure, the ABR is using the **area 1 nssa no-summary** command. This command works the same as the totally stubby technique. A single default route replaces both inbound-external (type 5) LSAs and summary (type 3 and 4) LSAs into the area.

The NSSA ABR, which is router ABR, automatically generates the O *N2 default route into the NSSA area with the **no-summary** option configured at the ABR, so the **default-information-originate** option is not required.

All other routers in the NSSA area only require the **area 1 nssa** command.

Example of Different Areas



The figure shows different area types as follows:

- A normal area accepts link updates, summaries, and external routes.
- A stub area does not accept external LSAs.
- A totally stubby area does not accept summary or external LSAs.
- An NSSA area does not accept external LSAs but allows ASBRs.
- A totally NSSA area does not accept summary or external LSAs but allows ASBRs.

Verifying All OSPF Area Types

This topic describes how to verify all types of OSPF stub areas.

show Commands for Stub and NSSA

R1#
`show ip ospf`

- Displays which areas are normal, stub, or NSSA

R1#
`show ip ospf database`

- Displays the details of the LSAs

R1#
`show ip ospf database nssa-external`

- Displays the details of each type 7 LSA update in the database

R1#
`show ip route`

- Displays all routes

© 2009 Cisco Systems, Inc. All rights reserved. ROUTE v1.0-3-22

The **show** commands in the figure are used to display information about all the types of stub areas that may be configured.

Summary

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

Summary

- There are several OSPF area types: normal, backbone, stub, totally stubby, NSSA, and totally NSSA.
- Use the **area area-id stub** command to define an area as stubby.
- Use the **area area-id stub** command with the **no-summary** keyword on the ABR only to define an area as totally stubby.
- For stub areas, external routes are not visible in the routing table, but are accessible via the intra-area default route. For totally stubby areas, interarea and external routes are not visible in the routing table, but are accessible via the intra-area default route.
- Use the **area area-id nssa** command to define an area as NSSA.
- Use the **show ip ospf**, **show ip ospf database**, and **show ip route** commands to verify all types of stub areas. Use the **show ip ospf database nssa-external** command to display details of type 7 LSAs.

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—3-23

Lab 3-4 Debrief

Overview

In Lab 3-4, you configured and verified OSPF special area types. After performing an initial examination of OSPF, the routing information, and the IP routing table, you discovered that the routers in each OSPF area are announcing that their own subnets and external routes are included in OSPF. To minimize the OSPF routing information in Area 24, you had to convert the area into an OSPF stub area. The results did not meet expectations, so you had to perform additional optimization by converting the area into a totally stub area. At the same time, you also had to reduce OSPF information in Area 3 by converting Area 3 into a totally NSSA.

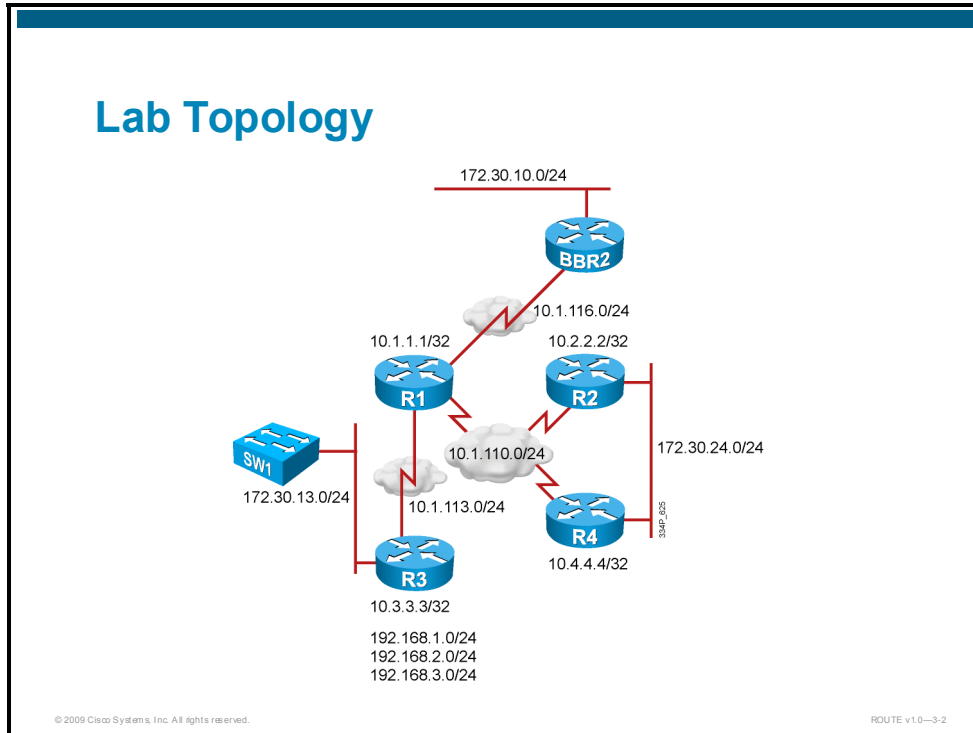
Objectives

Upon completing this lesson, you will be able to configure and verify the OSPF special area types. This ability includes being able to meet these objectives:

- Identify the implementation and verification tasks for implementing the OSPF special area types, including stub, totally stubby, and totally NSSA
- Present a sample solution and identify possible alternative solutions

Lab Overview and Verification

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



The figure presents the physical lab topology that is used for Lab 3-4: “Configure and Verify OSPF Special Area Types.” The topology uses four pod routers and one backbone router.

Based on this topology and the OSPF configuration, several internal and external OSPF routes exist in the IP routing table in all routers in the different areas. Optimization is required to make the IP routing tables smaller. One of the solutions is to change the area types, which will limit the amount of information that is exchanged between the areas and replace many updates with a summary route.

Lab Review: What Did You Accomplish?

- **Task 1:** Examine OSPF routing information.
 - How can you verify the operation of an OSPF routing protocol?
 - What you can see by observing OSPF neighbors, the OSPF database, OSPF interfaces, and the IP routing table?
- **Task 2:** Optimize OSPF routing for Area 24.
 - Which steps are required to restrict OSPF from announcing information about OSPF external routers, while preserving the insertion of internal routes from other areas?
- **Task 3:** Minimize OSPF information in Area 24.
 - Which steps are also required to restrict OSPF from announcing information about OSPF internal routers from other OSPF areas?

Lab Review: What Did You Accomplish? (Cont.)

- **Task 4:** Reduce OSPF information in Area 3.
 - Which steps are required to minimize the sizes of the OSPF link-state database and IP routing table on a router inside the area, in such a way that only information about the area announced routes (internal or external) is allowed and redistribution of external routes is preserved?

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0--3-4

In the first task, you examined OSPF routing information. Your pod was preconfigured with the OSPF configuration. The steps that you needed to follow to optimize OSPF are in the following tasks.

In the second task, you configured the optimization for OSPF routing in Area 24. Optimization prevents OSPF from announcing information about the OSPF external routers while preserving the insertion of the internal routes from other areas. You did this by configuring a stub area.

In the third task, you minimized the OSPF information in Area 24. You needed to optimize the IP routing table by preventing OSPF from announcing information about the OSPF internal routers from other OSPF areas. You did this by configuring a totally stubby area.

In the fourth task, you reduced the size of the OSPF link-state database and IP routing table on the router inside Area 3. You needed to allow only information about the routes that are announced by the area (internal and external). At the same time, you had to preserve the redistribution of external routes. You did this by configuring a totally NSSA.

Verification

- Did you have enough information to create an implementation plan?
- Were you able to define the OSPF topology and content of the IP routing table?
- After the second task, were external routes from other areas suppressed, while internal routes remained in the IP routing table and were injectable?
- After the third task, were external and internal routes from other areas suppressed?
- After the fourth task, were external and internal routes from other areas prevented from being injected, while redistribution of external routes was allowed?

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-3-5

To verify the implementation process for configuring OSPF operations, answer the following questions:

- Did you have enough information to create an implementation plan?
- Were you able to define the OSPF topology and content of the IP routing table?
- After the second task, were external routes from other areas suppressed, while internal routes remained in the IP routing table and were injectable?
- After the third task, were external and internal routes from other areas suppressed?
- After the fourth task, were external and internal routes from other areas prevented from being injected, while redistribution of external routes was allowed?

Checkpoints

- Examine the IP routing information that is exchanged by routers configured with OSPF.
- Change the area type to suppress external routes from other areas to be injected.
- Check the IP routing table and OSPF database for verification.
- Change the area type to suppress external and internal routes from other areas to be injected.
- Check the IP routing table and OSPF database for verification.
- Change the area type to suppress external and internal routes but allow injection of external routes into the area.
- Check the IP routing table and OSPF database for verification.

© 2009 Cisco Systems, Inc. All rights reserved.

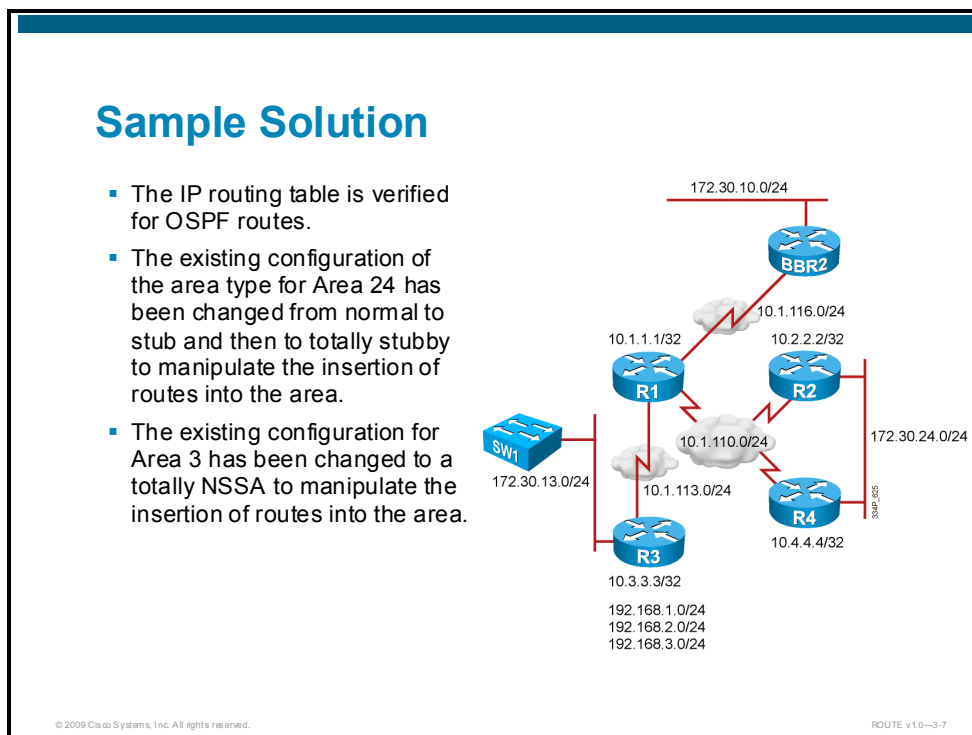
ROUTE v1.0--3-6

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

- Examine the IP routing information that is exchanged by routers that are configured with OSPF.
- Change the area type to suppress external routes from other areas to be injected.
- Check the IP routing table and OSPF database for verification.
- Change the area type to suppress external and internal routes from other areas to be injected.
- Check the IP routing table and OSPF database for verification.
- Change the area type to suppress external and internal routes but allow injection of external routes into the area.
- Check the IP routing table and OSPF database for verification.

Sample Solution and Alternatives

This topic describes a sample solution and possible alternatives.




A sample solution includes the implementation details and the details for each task of the implementation plan. Different solutions are possible; the figure shows a few details of a successful configuration.

The proper implementation of route redistribution between multiple IP routing protocols includes the following details:

- The IP routing table is verified for the OSPF routes.
- The configuration of the area type for Area 24 has been changed from normal to stub and then to totally stubby to manipulate the insertion of routes into the area.
- The configuration for Area 3 has been changed to a totally NSSA to manipulate the insertion of routes into the area.

Alternative Solutions

- Summarization and filtering can be used to manipulate the insertion of routes into a specific area.
- Changing the routing protocol is not a realistic solution, but you can implement static and default routes with filtering instead.



© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-3-8

To decrease the size of the routing table and prevent the insertion of routes, you can use summarization together with the filtering of routes.

To provide reachability in the network, you can use several routing protocols. Changing the routing protocol is not a realistic solution, because network system administrators will not redesign and reconfigure the entire network for another routing protocol. Sometimes, reconfiguration is not even possible, because not all devices may support the desired routing protocol. Instead, static and default routes can be configured together with route filtering. Filtering prevents routing updates from being exchanged, and the static and default routes provide reachability. Static routes are not a scalable solution.

Q and A

1. How you can verify the OSPF routes and OSPF topology?
2. What can you change in OSPF to manipulate which routes are inserted into an area?
3. Which OSPF area type suppresses external routes from other areas to be inserted?
4. Which OSPF area type suppresses external and internal routes from other areas to be inserted?
5. Which area type suppresses external routes from other areas to be inserted but allows redistribution of external routes?
6. Which area type suppresses external and internal routes from other areas to be inserted but allows redistribution of external routes?

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-3-9

1. You can verify the IP routing table by observing the contents of the IP routing table for the OSPF routes. You can verify the operation of OSPF by checking the information of OSPF neighbors, the OSPF database, OSPF interfaces, and the IP routing table.
2. You can use summarization and OSPF area types to manipulate which routes are inserted into the area.
3. The stub OSPF area type suppresses external routes from other areas to be inserted.
4. The totally stub OSPF area type suppresses external and internal routes from other areas to be inserted.
5. The NSSA OSPF area type suppresses external routes from other areas to be inserted, but allows redistribution of external routes.
6. The totally NSSA OSPF area type suppresses external and internal routes from other areas to be inserted, but allows redistribution of external routes.

Summary

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

Summary

- The IP routing table provides information that can be used to verify the proper configuration of different OSPF area types.
- To optimize OSPF, you must configure Area 24 as an OSPF stub area.
- To minimize OSPF information in Area 24, you must configure an OSPF totally stubby area.
- To reduce OSPF information in Area 3, you must configure an OSPF totally NSSA.

Configuring and Verifying OSPF Authentication

Overview

You can prevent your router from receiving fraudulent route updates by configuring neighbor authentication. You can configure Open Shortest Path First (OSPF) neighbor authentication (also called neighbor router authentication or route authentication) in such a way that routers can participate in routing based on predefined passwords.

This lesson describes the two types of OSPF authentication as well as how to configure and verify them.

Objectives

Upon completing this lesson, you will be able to implement authentication in an OSPF network. This ability includes being able to meet these objectives:

- Define types of OSPF authentication
- Implement simple password authentication for OSPF
- Configure MD5 authentication for OSPF
- Troubleshoot authentication for OSPF

Types of OSPF Authentication

This topic describes the two types of OSPF authentication.

OSPF Authentication Types

- OSPF supports two types of authentication:
 - Simple password (or plaintext) authentication
 - MD5 authentication
- The router generates and checks every OSPF packet.
- The source of each routing update packet that is received is authenticated.
- Each participating neighbor must have the same key (password) configured.

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—9-2

OSPF neighbor authentication (also called neighbor router authentication or route authentication) can be configured in such a way that routers participate in routing based on predefined passwords.

Once you have configured neighbor authentication on a router, the router authenticates the source of each routing update packet that it receives. It does this authentication by exchanging an authenticating key (sometimes referred to as a password) that is known to both the sending and the receiving router.

By default, OSPF is not using authentication, which means that routing exchanges over a network are not authenticated. OSPF supports two authentication methods:

- Simple password or plaintext authentication
- Message Digest 5 (MD5) authentication

OSPF MD5 authentication includes a nondecreasing sequence number in each OSPF packet to protect against replay attacks.

Implementing Simple Password Authentication for OSPF

This topic describes how to configure simple password authentication.

Configuring Simple Password Authentication for OSPF

```
R1(config-if)#ip ospf authentication-key mykey
```

- This command defines a password to be used with a neighboring router.
- The neighboring router must have the same password configured.

```
R1(config-if)#ip ospf authentication
OR
R1(config-router)#area 0 authentication
```

- Specifies the authentication type for an interface or the authentication type for an area.

© 2009 Cisco Systems, Inc. All rights reserved. ROUTE v10-3-3

To implement OSPF authentication, you must create an implementation plan. The implementation plan includes two main steps for the configuration of authentication:

- Key (password) configuration
- Authentication-type configuration

The first step to configure OSPF simple password authentication is to configure the key (password). Assign a password that is to be used with the neighboring routers that use OSPF simple password authentication by issuing the **ip ospf authentication-key** interface configuration command, as shown in the figure.

Note In Cisco IOS Release 12.4, the router will give a warning message if you try to configure a password that is longer than eight characters; only the first eight characters will be used. Some earlier Cisco IOS Software releases did not provide this warning.

The password that is created by this command is used as a “key” that is inserted directly into the OSPF header when Cisco IOS Software originates routing protocol packets. A separate password can be assigned to each network on a per-interface basis. All neighboring routers on the same network must have the same password to be able to exchange OSPF information.

Note If the **service password-encryption** command is not used when configuring OSPF authentication, the key will be stored as plaintext in the router configuration. If you configure the **service password-encryption** command, the key will be stored and displayed in an encrypted form. When it is displayed, there will be an *encryption-type* of **7** that is specified before the encrypted key.

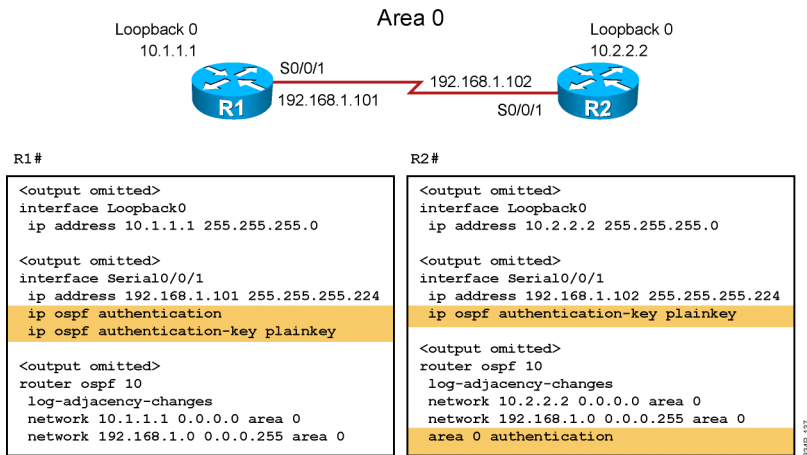
The second step to configure OSPF simple password authentication is to configure the authentication type. Specify the authentication type using the **ip ospf authentication** interface configuration command, as shown in the figure.

The **ip ospf authentication** command was introduced in Cisco IOS Release 12.0. For backward compatibility, the authentication type for an area is still supported. If the authentication type is not specified for an interface, the authentication type for the area will be used (the area default is null authentication). To enable authentication for an OSPF area, use the **area authentication** router configuration command.

For more details about the **ip ospf authentication-key**, **ip ospf authentication**, and **area authentication** commands, go to the Cisco IOS IP Routing: OSPF Command Reference via the following link:

http://www.cisco.com/en/US/docs/ios/iproute_ospf/command/reference/iro_book.html

Simple Password Authentication Configuration Example




The figure shows the network that is used to illustrate the configuration, verification, and troubleshooting of simple password authentication.

The configuration of router R1 is shown in this figure. Simple password authentication is configured on interface Serial0/0/1. The **ip ospf authentication** command is used without the additional keywords. The interface is configured with an authentication key of “plainkey.”

The configuration of router R2 is also shown in this figure. Notice that the connecting interfaces on both R1 and R2 are configured with the same authentication key. The only difference is that R2 authentication is configured for the whole Area 0, while authentication on R1 is configured on the Serial0/0/1 interface only.

Simple Password Authentication Configuration for Virtual Links



```
R1#  
router ospf 10  
network 4.0.0.0 0.255.255.255 area 0  
network 5.0.0.0 0.255.255.255 area 1  
area 0 authentication  
!  
area 1 virtual-link 3.3.3.3 authentication-key cisco
```

```
R3#  
router ospf 10  
network 12.0.0.0 0.255.255.255 area 2  
network 6.0.0.0 0.255.255.255 area 1  
area 0 authentication  
!  
area 1 virtual-link 1.1.1.1 authentication-key cisco
```

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-3-5

The figure shows the network that is used to illustrate the configuration of simple password authentication for virtual links.

The configuration of R1 and R3 is similar. Plaintext authentication is configured for the whole Area 0, and the virtual link to Area 0 is created via transit Area 1 with plaintext authentication enabled. The authentication key “cisco” is configured on both routers.

Verifying Simple Password Authentication

```
R1#show ip ospf interface
Serial2/0 is up, line protocol is up
  Internet Address 192.168.1.101/27, Area 0
  Process ID 10, Router ID 10.1.1.1, Network Type POINT_TO_POINT, Cost: 64
  Transmit Delay is 1 sec, State POINT_TO_POINT
<output omitted>
  Neighbor Count is 1, Adjacent neighbor count is 1
  Adjacent with neighbor 10.2.2.2
  Suppress hello for 0 neighbor(s)
  Simple password authentication enabled
Loopback0 is up, line protocol is up
  Internet Address 10.1.1.1/24, Area 0
  Process ID 10, Router ID 10.1.1.1, Network Type LOOPBACK, Cost: 1
  Loopback interface is treated as a stub Host

R1#ping 10.2.2.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.2.2.2, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/29/32 ms
```

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v10-3-6

The figure shows the output of the **show ip ospf interface** and **ping** commands.

The **show ip ospf interface** command shows that R1 has one adjacent neighbor and that simple password authentication is enabled.

The results of the **ping** command to the R2 loopback interface address are also displayed, to illustrate that the link is working; this proves that devices are accessible.

For more details about the **show ip ospf interface** command, go to the Cisco IOS IP Routing: OSPF Command Reference via the following link:

http://www.cisco.com/en/US/docs/ios/iproute_ospf/command/reference/iro_book.html

For more details about the **ping** command, go to the Cisco IOS Configuration Fundamentals Command Reference via the following link:

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

Configuring MD5 Authentication for OSPF

This topic describes how to configure MD5 authentication.

Configuring OSPF MD5 Authentication

```
R1(config-if)#ip ospf message-digest-key 1 md5 mysecretkey
```


- Defines a key ID and key to be used with a neighboring router.
- Neighboring router must have the same combination of key ID and key configured.

```
R1(config-if)#ip ospf authentication message-digest
```

OR

```
R1(config-router)#area 0 authentication message-digest
```

- Specifies the authentication type for an interface or the authentication type for an area.



© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—3-7

The first step when configuring OSPF MD5 authentication is to configure the key ID and the key (password). You can assign a key ID and key to use with neighboring routers that use OSPF MD5 authentication by entering the **ip ospf message-digest-key** command, as shown in the figure.

Note In Cisco IOS Release 12.4, the router will give a warning message if you try to configure a password that is longer than 16 characters; only the first 16 characters will be used. Some earlier Cisco IOS Software releases did not provide this warning.

The key and the key ID that are specified in this command are used to generate a message digest (also called a hash) of each OSPF packet; the message digest is appended to the packet. A separate password can be assigned to each network on a per-interface basis.

Usually, one key per interface is used to generate authentication information when sending packets and to authenticate incoming packets. All neighboring routers on the same network must have the same password to be able to exchange OSPF information; in other words, the same key ID on the neighbor router must have the same key value.

The key ID allows for uninterrupted transitions between keys, which is helpful for administrators who wish to change the OSPF password without disrupting communication. If an interface is configured with a new key, the router will send multiple copies of the same packet, each authenticated by different keys. The router will stop sending duplicate packets when it detects that all its neighbors have adopted the new key.

For illustrating the process of changing keys, suppose that the current configuration is as follows:

```
interface FastEthernet 0/0
  ip ospf message-digest-key 100 md5 OLD
```

Then, you change the configuration to the following:

```
interface FastEthernet 0/0
  ip ospf message-digest-key 101 md5 NEW
```

The system assumes that its neighbors do not have the new key yet, so it begins a rollover process. It sends multiple copies of the same packet, each authenticated by different keys. In this example, the system sends out two copies of the same packet—the first one authenticated by key 100 and the second one authenticated by key 101.

Rollover allows neighboring routers to continue communication while the network administrator is updating them with the new key. Rollover stops when the local system finds that all its neighbors know the new key. The system detects that a neighbor has the new key when it receives packets from the neighbor that are authenticated by the new key.

After all neighbors have been updated with the new key, the old key should be removed. In this example, you would enter the following:

```
interface FastEthernet 0/0
  no ip ospf message-digest-key 100
```

Then, only key 101 will be used for authentication on the FastEthernet 0/0 interface.

Note You should not keep more than one key per interface. Each time that you add a new key, you should remove the old key to prevent the local system from continuing to communicate with a hostile system that knows the old key. If you do not use the **service password-encryption** command when configuring OSPF authentication, the key will be stored as plaintext in the router configuration. If you use the **service password-encryption** command, the key will be stored and displayed in an encrypted form; when it is displayed, there will be an *encryption-type* of 7 that is specified before the encrypted key.


The second step when configuring OSPF MD5 authentication is to configure the authentication type. Specify the authentication type using the **ip ospf authentication** command, as shown in the figure. For MD5 authentication, enter the **ip ospf authentication** command with the **message-digest** parameter. Before using this command, configure the message digest key for the interface with the **ip ospf message-digest-key** command.

Recall that the **ip ospf authentication** command was introduced in Cisco IOS Release 12.0. As with simple password authentication, the MD5 authentication type for an area is still supported using the **area authentication** router configuration command, for backward compatibility.

For more details about the **ip ospf message-digest-key**, **ip ospf authentication**, and **area authentication** commands, go to the Cisco IOS IP Routing: OSPF Command Reference via the following link:

http://www.cisco.com/en/US/docs/ios/iproute_ospf/command/reference/iro_book.html

OSPF MD5 Authentication Configuration Example



R1#	R2#
<pre><output omitted> interface Loopback0 ip address 10.1.1.1 255.255.255.0 <output omitted> interface Serial0/0/1 ip address 192.168.1.101 255.255.255.224 ip ospf authentication message-digest ip ospf message-digest-key 1 md5 mysecret <output omitted> router ospf 10 log-adjacency-changes network 10.1.1.1 0.0.0.0 area 0 network 192.168.1.0 0.0.0.255 area 0</pre>	<pre><output omitted> interface Loopback0 ip address 10.2.2.2 255.255.255.0 <output omitted> interface Serial0/0/1 ip address 192.168.1.102 255.255.255.224 ip ospf message-digest-key 1 md5 mysecret <output omitted> router ospf 10 log-adjacency-changes network 10.2.2.2 0.0.0.0 area 0 network 192.168.1.0 0.0.0.255 area 0 area 0 authentication message-digest</pre>

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-3-8

The figure shows the network that is used to illustrate the configuration, verification, and troubleshooting of MD5 authentication.

The configuration of R1 is shown in this figure. MD5 authentication is configured on interface Serial 0/0/1. The **ip ospf authentication** command is used with the **message-digest** keyword. The interface is configured with authentication key number 1 set to “mysecret.”

The configuration of R2 is also shown in this figure. Notice that the connecting interfaces on both R1 and R2 are configured with the same authentication key ID and the key. The authentication configuration on R2 is enabled for the whole Area 0; authentication on R1 is configured on the Serial0/0/1 interface only.

Verifying MD5 Authentication

```
R1#show ip ospf interface
Serial2/0 is up, line protocol is up
 Internet Address 192.168.1.101/27, Area 0
 Process ID 10, Router ID 10.1.1.1, Network Type POINT_TO_POINT, Cost: 64
 Transmit Delay is 1 sec, State POINT_TO_POINT
 Neighbor Count is 1, Adjacent neighbor count is 1
  Adjacent with neighbor 10.2.2.2
 Suppress hello for 0 neighbor(s)
 Message digest authentication enabled
 Youngest key id is 1
 Loopback0 is up, line protocol is up
 Internet Address 10.1.1.1/24, Area 0
 Process ID 10, Router ID 10.1.1.1, Network Type LOOPBACK, Cost: 1
 Loopback interface is treated as a stub Host

R1#ping 10.2.2.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.2.2.2, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/28/32 ms
```

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v10-3-9

The figure shows the output of the **show ip ospf interface** and **ping** commands.

The **show ip ospf interface** command shows that R1 has one adjacent neighbor and that message digest authentication is enabled.

The results of a **ping** command to the R2 loopback interface address are also displayed, to illustrate that the link is working; this proves that devices are accessible.

Troubleshooting Authentication for OSPF

This topic describes how to verify and troubleshoot authentication and focuses on simple password authentication.

Authentication Verification


Problems include the following:

- Authentication problems:
 - Authentication is not configured on both sides.
 - A different authentication type is configured on either side.
- Different passwords are configured on either side.

R1#

```
debug ip ospf adj
```

- This command displays the OSPF adjacency-related events.



© 2009 Cisco Systems, Inc. All rights reserved.

PD UTE v1.0-3-10

Authentication for OSPF routing packets prevents your router from receiving fraudulent route updates. Once authentication is configured, all the neighbors must use the same password and authentication type. If the authentication type is not the same for both neighbors, or one neighbor is not configured at all, OSPF will not be able to exchange routing updates successfully. The same applies for the password, which must be the same for both neighboring routers.

After OSPF configuration, the authentication verification, authentication type, and password must be observed. Once the authentication configuration is correct, an OSPF neighbor relationship is established, routing updates are exchanged, and OSPF routes are entered in the IP routing table. If the destinations are not accessible, or there are no neighbors or OSPF routes or both in the IP routing table, you can use debugging to obtain additional information about OSPF authentication problems.

To display information on OSPF-related adjacency events, such as packets being dropped because of authentication problems, use the **debug ip ospf adj** command from privileged EXEC mode.

For more details about the **debug ip ospf adj** command, go to the Cisco IOS Debug Command Reference via the following link:

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

Successful Simple Password Authentication Verification

- Authentication is configured correctly.

```
R1#debug ip ospf adj
OSPF adjacency events debugging is on

<output omitted>
*Feb 17 18:42:01.250: OSPF: 2 Way Communication to 10.2.2.2 on Serial0/0/1,
state 2WAY
*Feb 17 18:42:01.250: OSPF: Send DBD to 10.2.2.2 on Serial0/0/1 seq 0x9B6 opt
0x52 flag 0x7 len 32
*Feb 17 18:42:01.262: OSPF: Rcv DBD from 10.2.2.2 on Serial0/0/1 seq 0x23ED
opt0x52 flag 0x7 len 32 mtu 1500 state EXSTART
*Feb 17 18:42:01.262: OSPF: NBR Negotiation Done. We are the SLAVE
*Feb 17 18:42:01.262: OSPF: Send DBD to 10.2.2.2 on Serial0/0/1 seq 0x23ED opt
0x52 flag 0x2 len 72
<output omitted>

R1#show ip ospf neighbor
Neighbor ID      Pri  State      Dead Time   Address      Interface
10.2.2.2         0    FULL/ -    00:00:34   192.168.1.102 Serial0/0/1
```

The following output of the **debug ip ospf adj** command (part of which is also shown in the figure) illustrates successful simple password authentication on R1 after the Serial0/0/1 interface, on which authentication has been configured, comes up:

```
R1#debug ip ospf adj
OSPF adjacency events debugging is on
*Feb 17 18:41:51.242: OSPF: Interface Serial0/0/1 going Up
*Feb 17 18:41:51.742: OSPF: Build router LSA for area 0,
router ID 10.1.1.1, seq 0x80000013
*Feb 17 18:41:52.242: %LINEPROTO-5-UPDOWN: Line protocol on
Interface Serial0/0/1, changed state to up
*Feb 17 18:42:01.250: OSPF: 2 Way Communication to 10.2.2.2 on
Serial0/0/1, state 2WAY
*Feb 17 18:42:01.250: OSPF: Send DBD to 10.2.2.2 on
Serial0/0/1 seq 0x9B6 opt 0x52 flag 0x7 len 32
*Feb 17 18:42:01.262: OSPF: Rcv DBD from 10.2.2.2 on
Serial0/0/1 seq 0x23ED opt0x52 flag 0x7 len 32 mtu 1500 state
EXSTART
*Feb 17 18:42:01.262: OSPF: NBR Negotiation Done. We are the
SLAVE
*Feb 17 18:42:01.262: OSPF: Send DBD to 10.2.2.2 on
Serial0/0/1 seq 0x23ED opt 0x52 flag 0x2 len 72
*Feb 17 18:42:01.294: OSPF: Rcv DBD from 10.2.2.2 on
Serial0/0/1 seq 0x23EE opt0x52 flag 0x3 len 72 mtu 1500 state
EXCHANGE
*Feb 17 18:42:01.294: OSPF: Send DBD to 10.2.2.2 on
Serial0/0/1 seq 0x23EE opt 0x52 flag 0x0 len 32
*Feb 17 18:42:01.294: OSPF: Database request to 10.2.2.2
*Feb 17 18:42:01.294: OSPF: sent LS REQ packet to
192.168.1.102, length 12
```

```
*Feb 17 18:42:01.314: OSPF: Rcv DBD from 10.2.2.2 on
Serial0/0/1 seq 0x23EF opt0x52 flag 0x1 len 32 mtu 1500 state
EXCHANGE
*Feb 17 18:42:01.314: OSPF: Exchange Done with 10.2.2.2 on
Serial0/0/1
*Feb 17 18:42:01.314: OSPF: Send DBD to 10.2.2.2 on
Serial0/0/1 seq 0x23EF opt 0x52 flag 0x0 len 32
*Feb 17 18:42:01.326: OSPF: Synchronized with 10.2.2.2 on
Serial0/0/1, state FULL
*Feb 17 18:42:01.330: %OSPF-5-ADJCHG: Process 10, Nbr 10.2.2.2
on Serial0/0/1 from LOADING to FULL, Loading Done
*Feb 17 18:42:01.830: OSPF: Build router LSA for area 0,
router ID 10.1.1.1, seq 0x80000014
```

The output of the **show ip ospf neighbor** command in the figure shows that R1 has successfully formed an adjacency with R2.

Troubleshooting Simple Password Authentication Problems

- Simple authentication is not configured on R2.

```
R1#
*Feb 17 18:51:31.242: OSPF: Rcv pkt from 192.168.1.102, Serial0/0/1 :
Mismatch Authentication type. Input packet specified type 0, we use type 1

R2#
*Feb 17 18:50:43.046: OSPF: Rcv pkt from 192.168.1.101, Serial0/0/1 :
Mismatch Authentication type. Input packet specified type 1, we use type 0
```

3349-148

- Different keys on R1 and R2.

```
R1#
*Feb 17 18:54:01.238: OSPF: Rcv pkt from 192.168.1.102, Serial0/0/1 :
Mismatch Authentication Key - Clear Text

R2#
*Feb 17 18:53:13.050: OSPF: Rcv pkt from 192.168.1.101, Serial0/0/1 :
Mismatch Authentication Key - Clear Text
```

3349-151

© 2009 Cisco Systems, Inc. All rights reserved.

RO UTE v1.0-3-12

If simple password authentication is configured on the R1 Serial0/0/1 interface, but no authentication is configured on the R2 Serial0/0/1 interface, the routers will not be able to form an adjacency over that link. The output of the **debug ip ospf adj** command in the upper portion of the figure shows that the routers report a mismatch in authentication type; no OSPF packets will be sent between the neighbors.

Note The different types of authentication have these type codes: null—type 0, simple password—type 1, and MD5—type 2.

If simple password authentication is configured on the Serial0/0/1 interface on R1 and R2, but with different passwords, the routers will not be able to form an adjacency over that link.

The output of the **debug ip ospf adj** command in the lower portion of the figure shows that the routers report an authentication key mismatch; no OSPF packets will be sent between the neighbors.

Successful MD5 Authentication Verification

- Authentication is configured correctly.

```
R1#debug ip ospf adj
OSPF adjacency events debugging is on
<output omitted>
*Feb 17 17:14:06.530: OSPF: Send with youngest Key 1
*Feb 17 17:14:06.546: OSPF: 2 Way Communication to 10.2.2.2 on Serial0/0/1,
state 2WAY
*Feb 17 17:14:06.546: OSPF: Send DBD to 10.2.2.2 on Serial0/0/1 seq 0xB37 opt
0x52 flag 0x7 len 32
*Feb 17 17:14:06.546: OSPF: Send with youngest Key 1
*Feb 17 17:14:06.562: OSPF: Rcv DBD from 10.2.2.2 on Serial0/0/1 seq 0x32F opt
0x52 flag 0x7 len 32 mtu 1500 state EXSTART
*Feb 17 17:14:06.562: OSPF: NBR Negotiation Done. We are the SLAVE
*Feb 17 17:14:06.562: OSPF: Send DBD to 10.2.2.2 on Serial0/0/1 seq 0x32F opt
0x52 flag 0x2 len 72
*Feb 17 17:14:06.562: OSPF: Send with youngest Key 1
<output omitted>

R1#show ip ospf neighbor
Neighbor ID Pri State Dead Time Address Interface
10.2.2.2 0 FULL/- 00:00:35 192.168.1.102 Serial0/0/1
```

The following output of the **debug ip ospf adj** command (part of which is also shown in the figure) illustrates a successful MD5 authentication on R1 after the Serial0/0/1 interface, on which authentication has been configured, comes up:

```
R1#debug ip ospf adj
OSPF adjacency events debugging is on
*Feb 17 17:13:56.530: %LINK-3-UPDOWN: Interface Serial0/0/1,
changed state to up
*Feb 17 17:13:56.530: OSPF: Interface Serial0/0/1 going Up
*Feb 17 17:13:56.530: OSPF: Send with youngest Key 1
*Feb 17 17:13:57.030: OSPF: Build router LSA for area 0,
router ID 10.1.1.1, seq 0x80000009
*Feb 17 17:13:57.530: %LINEPROTO-5-UPDOWN: Line protocol on
Interface Serial0/0/1, changed state to up
*Feb 17 17:14:06.530: OSPF: Send with youngest Key 1
*Feb 17 17:14:06.546: OSPF: 2 Way Communication to 10.2.2.2 on
Serial0/0/1, state 2WAY
*Feb 17 17:14:06.546: OSPF: Send DBD to 10.2.2.2 on
Serial0/0/1 seq 0xB37 opt 0x52 flag 0x7 len 32
*Feb 17 17:14:06.546: OSPF: Send with youngest Key 1
*Feb 17 17:14:06.562: OSPF: Rcv DBD from 10.2.2.2 on
Serial0/0/1 seq 0x32F opt 0
x52 flag 0x7 len 32 mtu 1500 state EXSTART
*Feb 17 17:14:06.562: OSPF: NBR Negotiation Done. We are the
SLAVE
*Feb 17 17:14:06.562: OSPF: Send DBD to 10.2.2.2 on
Serial0/0/1 seq 0x32F opt 0x52 flag 0x2 len 72
*Feb 17 17:14:06.562: OSPF: Send with youngest Key 1
*Feb 17 17:14:06.602: OSPF: Rcv DBD from 10.2.2.2 on
Serial0/0/1 seq 0x330 opt 0x52 flag 0x3 len 72 mtu 1500 state
EXCHANGE
```

```
*Feb 17 17:14:06.602: OSPF: Send DBD to 10.2.2.2 on
Serial0/0/1 seq 0x330 opt 0x52 flag 0x0 len 32
*Feb 17 17:14:06.602: OSPF: Send with youngest Key 1
*Feb 17 17:14:06.602: OSPF: Database request to 10.2.2.2
*Feb 17 17:14:06.602: OSPF: Send with youngest Key 1
*Feb 17 17:14:06.602: OSPF: sent LS REQ packet to
192.168.1.102, length 12
*Feb 17 17:14:06.614: OSPF: Send with youngest Key 1
*Feb 17 17:14:06.634: OSPF: Rcv DBD from 10.2.2.2 on
Serial0/0/1 seq 0x331 opt 0x52 flag 0x1 len 32 mtu 1500 state
EXCHANGE
*Feb 17 17:14:06.634: OSPF: Exchange Done with 10.2.2.2 on
Serial0/0/1
*Feb 17 17:14:06.634: OSPF: Send DBD to 10.2.2.2 on
Serial0/0/1 seq 0x331 opt 0x52 flag 0x0 len 32
*Feb 17 17:14:06.634: OSPF: Send with youngest Key 1
*Feb 17 17:14:06.650: OSPF: Synchronized with 10.2.2.2 on
Serial0/0/1, state FULL
*Feb 17 17:14:06.650: %OSPF-5-ADJCHG: Process 10, Nbr 10.2.2.2
on Serial0/0/1 from LOADING to FULL, Loading Done
*Feb 17 17:14:07.150: OSPF: Send with youngest Key 1
*Feb 17 17:14:07.150: OSPF: Build router LSA for area 0,
router ID 10.1.1.1, seq 0x8000000A
*Feb 17 17:14:09.150: OSPF: Send with youngest Key 1
```

The output of the **show ip ospf neighbor** command in the figure shows that R1 has successfully formed an adjacency with R2.

Troubleshooting MD5 Authentication Problems

- MD5 authentication is configured on both routers.
- R1 has key 1 and R2 has key 2, both with the same passwords:

```
R1#
*Feb 17 17:56:16.530: OSPF: Send with youngest Key 1
*Feb 17 17:56:26.502: OSPF: Rcv pkt from 192.168.1.102, Serial0/0/1 :
Mismatch Authentication Key - No message digest key 2 on interface
*Feb 17 17:56:26.530: OSPF: Send with youngest Key 1

R2#
*Feb 17 17:55:28.226: OSPF: Send with youngest Key 2
*Feb 17 17:55:28.286: OSPF: Rcv pkt from 192.168.1.101, Serial0/0/1 :
Mismatch Authentication Key - No message digest key 1 on interface
*Feb 17 17:55:38.226: OSPF: Send with youngest Key 2
```

3349_106

© 2009 Cisco Systems, Inc. All rights reserved.

RO UTE v1.0-3-14

If MD5 authentication is configured on the Serial0/0/1 interface on R1 and R2, but R1 has key 1 and R2 has key 2, the routers will not be able to form an adjacency over that link, even though both have the same passwords configured. The output of the **debug ip ospf adj** command in the figure shows that the routers report an authentication key mismatch; no OSPF packets will be sent between the neighbors.

Summary

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

Summary

- When authentication is configured, the router generates and checks every OSPF packet and authenticates the source of each routing update packet that it receives. OSPF supports two types of authentication:
 - **Simple password (or plaintext) authentication:** The router sends an OSPF packet and key.
 - **MD5 authentication:** The router generates a message digest, or hash, of the key, key ID, and message. The message digest is sent with the packet; the key is not sent.
- To configure simple password authentication, use the **ip ospf authentication-key** *password* command and the **ip ospf authentication** command.

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—3-15

Summary (Cont.)

- To configure MD5 authentication, use the **ip ospf message-digest-key** *key-id md5 key* command and the **ip ospf authentication message-digest** command.
- Use the **show ip ospf neighbor**, **show ip route**, **ping**, and **debug ip ospf adj** commands to verify and troubleshoot both types of authentication. With MD5 authentication, the **debug ip ospf adj** command output indicates the key ID that is sent.

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—3-16

Lab 3-5 Debrief

Overview

In Lab 3-5, you configured and verified OSPF authentication.

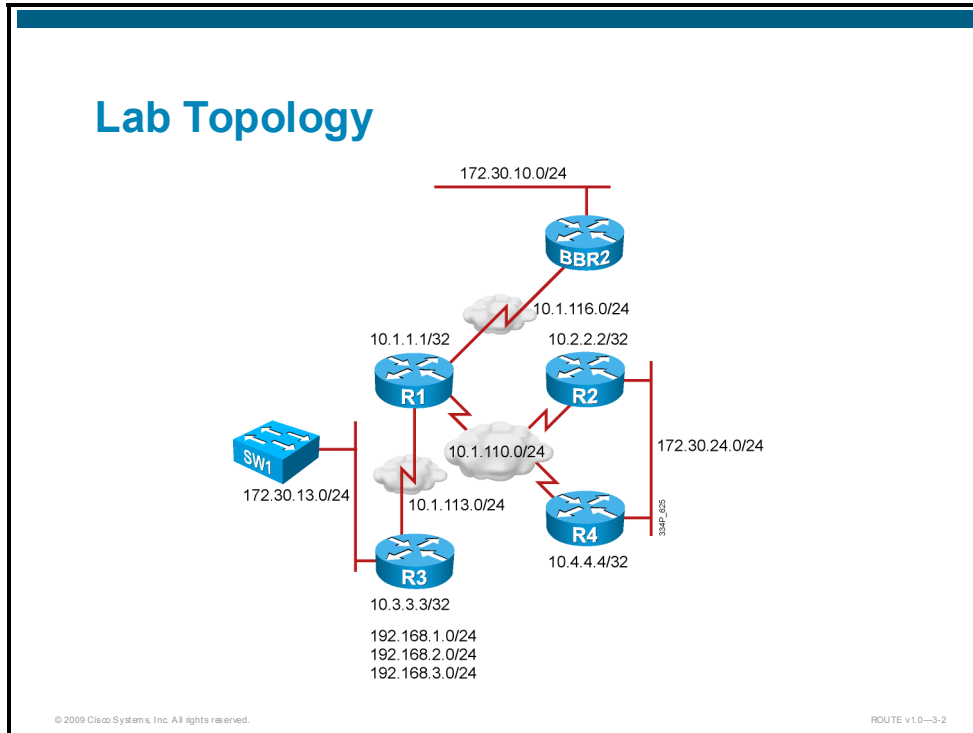
Objectives

Upon completing this lesson, you will be able to configure and verify OSPF authentication. This ability includes being able to meet these objectives:

- Identify the implementation and verification tasks for implementing OSPF authentication
- Present a sample solution and identify possible alternative solutions

Lab Overview and Verification

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



The figure presents the physical lab topology that is used for Lab 3-5: “Configure and Verify OSPF Authentication.” The topology uses four pod routers and one backbone router.

Based on the topology, you need to identify the required parameters and configure OSPF authentication to establish a neighbor relationship and secure exchange of routing packets. The backbone routers are preconfigured with authentication, and the configuration of the pod routers must match. For the rest of the pod routers, the proper configuration must be applied.

Lab Review: What Did You Accomplish?

- **Task 1:** Examine OSPF routing information.
 - How can you verify the operation of an OSPF routing protocol?
 - What can you see by observing the OSPF neighbors, OSPF database, OSPF interfaces, and IP routing table?
- **Task 2:** Enable OSPF link authentication.
 - How is OSPF link authentication between two routers implemented?
- **Task 3:** Enable OSPF area authentication.
 - How is OSPF area authentication implemented on a router?

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0–3-3

In the first task, you examined the OSPF routing information. Your pod was preconfigured with the OSPF configuration, and the proper configuration steps were needed to configure authentication in the following tasks.

In the second task, you needed to enable OSPF link authentication on the LAN segment between routers R2 and R4. You needed to use the authentication that provides the highest level of security.

In the third task, you needed to enable OSPF area authentication in the entire area. Again, you needed to use the authentication that provides the highest level of security.

Verification

- Did you have enough information to create an implementation plan?
- Were you able to define the OSPF topology and the contents of the IP routing table?
- After link authentication configuration, were all the adjacencies, the IP routing table, and the OSPF database included in OSPF?
- What did you see when you examined the OSPF link authentication process?
- After area authentication configuration, were all the adjacencies, the IP routing table, and the OSPF database included in OSPF?
- What did you see when you examined the OSPF area authentication process?
- When a router with an improper or missing OSPF authentication is connected to the area, is it prevented from establishing an OSPF adjacency?

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—3-4

To verify the implementation process for configuring OSPF authentication, answer the following questions:

- Did you have enough information to create an implementation plan?
- Were you able to define the OSPF topology and the contents of the IP routing table?
- After link authentication configuration, were all the adjacencies, the IP routing table, and the OSPF database included in OSPF?
- What did you see when you examined the OSPF link authentication process?
- After area authentication configuration, were all the adjacencies, the IP routing table, and the OSPF database included in OSPF?
- What did you see when you examined the OSPF area authentication process?
- When a router with an improper or missing OSPF authentication is connected to the area, is it prevented from establishing an OSPF adjacency?

Checkpoints

- Examine the IP routing information that is exchanged by routers configured with OSPF.
- Configure OSPF link authentication on the LAN segment.
- Check the IP routing table and OSPF database to verify that the adjacencies on the links are up and the links are included in the OSPF process.
- Configure OSPF area authentication.
- Check the IP routing table and OSPF database to verify that the adjacencies in the area are up and the links are included in the OSPF process.
- Check for OSPF information in the router where area authentication is not configured.

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-3-5

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


- Examine the IP routing information that is exchanged by routers that are configured with OSPF.
- Configure OSPF link authentication on a LAN segment.
- Check the IP routing table and OSPF database to verify that the adjacencies on the links are up and the links are included in the OSPF process.
- Configure OSPF area authentication.
- Check the IP routing table and OSPF database to verify that the adjacencies in the area are up and the links are included in the OSPF process.
- Check for OSPF information in the router where area authentication is not configured.

Sample Solution and Alternatives

This topic describes a sample solution and possible alternatives.

Sample Solution

- The OSPF topology and OSPF operation are verified, along with the IP routing table, which shows the OSPF routes.
- OSPF link authentication is configured on the LAN segment between R2 and R4.
- OSPF area authentication is configured in the entire area.
- MD5 authentication is used as the highest security provided in OSPF.



© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—3-6


A sample solution includes the implementation details and the details for each task of the implementation plan. Different solutions are possible; the figure shows a few details of a successful configuration.

The proper implementation of route redistribution between multiple IP routing protocols includes the following details:

- The OSPF topology and OSPF operation are verified, as well as the IP routing table, which shows the OSPF routes.
- OSPF link authentication is configured on the LAN segment between R2 and R4.
- OSPF area authentication is configured in the entire area.
- MD5 authentication is used as the highest security that is provided in OSPF.

Alternative Solutions

- Different authentication types can be used (simple instead of MD5).
- Different areas and redistribution can be used to control the updates, but they do not provide authentication.



To provide reachability and authentication, you can select between MD5 and simple authentication. In both cases, authentication provides a level of security in the network, but MD5 authentication is more secure.

To control the routing updates, you can use several areas and redistribution. The use of different areas, by itself, does not provide authentication, but you can use a different authentication key in each area. Redistribution does not provide authentication either, but it provides a way for you to control routing updates, if that is a concern.

Q and A

1. How can you verify the OSPF routes in an IP routing table?
2. How many authentication types are supported by OSPF?
3. How do you configure OSPF link authentication?
4. How do you configure OSPF area authentication?
5. Is OSPF area authentication applied to all OSPF-enabled interfaces on the router where it is configured?

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—3-8

1. In the IP routing table, all the routes with the letter “O” in front of them represent the OSPF routes.
2. OSPF supports two types of authentication: simple and MD5 authentication.
3. You can configure link authentication on the interface, along with the authentication key that is used on both neighboring routers.
4. You can configure area authentication in OSPF router configuration mode and apply it to the entire area. You need an authentication key on each interface that is facing the OSPF neighbor.
5. Area authentication is enabled for the entire area. The interfaces that are used must also be configured with the authentication key.

Summary

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

Summary

- IP routing information that is exchanged by routers configured with OSPF was examined.
- Simple OSPF authentication is configured between routers on a LAN segment.
- You not only deploy OSPF area authentication on each link, but also in the entire area for which all OSPF interfaces are involved.

Module Summary

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

Module Summary

- OSPF is one of the most commonly used link-state IP routing protocols in networking. It is an open standard and offers quick convergence and the ability to scale large networks.
- OSPF uses five types of routing protocol packets and six common LSAs. The routing protocol packets are hello, database description, link-state request, link-state update, and link-state acknowledgement. LSAs are router link (type 1 LSA), network link (type 2 LSA), network summary (type 3 LSA), ASBR summary (type 4 LSA), external (type 5 LSA), and NSSA external (type 7 LSA).
- OSPF has the following three network types: point-to-point, broadcast, and NBMA. NBMA supports five modes of OSPF operation: NBMA, point-to-multipoint, point-to-multipoint nonbroadcast, broadcast, and point-to-point.

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-3-1

Module Summary (Cont.)

- The configuration of OSPF is a two-step process:
 - Enter the OSPF configuration with the **router ospf** command.
 - Use the **network** command to describe which interfaces will run OSPF in which area.
- Route summarization reduces OSPF LSA flooding and the routing table size, which reduces memory and CPU utilization on routers.
- The OSPF area types that are supported are backbone area, normal area, stub area, totally stubby area, and NSSA. Stub area techniques improve OSPF performance by reducing the amount of LSA flooding.
- OSPF supports two types of authentication:
 - Simple password (or plaintext) authentication
 - MD5 authentication

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-3-2

Open Shortest Path First (OSPF) is one of the most commonly used interior gateway protocols in IP networking. OSPF is a complex, open-standard protocol that is made up of several protocol handshakes, database advertisements, and packet types.

Module Self-Check

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

- Q1) All these tables are maintained by a link-state routing protocol except which one? (Source: Planning Routing Implementations with OSPF as a Scalable Routing Protocol)
- A) routing
 - B) topology
 - C) update
 - D) neighbor
- Q2) The memory that is needed to maintain tables is one disadvantage of link-state protocols. (Source: Planning Routing Implementations with OSPF as a Scalable Routing Protocol)
- A) true
 - B) false
- Q3) Match each table to its function. (Source: Planning Routing Implementations with OSPF as a Scalable Routing Protocol)
- A) routing
 - B) topology
 - C) neighbor
- _____ 1. stores LSAs
- _____ 2. stores adjacencies
- _____ 3. stores best paths
- Q4) Which term refers to the router that connects Area 0 to a nonbackbone area? (Source: Planning Routing Implementations with OSPF as a Scalable Routing Protocol)
- A) Area Boundary Router
 - B) Area Border Router
 - C) Autonomous System Boundary Router
 - D) backbone router
- Q5) What is the recommended guideline for the maximum number of routers per OSPF area? (Source: Planning Routing Implementations with OSPF as a Scalable Routing Protocol)
- A) 50
 - B) 10
 - C) 200
 - D) 500
- Q6) Which OSPF packet helps form neighbor adjacencies? (Source: Planning Routing Implementations with OSPF as a Scalable Routing Protocol)
- A) exchange packet
 - B) hello packet
 - C) neighbor discovery packet
 - D) adjacency packet

- Q7) Which criterion does SPF use to determine the best path? (Source: Planning Routing Implementations with OSPF as a Scalable Routing Protocol)
- A) lowest delay
 - B) highest bandwidth
 - C) lowest total cost of the route
 - D) total bandwidth of the route
- Q8) Which table is populated as a result of SPF calculations? (Source: Planning Routing Implementations with OSPF as a Scalable Routing Protocol)
- A) topology
 - B) routing
 - C) adjacency
 - D) neighbor
- Q9) Cisco recommends no more than _____ area or areas per ABR in addition to Area 0. (Source: Planning Routing Implementations with OSPF as a Scalable Routing Protocol)
- A) one
 - B) two
 - C) four
 - D) eight
- Q10) An area border router maintains _____. (Source: Planning Routing Implementations with OSPF as a Scalable Routing Protocol)
- A) a separate database for each area with which it is connected
 - B) a single database for all areas
 - C) two databases: one for the backbone and one for all other areas
 - D) a separate routing table for each area
- Q11) In a multiarea network, any area can be the backbone area, although it is most often Area 0. (Source: Planning Routing Implementations with OSPF as a Scalable Routing Protocol)
- A) true
 - B) false
- Q12) When an OSPF router receives an LSA, it is installed in the _____. (Source: Planning Routing Implementations with OSPF as a Scalable Routing Protocol)
- A) neighbor table
 - B) topology table
 - C) routing table
 - D) update table
- Q13) An OSPF router receives an LSA and checks the sequence number of the LSA. This number matches the sequence number of an LSA that the receiving router already has. What does the receiving router do with the received LSA? (Source: Planning Routing Implementations with OSPF as a Scalable Routing Protocol)
- A) ignores the LSA
 - B) adds the LSA to the database
 - C) sends the newer LSU to the source router
 - D) floods the LSA to the other routers

- Q14) An OSPF router receives an LSA. The router checks the sequence number of the LSA and finds that this number is *higher* than the sequence number that it already has. Which two tasks does the router perform with the LSA? (Choose two.) (Source: Planning Routing Implementations with OSPF as a Scalable Routing Protocol)
- A) ignores the LSA
 - B) adds the LSA to the database
 - C) sends the newer LSU to the source router
 - D) floods the LSA to the other routers
- Q15) An OSPF router receives an LSA. The router checks the sequence number of the LSA and finds that this number is *lower* than the sequence number that it already has. What does the router do with the LSA? (Source: Planning Routing Implementations with OSPF as a Scalable Routing Protocol)
- A) ignores the LSA
 - B) adds the LSA to the database
 - C) sends the newer LSU to the source router
 - D) floods the LSA to the other routers
- Q16) Each LSA has its own age timer. By default, how long does an LSA wait before requiring an update? (Source: Planning Routing Implementations with OSPF as a Scalable Routing Protocol)
- A) 30 seconds
 - B) 1 minute
 - C) 30 minutes
 - D) 1 hour
- Q17) A network engineer must configure the OSPF routing protocol in the network. The engineer needs all these to implement the OSPF routing protocol except which one? (Source: Planning Routing Implementations with OSPF as a Scalable Routing Protocol)
- A) implementation plan
 - B) verification plan
 - C) the Cisco IOS Router Configuration Guide for the OSPF routing protocol
 - D) documentation of the OSPF implementation
- Q18) During planning for basic OSPF implementation, the network engineer must go through the following steps: define the network requirements, gather the required parameters, define the OSPF areas and routing, configure basic OSPF, verify the OSPF configuration, and complete the documentation. (Source: Planning Routing Implementations with OSPF as a Scalable Routing Protocol)
- A) true
 - B) false
- Q19) What is the IP protocol number for OSPF packets? (Source: How OSPF Packet Processes Work)
- A) 89
 - B) 86
 - C) 20
 - D) 76

- Q20) All these are OSPF packet types except which one? (Source: How OSPF Packet Processes Work)
- A) LSU
 - B) LSR
 - C) DBD
 - D) LSAck
 - E) hello
 - F) query
- Q21) Which multicast address does the OSPF Hello protocol use? (Source: How OSPF Packet Processes Work)
- A) 224.0.0.5
 - B) 224.0.0.6
 - C) 224.0.0.7
 - D) 224.0.0.8
- Q22) The Hello protocol sends periodic messages to ensure that a neighbor relationship is maintained between adjacent routers. (Source: How OSPF Packet Processes Work)
- A) true
 - B) false
- Q23) Place the exchange protocol states in the correct order. (Source: How OSPF Packet Processes Work)
- A) _____
 - B) _____
 - C) _____
 - D) _____
 - E) _____
 - F) _____
 - G) _____
1. two-way
 2. loading
 3. down
 4. full
 5. exchange
 6. init
 7. exstart
- Q24) DBD packets are involved during which two states? (Choose two.) (Source: How OSPF Packet Processes Work)
- A) exstart
 - B) loading
 - C) exchange
 - D) two-way
- Q25) At which interval does OSPF refresh LSAs? (Source: How OSPF Packet Processes Work)
- A) 10 seconds
 - B) 30 seconds
 - C) 30 minutes
 - D) 1 hour

- Q26) All these fields are in an OSPF packet header except which one? (Source: How OSPF Packet Processes Work)
- A) packet length
 - B) router ID
 - C) authentication type
 - D) MaxAge time
- Q27) OSPF does not require a Hello protocol on point-to-point links, because adjacent routers are directly connected. (Source: Improving Routing Performance in a Complex Enterprise Network)
- A) true
 - B) false
- Q28) Three routers are connected to an Ethernet LAN. One is a small router that should not take on the role of DR or BDR. How do you ensure that it never will? (Source: Improving Routing Performance in a Complex Enterprise Network)
- A) Set the interface priority to 100.
 - B) Set the interface priority to 0.
 - C) Leave the interface priority set to 1, and set the priority of the other two routers to 10.
 - D) Use the **no designated-router** command on the Ethernet interface.
- Q29) When the DR fails, the BDR automatically builds new adjacencies, exchanges databases with other routers, and takes over as DR. (Source: Improving Routing Performance in a Complex Enterprise Network)
- A) true
 - B) false
- Q30) What is the default hello interval for NBMA interfaces? (Source: Improving Routing Performance in a Complex Enterprise Network)
- A) 10 seconds
 - B) 30 seconds
 - C) 120 seconds
 - D) 60 seconds
- Q31) An OSPF router automatically builds adjacencies with neighboring routers on an NBMA link. (Source: Improving Routing Performance in a Complex Enterprise Network)
- A) true
 - B) false
- Q32) Which mode of OSPF operation is RFC-compliant? (Source: Improving Routing Performance in a Complex Enterprise Network)
- A) point-to-multipoint nonbroadcast
 - B) point-to-multipoint
 - C) broadcast
 - D) point-to-point

- Q33) Match the OSPF over Frame Relay mode of operation with its description. (Source: Improving Routing Performance in a Complex Enterprise Network)
- A) broadcast
 - B) point-to-multipoint
 - C) nonbroadcast
- _____ 1. does not discover neighbors automatically
- _____ 2. discovers neighbors automatically and requires a DR and BDR election
- _____ 3. is used in partial-mesh topologies, discovers neighbors automatically, and does not require a DR and BDR election
- Q34) Which two OSPF over Frame Relay modes elect a DR? (Choose two.) (Source: Improving Routing Performance in a Complex Enterprise Network)
- A) broadcast
 - B) nonbroadcast
 - C) point-to-multipoint
 - D) point-to-point
- Q35) A point-to-point subinterface provides which two benefits for OSPF over Frame Relay? (Choose two.) (Source: Improving Routing Performance in a Complex Enterprise Network)
- A) works with multiple vendors
 - B) does not require manual configuration of neighbors
 - C) does not require a DR and BDR
 - D) saves on subnets
- Q36) Which two statements are true for an OSPF adjacency that is built over a Layer 2 MPLS VPN backbone? (Choose two.) (Source: Improving Routing Performance in a Complex Enterprise Network)
- A) The DR and BDR are elected.
 - B) The OSPF parameters must be agreed upon with service provider.
 - C) Provider edge routers appear as additional routers in the customer network.
 - D) The OSPF network type is multiaccess broadcast.
- Q37) Which destination IP address does OSPF use when advertising to all SPF routers? (Source: Improving Routing Performance in a Complex Enterprise Network)
- A) 224.0.0.6
 - B) 224.0.0.5
 - C) 255.255.255.255
 - D) the IP address of the output interface
- Q38) What are the three types of networks that are defined by OSPF? (Choose three.) (Source: Improving Routing Performance in a Complex Enterprise Network)
- A) point-to-point
 - B) broadcast
 - C) point-to-multipoint
 - D) point-to-multipoint nonbroadcast
 - E) nonbroadcast multiaccess

- Q39) When an OSPF adjacency is built over a Layer 3 MPLS VPN backbone, customer routers are unaware of the MPLS VPN topology. (Source: Improving Routing Performance in a Complex Enterprise Network)
- A) true
 - B) false
- Q40) The BDR, like the DR, maintains a complete set of adjacencies on a broadcast link. (Source: Improving Routing Performance in a Complex Enterprise Network)
- A) true
 - B) false
- Q41) Which two OSPF over Frame Relay modes require a DR? (Choose two.) (Source: Improving Routing Performance in a Complex Enterprise Network)
- A) point-to-point
 - B) broadcast
 - C) point-to-multipoint
 - D) nonbroadcast
- Q42) Which two statements regarding OSPF nonbroadcast mode are correct? (Choose two.) (Source: Improving Routing Performance in a Complex Enterprise Network)
- A) This mode requires manual **neighbor** commands.
 - B) This mode does not use a DR and BDR.
 - C) This mode uses a DR and BDR.
 - D) This mode requires multiple subnets.
- Q43) When you are using an OSPF **neighbor** command, you must configure it under an interface. (Source: Improving Routing Performance in a Complex Enterprise Network)
- A) true
 - B) false
- Q44) Which OSPF mode requires **neighbor** commands? (Source: Improving Routing Performance in a Complex Enterprise Network)
- A) broadcast
 - B) point-to-point
 - C) point-to-multipoint
 - D) point-to-multipoint nonbroadcast
- Q45) Which two commands are required for a basic OSPF configuration? (Choose two.) (Source: Configuring and Verifying OSPF Routing)
- A) **network ip-address mask area area-id**
 - B) **network ip-address wildcard-mask area area-id**
 - C) **router ospf process-id**
 - D) **ip router ospf**
- Q46) Which OSPF **show** command describes a list of OSPF adjacencies? (Source: Configuring and Verifying OSPF Routing)
- A) **show ip ospf interface**
 - B) **show ip ospf**
 - C) **show ip route**
 - D) **show ip ospf neighbor**

- Q47) All these techniques are used for router ID selection except which one? (Source: Configuring and Verifying OSPF Routing)
- A) highest IP address on an interface
 - B) IP address on a loopback interface
 - C) lowest IP address when multiple loopback interfaces are used
 - D) the **router-id** command
- Q48) When you use the **router-id** command, the router ID immediately changes to the IP address that has been entered. (Source: Configuring and Verifying OSPF Routing)
- A) true
 - B) false
- Q49) Which network statement is used to configure OSPF on an interface with an IP address of 172.16.1.1 in Area 0? (Source: Configuring and Verifying OSPF Routing)
- A) **network 172.16.0.0 0.0.0.255 area 0**
 - B) **network 172.16.1.1 0.0.0.0 area 0**
 - C) **network 172.16.1.1 255.255.255.255 area 0**
 - D) **network 172.16.0.0 0.0.255.255 area 0**
- Q50) Only one OSPF process can run on a Cisco router at one time. (Source: Configuring and Verifying OSPF Routing)
- A) true
 - B) false
- Q51) A router has a Fast Ethernet interface with an IP address of 172.16.45.1, a loopback 0 interface with an IP address of 10.3.3.3, a loopback 1 interface with an IP address of 10.2.2.2, and a **router-id** command with an IP address of 10.1.1.1. Which router ID will be selected? (Source: Configuring and Verifying OSPF Routing)
- A) 172.16.45.1
 - B) 10.3.3.3
 - C) 10.2.2.2
 - D) 10.1.1.1
- Q52) The **show ip ospf neighbor** command shows the FULL state on one of the two neighbors in its table. Which neighbor or neighbors will successfully exchange LSDB information? (Source: Configuring and Verifying OSPF Routing)
- A) a neighbor that is in the FULL state
 - B) a neighbor that is not in FULL state
 - C) any neighbor, regardless of whether or not it is in the FULL state
 - D) no neighbor, regardless of whether or not it is in the FULL state
- Q53) Which two **show** commands can be used to verify the OSPF router ID of a router? (Choose two.) (Source: Configuring and Verifying OSPF Routing)
- A) **show ip ospf interface**
 - B) **show ip ospf neighbor**
 - C) **show ip ospf**
 - D) **show ip route**

- Q54) When you configure a loopback interface, you choose an IP address that is not going to be advertised by OSPF. This loopback address _____. (Source: Configuring and Verifying OSPF Routing)
- A) cannot be a router ID, because it cannot be pinged
 - B) can be the router ID, even though it cannot be pinged
 - C) can be the router ID and can be pinged if a private address is selected
 - D) cannot be the router ID; you should always advertise loopback addresses
- Q55) Which statement describes the process ID that is used in the **router ospf** command? (Source: Configuring and Verifying OSPF Routing)
- A) All OSPF routers in a network must have the same OSPF process ID.
 - B) The OSPF process ID is an internal number; process IDs on different routers do not need to match.
 - C) The OSPF process ID is similar to an AS number.
 - D) There can be only one OSPF process ID in a router configuration.
- Q56) Which three benefits are derived from a multiarea design in OSPF? (Choose three.) (Source: Configuring and Verifying OSPF Routing)
- A) reduced amount LSA flooding
 - B) reduced number SPF calculations
 - C) reduced size of the neighbor table
 - D) reduced size of the routing table
- Q57) List the four link types that a type 1 LSA defines. (Source: Configuring and Verifying OSPF Routing)
- A) _____
 - B) _____
 - C) _____
 - D) _____
- Q58) Match each LSA name with the number that corresponds to its LSA type. (Source: Configuring and Verifying OSPF Routing)
- A) external
 - B) network
 - C) summary
 - D) multicast
 - E) router
 - F) opaque
 - G) NSSA
- _____ 1
 - _____ 2
 - _____ 3 and 4
 - _____ 5
 - _____ 6
 - _____ 7
 - _____ 9 to 11

- Q59) If the OSPF routing table shows an O E1 route, what does this mean? (Source: Configuring and Verifying OSPF Routing)
- A) It is an interarea route that uses the external cost plus the interarea cost.
 - B) It is an interarea route that uses the external cost only.
 - C) It is an external route that uses the external cost only.
 - D) It is an external route that uses the external cost plus the internal cost.
- Q60) Which two LSAs describe intra-area routing information? (Choose two.) (Source: Configuring and Verifying OSPF Routing)
- A) summary
 - B) external 1
 - C) external 2
 - D) router
 - E) network
- Q61) Where will a type 3 LSA be sent? (Source: Configuring and Verifying OSPF Routing)
- A) only within the area that it originates from
 - B) within the area that it originates from plus the backbone area
 - C) within the area that it originates from, and between all other areas
 - D) within the backbone area, and between all other areas
- Q62) An O E1 route sums up the external metric and the interarea metric, while the O E2 route uses the external metric only. The O E1 route is the default for OSPF; the router must be configured to support O E2. (Source: Configuring and Verifying OSPF Routing)
- A) true
 - B) false
- Q63) A network uses Gigabit Ethernet, and you want OSPF to correctly calculate the metric using bandwidth. Which command should you use to ensure that this happens? (Source: Configuring and Verifying OSPF Routing)
- A) **ip ospf cost** on the interface
 - B) **auto-cost reference-bandwidth** under the OSPF routing process
 - C) **bandwidth** under the interface
 - D) **bandwidth** under the OSPF routing process
- Q64) Looking at the routing table, you notice “[110/55].” What does this mean? (Source: Configuring and Verifying OSPF Routing)
- A) The O E1 cost is 110, and the O E2 cost is 55.
 - B) The administrative distance is 110, and the metric is 55.
 - C) The administrative distance is 55, and the metric is 110.
 - D) The total cost of the route is 165.
- Q65) What does it mean if a route in the routing table has an indicator of O? (Source: Configuring and Verifying OSPF Routing)
- A) It is intra-area.
 - B) It is interarea.
 - C) It is external.
 - D) It is a stub.

- Q66) What is the difference between an LSA 3 and an LSA 4? (Source: Configuring and Verifying OSPF Routing)
- A) LSA 3 is a summary LSA, and LSA 4 is E1.
 - B) LSA 3 is E1, and LSA 4 is a summary.
 - C) LSA 3 is a summary for networks, and LSA 4 is a summary for ASBRs.
 - D) LSA 3 is a summary for ASBRs, and LSA 4 is a summary for networks.
- Q67) By default, OSPF assigns a cost of 1 to a bandwidth of _____. (Source: Configuring and Verifying OSPF Routing)
- A) T1
 - B) 1 Gb/s
 - C) 100 Mb/s
 - D) 10 Gb/s
- Q68) The OSPF LSDB shows an LSA with an age of 1799. What does this mean? (Source: Configuring and Verifying OSPF Routing)
- A) The LSA is going to age out in 1 second.
 - B) It has been 1799 minutes since the last update.
 - C) The LSA will be refreshed in 1 second.
 - D) The LSA was just refreshed, and another refresh is coming in 29 minutes and 59 seconds.
- Q69) What are the two reasons why route summarization is important? (Choose two.) (Source: Configuring and Verifying OSPF Route Summarization)
- A) reduces type 1 LSA flooding
 - B) reduces type 3 LSA flooding
 - C) reduces the size of the routing table
 - D) reduces the size of the neighbor table
- Q70) Which two features play a key role in route summarization? (Choose two.) (Source: Configuring and Verifying OSPF Route Summarization)
- A) contiguous IP addressing
 - B) discontinuous IP addressing
 - C) FLSM
 - D) VLSM
- Q71) Which command would you use to summarize routes into Area 0 from the ABR? (Source: Configuring and Verifying OSPF Route Summarization)
- A) **summary-address**
 - B) **area x range**
 - C) **network**
 - D) **area x summary**
- Q72) Which command would you use to summarize routes into OSPF from the ASBR? (Source: Configuring and Verifying OSPF Route Summarization)
- A) **summary-address**
 - B) **area x range**
 - C) **network**
 - D) **area x summary**

- Q73) A default route is identified in the OSPF database as a _____. (Source: Configuring and Verifying OSPF Route Summarization)
- A) type 1 LSA
 - B) type 2 LSA
 - C) type 3 LSA
 - D) type 4 LSA
 - E) type 5 LSA
- Q74) The primary purpose of a default route is to reduce the sizes of the routing table and the LSDB. A default route avoids detailed updating of routes by inserting a single 0.0.0.0 route into the routing table, making this 0.0.0.0 route act as a gateway of last resort. (Source: Configuring and Verifying OSPF Route Summarization)
- A) true
 - B) false
- Q75) When should you use the **always** keyword with the **default-information originate** command? (Source: Configuring and Verifying OSPF Route Summarization)
- A) on by default; configuration not required
 - B) when you want to send summarized routes
 - C) when your default route is always in the routing table
 - D) when you want the default route advertised, even if it is not in the routing table
- Q76) A summary LSA (type 3 LSA) is designed to automatically summarize a network into blocks. (Source: Configuring and Verifying OSPF Route Summarization)
- A) true
 - B) false
- Q77) Route summarization reduces the flooding of which two of the following LSA types? (Choose two.) (Source: Configuring and Verifying OSPF Route Summarization)
- A) router
 - B) network
 - C) summary
 - D) external
 - E) NSSA
- Q78) You are at the ABR of Area 1 and want to classfully summarize network 172.16.32.0 through 172.16.63.0 into Area 0. Write the configuration command that you would use. (Source: Configuring and Verifying OSPF Route Summarization)
-
- Q79) You are at the ASBR between an OSPF Area 0 and an EIGRP network. EIGRP routes are being redistributed into OSPF. Write the correct summarization command to summarize the EIGRP block 172.16.32.0 through 172.16.63.0. (Source: Configuring and Verifying OSPF Route Summarization)
-
- Q80) It is important that you always summarize the routes from Area 0 into other areas. Suboptimal path selection can occur if you do not. (Source: Configuring and Verifying OSPF Route Summarization)
- A) true
 - B) false

- Q81) The **area range** command has an optional **not-advertise** parameter, which is used to prevent advertising of _____. (Source: Configuring and Verifying OSPF Route Summarization)
- A) all summary LSAs into Area 0
 - B) summary LSAs that match the **area range** command
 - C) all external LSAs
 - D) external LSAs that match the **area range** command
- Q82) Generally, a default route is described in the routing table as an _____. (Source: Configuring and Verifying OSPF Route Summarization)
- A) O route
 - B) O IA route
 - C) O *E1 route
 - D) O *E2 route
- Q83) Which command is best to use if you want to establish a default route from a router that has no default route in its routing table? (Source: Configuring and Verifying OSPF Route Summarization)
- A) **ip route 0.0.0.0 0.0.0.0 next hop address**
 - B) **default-information originate**
 - C) **default-information originate always**
 - D) **static route**
- Q84) The **area x range** and **network** commands are similar because both use inverse masks for configuration purposes. (Source: Configuring and Verifying OSPF Route Summarization)
- A) true
 - B) false
- Q85) All these are permitted in a stub area except which one? (Source: Configuring and Verifying OSPF Special Area Types)
- A) an ABR
 - B) an ASBR
 - C) summary routes
 - D) summary LSAs
- Q86) Which type of router advertises the default into a stub area? (Source: Configuring and Verifying OSPF Special Area Types)
- A) ASBR
 - B) backbone router
 - C) ABR
 - D) internal router
- Q87) What is the correct configuration for stub area 10? (Source: Configuring and Verifying OSPF Special Area Types)
- A) **area 10 stub-area**
 - B) **router ospf 10 stub**
 - C) **area 10 stub**
 - D) **area 10 stub no-summary**

- Q88) What is the meaning of the **no-summary** parameter of the **area x stub** command? (Source: Configuring and Verifying OSPF Special Area Types)
- A) There is no route summarization in the stub area.
 - B) No summary LSAs are sent into the stub area.
 - C) No type 5 LSAs are sent into the stub area.
 - D) There are no external LSAs in the stub area.
- Q89) The default route has a cost of 1 from the stub area ABR if no **area default-cost** command is used. (Source: Configuring and Verifying OSPF Special Area Types)
- A) true
 - B) false
- Q90) A disadvantage of NSSA is that it does not have a totally stubby feature, like a normal stub area. (Source: Configuring and Verifying OSPF Special Area Types)
- A) true
 - B) false
- Q91) Which characteristic is not a prerequisite for stub areas? (Source: Configuring and Verifying OSPF Special Area Types)
- A) virtual links not allowed
 - B) ASBRs not allowed
 - C) ABRs not allowed
 - D) one way in and out of the stub area
- Q92) Stub area design will not improve _____. (Source: Configuring and Verifying OSPF Special Area Types)
- A) CPU utilization on routers in the stub
 - B) memory requirements on routers in the stub
 - C) ability to reach outside networks
 - D) LSDB size on routers in the stub
- Q93) A type 7 LSA appears in the routing table as an _____. (Source: Configuring and Verifying OSPF Special Area Types)
- A) O E1 route
 - B) O E2 route
 - C) O N2 route
 - D) O I/A route
- Q94) What is the difference between a stub area and a totally stubby area configuration? (Source: Configuring and Verifying OSPF Special Area Types)
- A) **no-summary** option on the ABR
 - B) **area area-id totally-stubby** command on the internal routers
 - C) **area area-id nssa** command on the internal routers
 - D) **default-cost** command on the ABR
- Q95) A stub area blocks summary LSAs (type 3 and 4 LSAs). (Source: Configuring and Verifying OSPF Special Area Types)
- A) true
 - B) false

- Q96) Where should you configure the **area area-id stub** command when you are configuring a stub area? (Source: Configuring and Verifying OSPF Special Area Types)
- A) on all routers in the area
 - B) on the ABR
 - C) on the ASBR
 - D) on routers that require stub capability within the area
- Q97) In NSSA, the NSSA ABR translates type 7 LSAs into type 5 LSAs. (Source: Configuring and Verifying OSPF Special Area Types)
- A) true
 - B) false
- Q98) The ABR injects a default route into which three types of areas? (Choose three.) (Source: Configuring and Verifying OSPF Special Area Types)
- A) stub
 - B) totally stubby NSSA
 - C) totally stubby
 - D) Area 0
- Q99) Which two types of authentication are used in OSPF? (Choose two.) (Source: Configuring and Verifying OSPF Authentication)
- A) MD5
 - B) encrypted password
 - C) simple password
 - D) MD6
- Q100) When OSPF authentication is configured between two routers, each router has its own unique password. (Source: Configuring and Verifying OSPF Authentication)
- A) true
 - B) false
- Q101) Which three of these are used to generate the message digest when OSPF MD5 authentication is configured? (Choose three.) (Source: Configuring and Verifying OSPF Authentication)
- A) packet
 - B) sequence number
 - C) key ID
 - D) key
 - E) router ID
- Q102) Which command is used to specify that OSPF simple password authentication is to be used? (Source: Configuring and Verifying OSPF Authentication)
- A) **ip ospf authentication simple**
 - B) **ip ospf authentication**
 - C) **ip ospf authentication-key**
 - D) **ip ospf message-digest-key**
 - E) **ip ospf authentication message-digest**

- Q103) Which command is used to specify that OSPF MD5 authentication is to be used?
(Source: Configuring and Verifying OSPF Authentication)
- A) **ip ospf authentication simple**
 - B) **ip ospf authentication**
 - C) **ip ospf authentication-key**
 - D) **ip ospf message-digest-key**
 - E) **ip ospf authentication message-digest**
- Q104) When a new MD5 key is configured on a router for OSPF authentication, the router will use both the old and new keys until the new key is configured on neighboring routers. (Source: Configuring and Verifying OSPF Authentication)
- A) true
 - B) false
- Q105) Which command is used to troubleshoot OSPF authentication? (Source: Configuring and Verifying OSPF Authentication)
- A) **debug ip ospf adj**
 - B) **debug ip ospf adjacency events**
 - C) **debug ip ospf database**
 - D) **debug ip ospf packets**

Module Self-Check Answer Key

- Q1) C
- Q2) A
- Q3) A = 3, B = 1, C = 2
- Q4) B
- Q5) A
- Q6) B
- Q7) C
- Q8) B
- Q9) B
- Q10) A
- Q11) B
- Q12) B
- Q13) A
- Q14) B, D
- Q15) C
- Q16) C
- Q17) C
- Q18) B
- Q19) A
- Q20) F
- Q21) A
- Q22) A
- Q23) A = 3, B = 6, C = 1, D = 7, E = 5, F = 2, G = 4
- Q24) A, C
- Q25) C
- Q26) D
- Q27) B
- Q28) B
- Q29) B
- Q30) B
- Q31) B
- Q32) B
- Q33) A = 2, B = 3, C = 1
- Q34) A, B
- Q35) B, C
- Q36) A, D
- Q37) B
- Q38) A, B, E
- Q39) A
- Q40) A
- Q41) B, D

- Q42) A, C
- Q43) B
- Q44) D
- Q45) B, C
- Q46) D
- Q47) C
- Q48) B
- Q49) B
- Q50) B
- Q51) D
- Q52) A
- Q53) A, C
- Q54) B
- Q55) B
- Q56) A, B, D
- Q57) A = point-to-point
B = transit
C = stub
D = virtual link
- Q58) A = 5, B = 2, C = 3 and 4, D = 6, E = 1, F = 9 and 11, G = 7
- Q59) D
- Q60) D, E
- Q61) D
- Q62) B
- Q63) B
- Q64) B
- Q65) A
- Q66) C
- Q67) C
- Q68) C
- Q69) B, C
- Q70) A, D
- Q71) B
- Q72) A
- Q73) E
- Q74) A
- Q75) D
- Q76) B
- Q77) C, D
- Q78) area 1 range 172.16.0.0 255.255.0.0
- Q79) summary-address 172.16.32.0 255.255.224.0
- Q80) B
- Q81) B
- Q82) D

- Q83) C
- Q84) B
- Q85) B
- Q86) C
- Q87) C
- Q88) B
- Q89) A
- Q90) B
- Q91) C
- Q92) C
- Q93) C
- Q94) A
- Q95) B
- Q96) A
- Q97) A
- Q98) A, B, C
- Q99) A, C
- Q100) B
- Q101) A, C, D
- Q102) B
- Q103) E
- Q104) A
- Q105) A

Implementing an IPv4-Based Redistribution Solution

Overview

This module explains why it is necessary to manipulate routing information. Without manipulation, route redistribution between IP routing domains may result in suboptimal routing. There are also times when routing information would waste bandwidth on a router interface, because the routing information is not needed.

The implementation of redistribution in a multiprotocol network using Cisco IOS Software features according to a given network design and requirements is a primary learning objective of this module. The lessons describe prefix lists, distribute lists, and route maps first, and then describe the controlling of routing traffic with configuration and verification steps.

Module Objectives

Upon completing this module, you will be able to implement a redistribution solution in a multiprotocol network that uses Cisco IOS Software features to control the path selection and a loop-free topology according to a given network design and requirements. This ability includes being able to meet these objectives:

- Identify common network performance issues
- Define how prefix lists work
- Determine how to use a prefix list to control routing updates
- Determine how distribute lists work
- Use distribute lists to control routing updates
- Determine how route maps work
- Determine how to use route maps to control routing updates
- Use route maps to filter routes

Assessing Network Routing Performance and Security Issues

Overview

Routing updates compete with user data for bandwidth and router resources, but routing updates are critical, because they carry the information that routers need to make sound routing decisions. To ensure that the network operates efficiently, you must control and tune routing updates. Information about networks must be sent where it is needed and filtered from where it is not needed. There is no one type of route filter that is appropriate for every situation. Therefore, the more techniques that you have at your disposal, the better your chance of having a smooth, well-run network.

This lesson discusses common network performance issues and how to control the updates that are sent and received by dynamic routing protocols. Prefix lists, distribute lists, and route maps are described as the main mechanisms to filter and control routing update traffic.

Objectives

Upon completing this lesson, you will be able to describe prefix lists, distribute lists, and route maps. You will also be able to explain how to use them to filter and control routing update traffic. This ability includes being able to meet these objectives:

- Determine common network performance issues
- Identify how distribute lists work
- Use distribute lists to control routing updates
- Identify how prefix lists work
- Use a prefix list to control routing updates
- Identify how route maps work

- Use route maps to control routing updates
- Use route maps to filter routes
- Suppress routing updates using the **passive-interface** command

Common Network Performance Issues

This topic describes common network performance and security issues that are caused by update inefficiencies.

Common Factors Affecting Network Performance

- Routing factors that influence CPU utilization include:
 - The size of the routing information update
 - The frequency of the updates
 - The weaknesses in the design
 - The presence of any route maps or filters
 - Incorrectly configured route filters
- Running different protocols in different areas within the same autonomous system
 - The number of routing protocol processes receiving the updates

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-4-2

Routing updates can decrease network performance. If a Layer 3-enabled switch or router receives many updates, they must be processed. The CPU utilization spike during this processing depends on the following factors:


- The size of the routing update
- The frequency of the updates
- The number of routing protocol processes that are receiving the updates
- The presence of any route filters
- The weaknesses in the design

Some factors are also dependent on the routing protocol. The size of routing updates and their frequency is directly proportional to the amount of traffic that must be sent over the links and processed by the CPUs. The same applies to the number of routing processes per router, where more routing processes generate more updates and require more CPU resources.

Several solutions exist, and a good design and filtering may significantly improve network performance. However, the wrong design and incorrectly configured route filters will not improve network performance at all.

Routing Updates

- Qualities of routing updates that influence CPU utilization include:
 - The size of the routing information update
 - The frequency of the updates
 - A bad design



© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—4-3

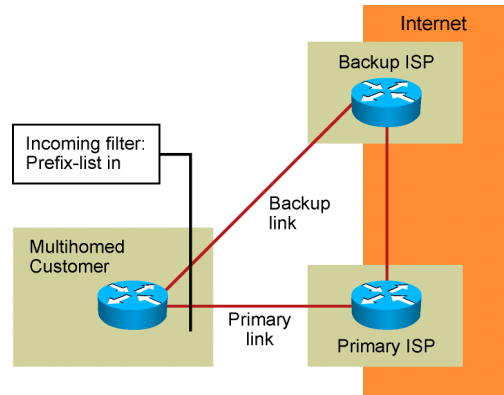
In some network deployments, high CPU utilization is normal. In general, larger Layer 2 or Layer 3 networks demand more CPU resources to process network-related traffic. Examples of operations with potentially high CPU utilization are as follows:

- IP routing table updates
- Cisco IOS Software commands (such as **debug** commands)

More devices and links inside an autonomous system (AS) produce more routing update traffic. Redistribution from an external AS adds to the size as well. When a router receives many routing updates, the routing information must be processed.

Filtering Routing Updates

- Routing update filters may improve network performance.
- Ensure that router filters are configured correctly.




Route filtering works by regulating the routes that are entered into or advertised out of the routing table. Filtering has different effects on link-state routing protocols compared with distance vector protocols. Distance vector protocols advertise routes based on what is in the protocol route table. Link-state protocols on the opposite side determine their routes based on the information in their link-state databases (LSDBs). Distance vector protocols filter routes that the router advertises to its neighbors; link-state protocols, in contrast, do not take into account the advertised route entries of the neighbors of each router.

Filters can be configured to prevent updates through router interfaces, to control the advertising of routes in routing updates, or to control the processing of routing updates. If filters are not configured correctly or if filters are applied to the wrong interfaces, network performance issues will result.

Running Multiple Routing Protocols

- You can run different protocols in different areas within the same AS:
 - If many routing protocol processes receive updates at the same time, performance will be affected.



© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—1-5

You can configure multiple routing protocols in a single router to connect networks that use different routing protocols. For example, you can run Enhanced Interior Gateway Routing Protocol (EIGRP) on one subnetted network and Open Shortest Path First (OSPF) on another subnetted network, and exchange routing information between them in a controlled fashion. Additionally, the same router can be connected to the ISP and exchange Border Gateway Protocol (BGP) routes. Router R1 in the figure runs EIGRP, OSPF, and BGP.

The available routing protocols were not designed to interoperate with one another, so each protocol collects different types of information and reacts to topology changes in its own way. This behavior produces a large amount of routing update traffic that must be processed by each protocol separately in a different way. For example, Routing Information Protocol (RIP) uses a hop-count metric, and OSPF uses link metric information. Maintaining all routing, topology, and database tables inside the routers requires not only a great amount of CPU capacity but also large amounts of memory resources.

Controlling Routing Updates

- Design change:
 - Limit the number of routing protocols used
- Passive interfaces
- Redistribution with route filtering:
 - Access lists
 - Prefix lists
 - Distribute lists
 - Route maps

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-4-6

To increase the network performance, routing updates must be exchanged in a controlled way, and the network designer must take into account the recommended design rules.

The passive interface technique prevents all routing updates from being advertised out of an interface. However, in many cases, you do not want to prevent all routing information from being advertised.

You might want to block the advertisement of only certain specific routes. For example, you might do so to prevent routing loops when you are implementing two-way route redistribution with dual redistribution points.

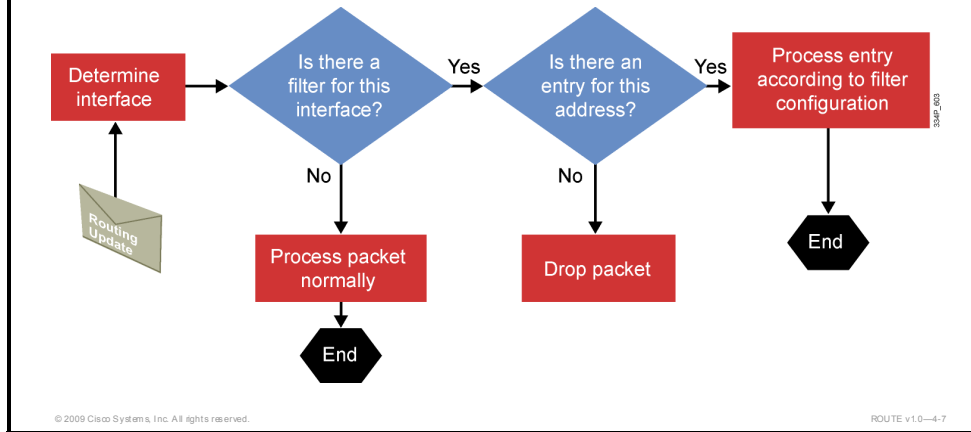
When routing information is being exchanged between different networks that use different routing protocols, you can use many configuration options to filter the routing information. A good design must limit the number of routing protocols that are running on one router, because filtering must be applied to the redistribution between different protocols. You can use various methods of filtering routes, with various effects, to increase network performance.

You can use the following tools to control routing updates and redistribution:

- Access lists
- Prefix lists
- Distribute lists
- Route maps

Using Route Filters

- A neighbor relationship is established.
- Adjacent routes exchange routing updates:
 - The process takes effect after multiple stages have been completed.

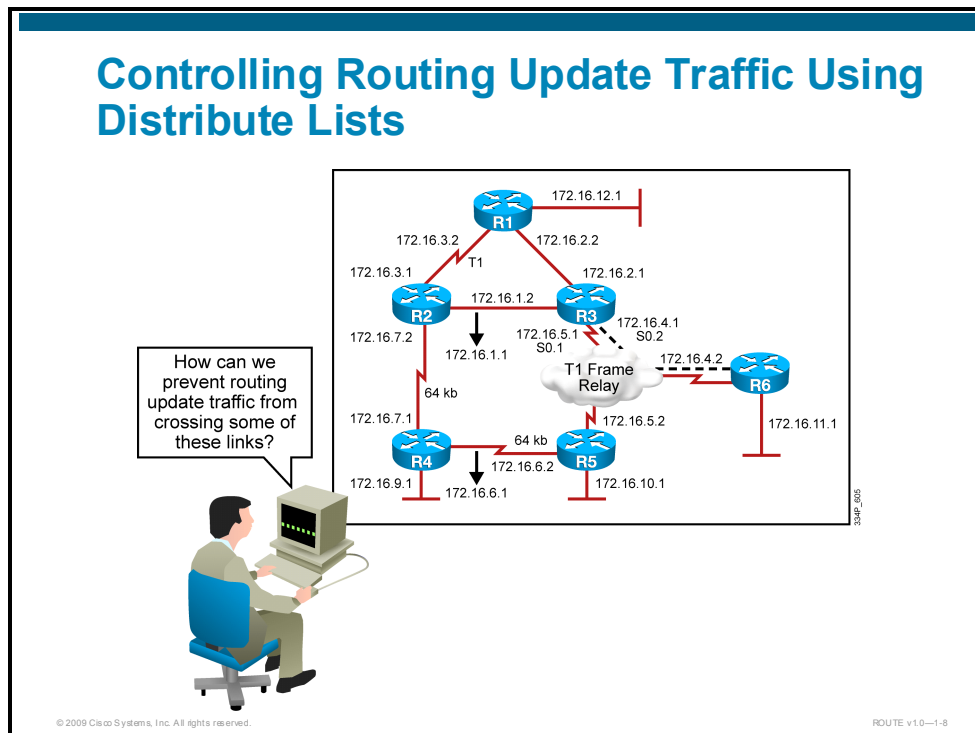


The figure shows how filtering is used to control incoming routing update traffic. Only inbound filtering is shown.

- Step 1** A routing update arrives at a router interface, and the router stores the packet in the interface buffer and triggers the CPU to make a decision.
- Step 2** The router CPU checks if there is an incoming filter that is applied to this interface. If no filter is applied, the routing update packet is processed normally. If a filter is applied, the next step is taken.
- Step 3** The router CPU checks if there is an entry for this address in the routing update packet. If there is no entry, the route is dropped. If an entry exists, the next step is taken.
- Step 4** The router CPU processes the routing update packet according to the filter configuration.

How Distribute Lists Work

This topic describes how distribute lists can be used to control routing updates.



The ways to control or prevent dynamic routing updates are as follows:

- **Passive interface:** As previously stated, this feature prevents all routing updates from being sent through an interface. For EIGRP and OSPF, this method includes Hello protocol packets.
- **Default routes:** This feature instructs the router that if it does not have a route for a given destination, it should send the packet to the default route. Therefore, no dynamic routing updates about the remote destinations are necessary.
- **Static routes:** This feature allows routes to remote destinations to be manually configured in the router. Therefore, no dynamic routing updates about the remote destinations are necessary.

Another way to control routing updates is a technique that is called a distribute list. A distribute list allows the application of an access list to the routing updates. You may be familiar with access lists that are associated with an interface and how they are used to control IP traffic. However, routers can have many interfaces, and route information can also be obtained through route redistribution, which does not involve an interface at all.

Additionally, access lists do not affect traffic that is originated by the router, so applying one to an interface would have no effect on the outgoing routing advertisements. When you link an access list to a distribute list, routing updates can be controlled no matter what their source is.

Configure access lists in global configuration mode, and then configure the associated distribute list under the routing protocol. The access list should permit the networks that will be advertised or redistributed and deny the networks that will remain hidden.

The router then applies the access list to the routing updates for that protocol. Options in the **distribute-list** command allow updates to be filtered based on three factors:

- Incoming interface
- Outgoing interface
- Redistribution from another routing protocol

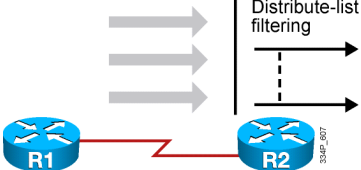
Using a distribute list gives the administrator great flexibility in determining which routes that the router distributes.

Using Distribute Lists to Control Routing Updates

This topic describes how to implement and verify the distribute list route-filtering technique.

Steps to Configure Distribute List Filters

- Define the traffic filtering requirements to permit or deny routes using one of these two methods:
 - Configure an access list (ACL).
 - Configure a route map.
- Configure a distribute list to use the ACL or a route map:
 - Apply it to the inbound or outbound updates.



© 2009 Cisco Systems, Inc. All rights reserved.ROUTE v1.0—4-9

When you are planning how to configure distribute list filters, you must complete the following steps and create an implementation plan for the filters:

- Define traffic filtering requirements to permit or deny routes using one of these two methods:
 - Configure an access list.
 - Configure a route map.
- Configure a distribute list to use the access list or route map:
 - Apply it to the inbound or outbound updates.

Configuring a Distribute List Filter

- A distribute list filter can be applied to transmitted, received, or redistributed routing updates.

R1(config)#

```
router rip
 redistribute ospf 1 metric 5
 distribute-list 10 out OSPF 1
```

- Filtering of updates being advertised from OSPF into RIP routing protocol according to access list 10.

R1(config)#

```
router EIGRP 100
 distribute-list 7 in Serial0
```

- Filtering of networks received in updates from interface Serial0 according to access list 7.

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-4-10

You can filter routing update traffic for any protocol by defining an access list and applying it to the specific routing protocol. Use the **distribute-list** command and link it to an access list to complete the filtering of routing update traffic. (The inbound **distribute-list** command allows the use of a route map instead of an access list.)

A distribute list enables you to filter routing updates that are coming into a specific interface from going out to neighboring routers using the same routing protocol. A distribute list also allows you to filter routes that are redistributed from other routing protocols or sources.

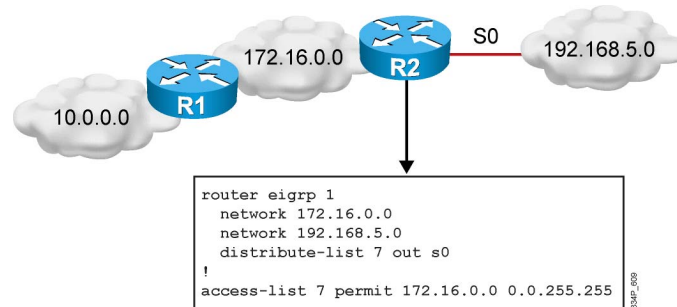
Use the **distribute-list out** command to assign the access list to filter outgoing routing updates or to assign it to routes that are being redistributed into the protocol. The **distribute-list out** command cannot be used with link-state protocols for blocking outbound link-state advertisements (LSAs) on an interface. The **distribute-list out** command in the figure filters the updates that are being redistributed from OSPF into RIP. Access list 10 is used to define the networks that are permitted or denied.

Use the **distribute-list in** command to cause the access list to filter incoming routing updates that are coming in through an interface. This command prevents most routing protocols from placing the filtered routes in their databases. When this command is used with OSPF, the routes are placed in the database but not in the routing table. The **distribute-list in** command in the figure filters the networks that are received in updates from interface Serial0 according to access list 7.

For more details about the **distribute-list out** and **distribute-list in** commands, go to the Cisco IOS IP Routing: Protocol-Independent Command Reference via the following link:
https://www.cisco.com/en/US/docs/ios/iproute_pi/command/reference/iri_book.html

Filtering Routing Updates with a Distribute List

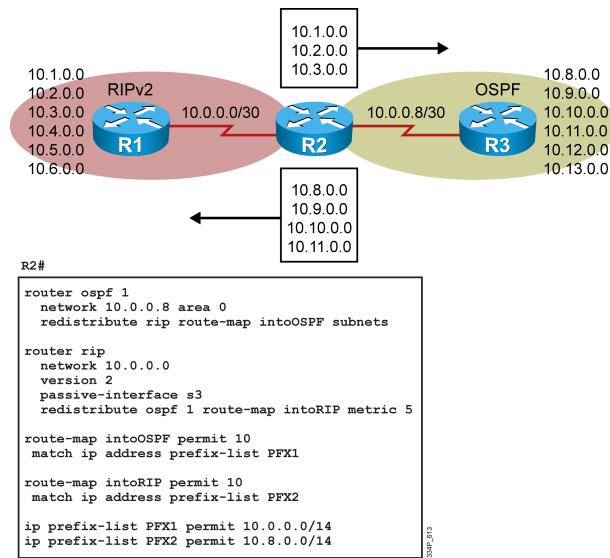
- Hides network 10.0.0.0 using interface filtering



The figure shows the configuration of router R2, which is hiding network 10.0.0.0 behind the Serial0 interface. The **distribute-list 7 out s0** command applies access list 7 to routing updates that are sent out from Serial0 to other routers that are running this routing protocol. This access list permits routing information about network 172.16.0.0 only.

The implicit “deny any” statement at the end of the access list prevents routing updates about any other networks from being advertised. As a result, network 10.0.0.0 is hidden from the rest of the network.

Controlling Redistribution with Prefix Lists



With redistribution, using a distribute list helps prevent all routes from being redistributed from RIP routing protocol into OSPF and vice versa.

As shown in the figure, only networks 10.1.0.0 to 10.3.0.0 are redistributed from RIP into OSPF. Other networks are not permitted with access list 2 and are denied. Access list 2 allows the subset of RIP routes and denies all others. The distribute list that is configured under OSPF refers to access list 2.

The same redistribution processing applies to networks 10.8.0.0 to 10.11.0.0, originated by OSPF. They can be redistributed from OSPF to RIP. Redistribution into RIP from OSPF is filtered with access list 3. The distribute list that is configured under RIP refers to access list 3 and permits only a subset of networks.

A distribute list hides network information, which may be considered a drawback in some circumstances. In a network with redundant paths, the goal of using a distribute list may be to prevent routing loops. The distribute list permits routing updates that enable only the desired paths to be advertised. Therefore, other routers in the network do not know about the other ways to reach the filtered networks.

How Prefix Lists Work

This topic describes how prefix lists can be used to control routing updates.

IP Prefix Filters

- Traditionally, IP prefix filters were implemented with IP access lists configured with the **distribute-list** command.
- Prefix lists:
 - Better performance than access lists
 - User-friendly CLI
 - Match routes in part of an address space with a subnet mask that is longer or shorter than a set number

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—4-13

Traditionally, IP prefix filters were implemented with IP access lists that were configured with the **distribute-list** command. IP access lists that are used as route filters have several drawbacks:

- The subnet mask cannot be easily matched.
- Access lists are evaluated sequentially for every IP prefix in the routing update.
- Extended access lists can be cumbersome to configure.

Using the **ip prefix-list** command has several benefits compared with using the **access-list** command. The intended use of prefix lists is limited to route filtering, where access lists were originally intended to be used for packet filtering and were then extended to route filtering.

A router transforms a prefix list into a tree structure, with each branch of the tree serving as a test. Cisco IOS Software determines a verdict of either “permit” or “deny” much faster this way than when sequentially interpreting access lists.

The command-line interface (CLI) that you use to configure the **ip prefix-list** command provides you with the ability to assign a line number to each line of the prefix list. The router will use this number to sort the entries in the prefix list. If you initially sequence the lines with some gaps in between the sequence numbers, you can insert additional lines later. You can also remove individual lines without removing the entire list.

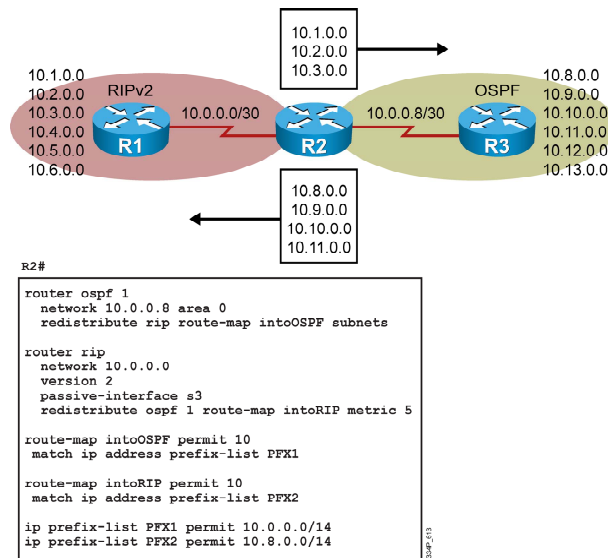
Routers match network numbers in a routing update against the prefix list using as many bits as indicated. For example, you can specify a prefix list to be 10.0.0.0/16, which will match 10.0.0.0 routes but not 10.1.0.0 routes.

The prefix list can specify the size of the subnet mask and can also indicate that the subnet mask must be in a specified range.

Prefix lists are similar to access lists in many ways. A prefix list can consist of any number of lines, each of which indicates a test and a result. The router can interpret the lines in the specified order, although Cisco IOS Software optimizes this for processing in a tree structure. When a router evaluates a route against the prefix list, the first line that matches will result in either a “permit” or “deny.” If none of the lines in the list match, the result is “implicitly deny.”

Testing is done using prefixes. The router compares the indicated number of bits in the prefix with the same number of bits in the network number in the update. If they match, testing continues with an examination of the number of bits set in the subnet mask. The prefix list line can indicate a range in which the number must be in order to pass the test. If you do not indicate a range in the prefix line, the subnet mask must match the prefix size.

Controlling Redistribution with Prefix Lists



With redistribution, a prefix list can be used instead of a distribute list.

As shown in the figure, networks 10.1.0.0 to 10.3.0.0 are redistributed only from RIP into OSPF. The other networks are not permitted with the prefix list PFX1 and are denied. PFX1 allows the subset of RIP routes and denies all others. The route map into OSPF matches the networks that are permitted by PFX1.

The same applies to networks 10.8.0.0 to 10.11.0.0, originated by OSPF. They can be redistributed from OSPF to RIP. The redistribution into RIP from OSPF is filtered with the prefix list PFX2. The route map into RIP matches networks that are permitted by PFX2.

In a network with redundant paths, the goal of using prefix lists and route maps may be to prevent routing loops. The prefix lists permit routing updates that enable only the desired paths to be advertised. Therefore, other routers in the network do not know about the other ways to reach the filtered networks.

Prefix List Matching Rules

- Filter by exact prefix length:
 - mask filtering “/”
- Filter within a range:
 - using **ge**
 - using **le**
 - using **ge** and **le**
- The matching process also considers the subnet mask.

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-4-15

Prefix lists are configured to filter traffic based on a match of an exact prefix length or a match within a range when the **ge** and **le** keywords are used. The **ge** and **le** keywords are used to specify a range of prefix lengths. They provide a more flexible configuration than can be obtained using only the network/length arguments. A prefix list is processed using an exact match when neither the **ge** nor **le** keyword is specified. If only the **ge** value is specified, the range is the value that is entered for the *ge-length* argument to a full 32-bit length. If only the **le** value is specified, the range is from the value that is entered for the network/length argument to the *le-length* argument. If both the **ge** and **le** values are specified, the range is between the values that are used for the *ge-length* and *le-length* arguments.

The following formula shows this behavior: $length < ge-length < le-length \leq 32$

Using a Prefix List to Control Routing Updates

This topic describes how to implement and verify a prefix list for route filtering and redistribution.

Prefix List Matching Without ge or le

- Similar to IP access lists with no wildcard bits

```
R2 (config) #  
ip prefix-list MyMatchList permit 192.168.0.0/16
```

- Which prefixes are matched?
 - 192.168.0.0/16: Match
 - 192.168.0.0/20: No Match
 - 192.168.2.0/24: No Match

© 2009 Cisco Systems, Inc. All rights reserved. RD UTE v1.0-4-16

To create a prefix list or to add a prefix list entry, use the **ip prefix-list** command in global configuration mode. Prefix lists are configured with **permit** or **deny** keywords to either permit or deny a prefix based on a matching condition. An “implicit deny” statement is applied to traffic that does not match any prefix list entry. A prefix list entry consists of an IP address and a bit mask. The IP address can be for a classful network, a subnet, or a single host route. The bit mask is a number from 1 to 32.

If the **seq** keyword is configured without a sequence number, the default sequence number is 5. In this scenario, the first prefix list entry (like in the figure) is assigned the number 5, and subsequent prefix list entries increment by 5. For example, the next two entries would have sequence numbers 10 and 15. If a sequence number is entered for the first prefix list entry but not for subsequent entries, the subsequent entry numbers increment by 5. For example, if the first configured sequence number is 3, subsequent entries will be 8, 13, and 18. Default sequence numbers can be suppressed by entering the **no ip prefix-list** command with the **seq** keyword.

The evaluation of a prefix list starts with the lowest sequence number and continues down the list until a match is found. When an IP address match is found, the permit or deny statement is applied to that network, and the remainder of the list is not evaluated.

Tip For best performance, the most frequently processed prefix list statements should be configured with the lowest sequence numbers. The **seq number** keyword and argument can be used for resequencing.

Prefix list entries without the **ge** or **le** option only match the route with the specified IP address and subnet mask. In the figure, the prefix list entry **permit 192.168.0.0/16** will not match the route 192.168.2.0/24, because of the mismatch in the IP address. It will also not match the route 192.168.0.0/20, because of the mismatch in the subnet mask.

For more details about the **ip prefix-list** command, go to the Cisco IOS IP Command Reference via the following link:

http://www.cisco.com/en/US/docs/ios/12_3/ipaddr/command/reference/ip1_i2g.html

Prefix List Matching with ge or le

- A prefix list entry with the **ge** or **le** keyword matches any prefix within a specified address space for which the subnet mask falls within the specified limits.

R2 (config) #

```
ip prefix-list List1 permit 192.168.0.0/16 le 20
ip prefix-list List2 permit 192.168.0.0/16 ge 18
```

- Which prefixes are matched?
 - 192.168.0.0/16, List1: Match
 - 192.168.0.0/16, List2: No Match

© 2009 Cisco Systems, Inc. All rights reserved.

RO UTE v1.0—4-17

Prefix list entries for which you specify the **ge** or **le** option match any prefix within the address space that you specify by the network/length parameter, as long as the subnet mask length of the route falls within the range that you specify by the **le** and **ge** parameters.

In the first example—the prefix list named List1—the route 192.168.2.0/24 is not matched by the prefix list entry **permit 192.168.0.0/16 le 20**, even though the IP address falls within the specified address range, because the subnet mask is too long.

In the second example—the prefix list named List2—the route 192.168.0.0/16 is not matched by the prefix list entry **permit 192.168.0.0/18 ge 18**, because the subnet mask is too short.

Configuring Prefix List Examples

What will be matched by?

1. **ip prefix-list A permit 0.0.0.0/0 ge 32**
2. **ip prefix-list B permit 128.0.0.0/2 ge 17**
3. **ip prefix-list C permit 0.0.0.0/0 le 32**
4. **ip prefix-list D permit 0.0.0.0/0**
5. **ip prefix-list E permit 0.0.0.0/1 le 24**

All host routes:

1. Only hosts
2. Any subnet in a Class B address space
3. All routes
4. Just the default route
5. Any prefix in a Class A address space covering at least 256 addresses

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-4-19

The figure contains some commonly used prefix list examples. They would match the following subnets:

1. **Only hosts:** The prefix 0.0.0.0/0 means that none of the bits needs to be matched; **ge 32** means that the mask needs to be exactly 32 bits, limiting this statement to hosts (and not subnets).
2. **Any subnet in a Class B address space:** The prefix 128.0.0.0/2 means that the first two bits of the IP address must be the same as the first two bits of 128 (10000000), thus limiting the entry to Class B addresses. **Ge 17** means that the mask should not be a Class B mask (16 bits), but 17 bits or greater, thus allowing any subnet address within the Class B space.
3. **All routes:** The prefix 0.0.0.0/0 means that none of the bits needs to be matched; **le 32** means that the mask needs to be 32 bits or less, thus matching all routes.
4. **Just the default route:** There is no indication of a **ge** or **le** argument in this statement, so the route has to match what is entered, 0.0.0.0/0, which is a default route.
5. **Any prefix in a Class A address space covering at least 256 addresses:** 0.0.0.0/1 means that the first bit must be 0, which indicates a Class A address. **Le 24** means that the mask must be 24 bits or less, which points to any address in Class A covering 256 addresses or more.

How Route Maps Work

This topic describes the functionality of route maps, which are powerful and flexible configuration tools.

Route Maps

- Route maps are similar to a scripting language for these reasons:
 - They work like access lists, but are more sophisticated.
 - They offer top-down processing.
 - When one of them finds a match, it stops searching.
 - Lines are sequence-numbered for easier editing.
 - Insertion of lines
 - Deletion of lines
 - Route maps are named, rather than numbered, for easier documentation.
 - Match criteria and set criteria can be used; similar to the if-then logic in scripting languages.

© 2009 Cisco Systems, Inc. All rights reserved. RD UTE v1.0-4-20

Route maps are complex access lists that allow conditions to be tested against a packet or route using the **match** commands. If the conditions match, actions can be taken to modify attributes of the packet or route. These actions are specified by the **set** commands.

A collection of route map statements that have the same route map name is considered one route map. Within a route map, each route map statement is numbered and can be edited individually.

The statements in a route map are analogous to the lines of an access list. Specifying the match conditions in a route map is similar to specifying the source and destination addresses and masks in an access list.

One major difference between route maps and access lists is that route maps can use the **set** commands to modify the packet or route.

Route Map Applications

The common uses of route maps are as follows:

- Redistribution route filtering:
 - A more sophisticated alternative to distribute lists
- PBR:
 - The ability to determine a routing policy based on criteria other than the destination network
- BGP policy implementation:
 - The primary tool for defining BGP routing policies

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-4-21

Network administrators use the route map tool for many purposes. Several of the more common applications for route maps are as follows:

- **Route filtering during redistribution:** Redistribution nearly always requires some amount of route filtering. Although distribute lists can be used for this purpose, route maps offer an added benefit of manipulating routing metrics through the use of the **set** commands.
- **Policy-based routing (PBR):** Route maps can be used to match source and destination addresses, protocol types, and end-user applications. When a match occurs, a **set** command describes the interface or next-hop address to which the packet should be sent. PBR allows the operator to define a routing policy other than basic destination-based routing using the routing table.
- **BGP:** Route maps are the primary tools for implementing a BGP policy. Network administrators assign route maps to specific BGP sessions (neighbors) to control which routes are allowed to flow into and out of the BGP process. In addition to their use in filtering, route maps allow for sophisticated manipulation of BGP path attributes.

Route Map Operation

- A list of statements constitutes a route map.
- The list is processed in a top-down manner, like access lists.
- The first match that is found for a route is applied.
- The sequence number is used for inserting or deleting specific route map statements.

```
route-map MyRouteMap permit 10
  { match statements }
  { match statements }
  { set statements }
  { set statements }
route-map MyRouteMap deny 20
:: :: ::
:: :: ::
route-map MyRouteMap permit 30
:: :: ::
:: :: ::
```

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-4-22

Route maps operate in a manner similar to access lists. When determining which routes will be redistributed from one protocol to the next, the router checks each route against the route map, beginning with the top line.

Each line is sequence-numbered, both for top-down processing purposes and for editing purposes. Lines can be added or removed from a route map as changes are required.


Each line has a permit or deny statement. If a route is matched in the matching statements and the line has a permit statement, the router sets the metrics or other defined conditions and permits the redistribution of that route. The route map stops processing at the first match.

If the packet is matched and the route map line has a deny statement, the router stops at the matched line in the map and does not redistribute that route. Routes are filtered using this method.

Routers search for matches on a line-by-line basis. If a router reaches the bottom of the route map without finding a match, it denies the route from being redistributed. There is always an implicit deny statement at the end of a route map.

Route Map Operation (Cont.)

- The match statement may contain multiple references.
- Multiple match criteria in the same line: logical OR.
- At least one reference must permit the route for it to be a candidate for redistribution.



- Multiple match statements on separate lines: logical AND.
- All match statements must permit the route for it to remain a candidate for redistribution.
- Route map permit or deny statements determine if the candidate will be redistributed.

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-4-23

Matching statements in a route map can be complex. Multiple match criteria in the same line are processed with OR logic. Separate match criteria can also be applied vertically under a route map line. In this case, each match uses AND logic.

A route map may consist of multiple route map statements. The statements are processed top-down, like statements are processed for an access list. The first match that is found for a route is applied. The sequence number is used for inserting or deleting specific route map statements in a specific place in the route map.

The **match** route map configuration command defines the conditions that are to be checked. The **set** route map configuration command defines the actions that you should follow if there is a match.

The single-match statement may contain multiple conditions. At least one condition in the match statement must be true to consider the statement a match (logical OR). A route map statement may contain multiple match statements. All match statements in the route map statement must be true to consider the route map statement a match (logical AND).

The sequence number specifies the order in which conditions are checked. For example, if there are two statements in a route map named MyRouteMap, one with sequence number 10 and the other with sequence number 20, the one with sequence number 10 is checked first. If the match conditions in the statement with sequence number 10 are not met, the statement with sequence number 20 is checked.

As is true of an access list, there is an implicit “deny any” statement at the end of a route map. The consequences of this depend on how the route map is used.

Steps to Configure a Route Map

- Define the route map conditions:
 - Define the conditions to match.
 - Define the action to be taken on a match.
- Attach the route map to an interface.

© 2009 Cisco Systems, Inc. All rights reserved.

RO UTE v1.0—4-24

When you are planning to configure route maps, you must follow these steps and create an implementation plan for them:

- Define the route map conditions:
 - Define the conditions to match.
 - Define the action to be taken when each condition is matched.
- Attach the route map to an interface.

Configuring a Route Map

R1(config)#

```
route-map MyRouteMap permit 10
```

- Defines the route map with the name MyRouteMap conditions

R1(config-route-map)#

```
match ip address prefix-list MyList
```

- Matches based on the prefix list MyList when defining the conditions to match

R1(config-route-map)#

```
set interface ethernet 0
```

- Defines that interface Ethernet 0 be used to forward packets that pass a match clause

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-4-25

To define the conditions for redistributing routes from one routing protocol into another, or to enable policy routing, use the **route-map** command in global configuration mode and the **match** and **set** commands in route map configuration modes.

The commands in the figure show the route map configuration that is used to enable policy routing. The **match** and **set** commands are used with the **route-map** command to define the conditions for policy-routing packets. The **match** command specifies the condition under which policy routing occurs. The **set** command specifies the routing actions to perform if the criterion that is enforced by the **match** command is met. You might want to route packets some way other than along the obvious shortest path.

The *sequence-number* argument works as follows:

- If no entry is defined with the supplied tag, an entry is created with the *sequence-number* argument set to 10.
- If only one entry is defined with the supplied tag, that entry becomes the default entry for the following **route-map** command. The *sequence-number* argument of this entry is unchanged.
- If more than one entry is defined with the supplied tag, an error message is printed to indicate that a *sequence-number* argument is required.
- If the **no route-map map-tag** command is specified (with no *sequence-number* argument), the whole route map is deleted.

The **match** command is applied within a route map. The following is a general list of match criteria. Some criteria are used for BGP policy, some criteria are used for PBR, and some criteria are used for redistribution filtering.

- **match community**
- **match interface**
- **match ip address**
- **match ip next-hop**
- **match ip route-source**
- **match length**
- **match metric**
- **match route-type**
- **match tag**

The **set** commands are used within a route map to change or add characteristics, such as metrics, to any routes that have met a match criterion. The following is a general list of set actions:

- **set automatic-tag**
- **set default interface**
- **set interface**
- **set ip default next-hop**
- **set ip default next-hop verify-availability**
- **set ip global**
- **set ip next-hop**
- **set ip next-hop verify-availability**
- **set ip vrf**
- **set level (IP)**
- **set local-preference**
- **set metric (BGP-OSPF-RIP)**
- **set metric-type**
- **set next-hop**
- **set tag (IP)**

For more details about the **route-map**, **match**, and **set** commands, go to the Cisco IOS IP Routing: Protocol-Independent Command Reference via the following link:
https://www.cisco.com/en/US/docs/ios/iproute_pi/command/reference/iri_book.html

Attaching a Route Map to an Interface

R1(config-if)#

```
ip policy route-map MyRouteMap
```

- Defines a route map named MyRouteMap to be used for policy routing on an interface.

R1#

```
interface serial 0
  ip policy route-map MyRouteMap
!
route-map MyRouteMap
  match ip address 6 25
  set ip next-hop 172.21.16.18
!
access-list 6 permit 172.16.129 0.0.127
...
access-list 25 permit 10.10.8.0.0.0.0.255
access-list 25 permit 10.14.0.0.0.0.255.255
```

- Set the next hop to 172.21.16.18 for addresses specified by access list numbers 6 or 25

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-1-26

Use the **ip policy route-map** command, in addition to the **route-map** command, and the **match** and **set** commands to define the conditions for policy-routing packets. The **ip policy route-map MyRouteMap** command in the figure is used to define a route map named MyRouteMap to be used for policy routing on an interface. Each **route-map** command has a list of **match** and **set** commands that are associated with it. The **match** commands specify the match criteria—the conditions under which policy routing is allowed for the interface, based on the destination IP address of the packet. The **set** commands specify the set actions—the particular policy routing actions to perform if the criteria that are enforced by the **match** commands are met.

You might enable policy routing if you want your packets to take a route other than the obvious shortest path. In the figure, the **route-map MyRouteMap** command matches IP addresses defined by access control list (ACL) 6 or ACL 25 and sets the next hop to 172.21.16.18 when the condition under the **match** command is met. The configured route map is then attached to interface serial 0 using the **ip policy route-map MyRouteMap** command.

Note In the example provided in the figure, it might have been possible to merge both ACLs into one ACL. These two ACLs are provided as examples. In real configurations, ACLs may be used for different purposes. The **match** statement allows you to match several ACLs.

For more details about the **ip policy route-map** command, go to the Cisco IOS IP Routing: Protocol-Independent Command Reference via the following link:
http://www.cisco.com/en/US/docs/ios/iproute_pi/command/reference/iri_book.html

Using Route Maps to Control Routing Updates

This topic describes how to use the **route-map** command to define the conditions for route filtering and redistribution.

Steps to Configure Redistribution with Route Maps

- Define the route map:
 - Define match statements.
 - Define set statements.
- Define redistribution using the route map.

© 2009 Cisco Systems, Inc. All rights reserved.

RD UTE v1.0-4-27

When you plan to configure route maps, you must follow these steps and create an implementation plan for them:

- Define the route map:
 - Define match statements.
 - Define set statements.
- Define redistribution using route maps.

Route Map Redistribution Commands

R1#

```
route-map rip_to_eigrp deny 10
  match tag 88
route-map rip_to_eigrp permit 20
  set tag 77
```

- Define the route map used during redistribution.

R1(config)#

```
router eigrp 7
  redistribute rip route-map rip_to_eigrp metric 1 1 1 1 1
```

- Configure redistribution from RIP to the EIGRP routing protocol using the route map to filter updates.

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-4-28

When used for redistribution filtering, a route map is applied to the route redistribution process by adding the **route-map** command to the end of the **redistribute** (IP) routing protocol command.

Use the **route-map** command and the **match** and **set** route map configuration commands to define the conditions for redistributing routes from one routing protocol into another. Each **route-map** command has a list of **match** and **set** commands that are associated with it. The **match** commands specify the match criteria—the conditions under which redistribution is allowed for the current **route-map** command. The **set** commands specify the set actions—the particular redistribution actions to perform if the criteria that are enforced by the **match** commands are met. The **no route-map** command deletes the route map.

Use route maps when you want detailed control over how routes are redistributed between the routing processes. Specify the destination routing protocol with the **router** global configuration command. Specify the source routing protocol with the **redistribute** (IP) configuration command.

In the figure, the EIGRP process with AS number 7 is configured locally in R1. The **redistribute rip route-map rip_to_eigrp metric 1 1 1 1 1** command is used to configure redistribution of RIP routes into the EIGRP 7 routing process, where the route map named **rip_to_eigrp** is used to influence the redistribution of RIP routes into the EIGRP routing process. The route map is configured to match the routes that are tagged with **88**. These routes are denied from being redistributed into the EIGRP 7 routing process. The routes without the tag **88** reach the second route map statement and are tagged with **77**. The **set tag 77** command is used in the route map. These routes are allowed to be distributed into the EIGRP 7 routing process.

For more details about the **redistribute** command, go to the Cisco IOS IP Routing: Protocol-Independent Command Reference via the following link:
http://www.cisco.com/en/US/docs/ios/iproute_pi/command/reference/iri_book.html

Route Map Redistribution Commands Example

R1#

```
router ospf 10
 redistribute rip subnets
 route-map redisRIP
 !
 Route-map redisRIP permit 10
 match ip address 23 29
 set metric 500
 set metric-type type-1
```

```
route-map redisRIP deny 20
 match ip address 37
 route-map redisRIP permit 30
 set metric 5000
 set metric-type type-2
 !
 access-list 23 permit 10.1.0.0 0.0.255.255
 access-list 29 permit 172.16.1.0 0.0.0.255
 access-list 37 permit 10.0.0.0 0.255.255.255
```

- Routes matching either access list 23 or 29 are redistributed with an OSPF cost of 500, E1.
- Routes permitted by access list 37 are not redistributed.
- All other routes are redistributed with an OSPF cost metric of 5000, E2.

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-4-29

In this example, RIP version 1 (RIPv1) is being redistributed into OSPF 10. A route map named “redisRIP” has been attached to the **redistribute** (IP) command.

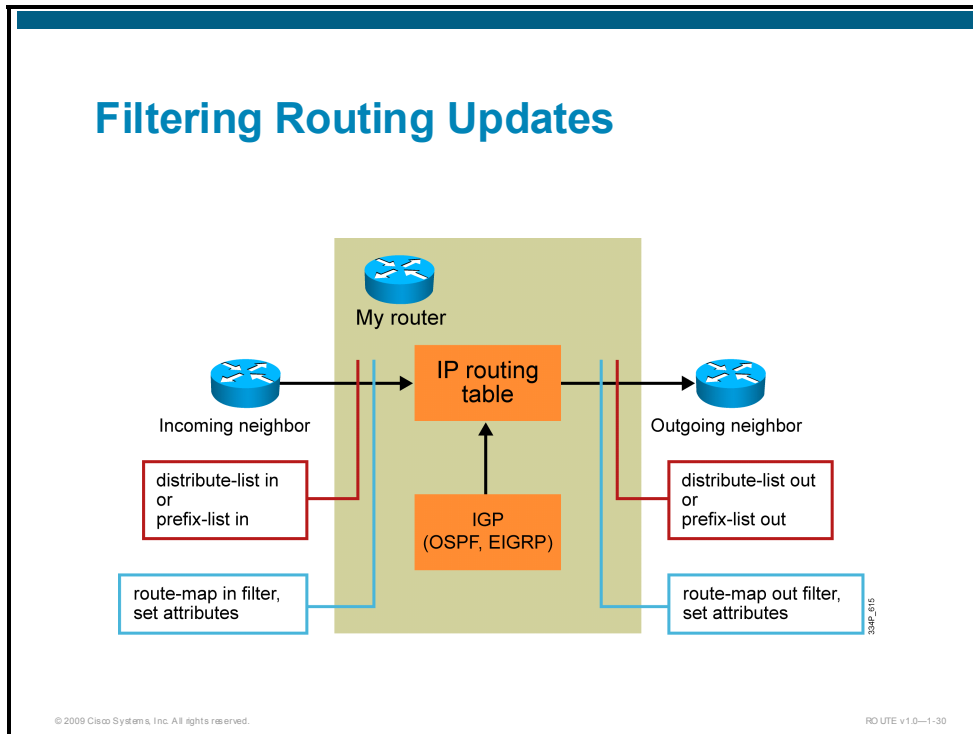
Sequence number 10 of the route map looks for an IP address match in access list 23 or access list 29. If a match is found, the router redistributes the route into OSPF with a cost metric of 500 and sets the new OSPF route to external type 1 (E1).

If there is no match to line 10, move to line 20. If there is a match in access list 37, do not let that route redistribute into OSPF, because the route map statement with sequence number 20 is a deny statement.

If there is no match to sequence number 20, move to 30. Because the route map statement with sequence number 30 is a permit statement and there is no match criterion, all remaining routes are redistributed into OSPF with a cost metric of 5000 and an external metric of type 2.

Using Route Maps to Filter Routes

This topic describes how to implement and verify route maps with route filtering.



You can apply prefix lists, distribute lists, and route maps on either incoming or outgoing information. The incoming prefix list, the incoming distribute list, and the incoming route map must all permit the routes that are received from a neighbor before they will be accepted into the IP routing table. Outgoing routes must pass the outgoing distribute list, the outgoing prefix list, and the outgoing route map before they will be transmitted to the neighbor.

Suppressing Routing Updates Using a Passive Interface

This topic describes how to configure dynamic routing protocol updates for passive interfaces and provides several examples where a passive interface can be used to limit overhead of routing protocols.

Passive Interface

- Routers can have many interfaces:
 - Not all are allowed to send and receive routing updates.
- The suppression of routing updates over some interfaces can be enabled.
- Passive interfaces are used to:
 - Suppress updates on an interface
 - Suppress updates on all interfaces
- Each routing protocol has different rules:
 - **OSPF**: Routing information is neither sent nor received.
 - **EIGRP**: Routing process is disabled on an interface.
 - **RIP**: The sending of updates is disabled, but listening is allowed.

© 2009 Cisco Systems, Inc. All rights reserved.

RD UTE v1.0—4-31

Routers can have many interfaces, and there are times when you must include an interface in a **network** command, although you do not want that interface to participate in the routing protocol. In such cases, the routing protocol might be configured differently to exclude the interface from routing purposes. A better solution is to suppress routing updates over all or some interfaces, which can be enabled globally in the routing process.

A passive interface can be configured to suppress updates on the specific interface or on all interfaces.

General rules for passive interfaces state that if you disable the sending of routing updates on an interface, the particular subnet will continue to be advertised to other interfaces and the interface will continue to receive and process updates from other routers. It is possible to set all interfaces as passive by default. You can then enable specific interfaces one by one. This is useful in ISP and large enterprise networks, where many of the distribution routers have more than 200 interfaces. This old solution meant entering 200 or more **passive-interface** statements.

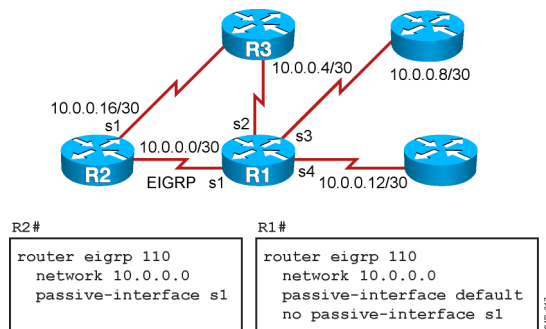
General rules are not always the same for different routing protocols. The passive interface behavior varies from one protocol to another.

For the OSPF protocol, OSPF routing information is neither sent nor received through the specified router interface that is configured as passive. The specified interface address appears as a stub network in the OSPF domain.

EIGRP is disabled on an interface that is configured as passive, although it advertises the route. For passive interfaces, EIGRP suppresses the exchange of hello packets between two routers, which results in the loss of their neighbor relationship. This not only stops routing updates from being advertised but also suppresses incoming routing updates.

With RIP, passive interfaces do not send multicast updates but can listen to incoming updates from other RIP-speaking neighbors. Routers are still able to receive updates on a passive interface and use them in the routing table.

Using the passive-interface Command



In the figure, R1 and R2 run EIGRP and have a **network** statement that encompasses all their interfaces; however, you should run EIGRP only on the link between R1 and R2.

R1 has several interfaces; the **passive-interface default** command was used to set all interfaces to passive, and the **no passive-interface** command was used to enable the one interface from which the EIGRP updates are desired. R2 has only two interfaces; the **passive-interface** command was used for the one interface that is not to participate in EIGRP routing.

It is important to understand how this configuration affects the information that is exchanged between R1 and R2, and between them and R3. Unless you configure another routing protocol and redistribute between it and EIGRP, R1 does not tell R3 that it has a way to reach the networks that are advertised by R2 via EIGRP.

Likewise, R2 does not tell R3 that it has a way to reach the networks that are advertised by R1 via EIGRP.

Redundancy is built into this network. However, the three routers are not able to use the redundancy effectively. For example, if the link between R3 and R1 fails, R3 does not know that it has an alternate route through R2.

For more details about the **passive-interface** command, go to the Cisco IOS IP Routing: Protocol-Independent Command Reference via the following link:
http://www.cisco.com/en/US/docs/ios/iproute_pi/command/reference/iri_book.html

Summary

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

Summary

- Sending an excessive number of routing updates, running different protocols in different areas within the same AS, and incorrectly configuring route filters are common causes of network performance issues.
- Distribute lists use access lists or route maps to define the traffic filtering requirements to permit or deny routes.
- The **distribute-list** command allows updates to be filtered based on the incoming interface, outgoing interface, or redistribution from another routing protocol.
- Prefix lists significantly improve the performance of complex filters relative to traditional IP prefix filters that are implemented with IP access lists.
- The matching process for prefix lists takes into account the subnet mask. A filter can be configured for the exact prefix length, or it can be configured within a range using the **ge** and **le** keywords.

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-4-33

Summary (Cont.)

- Route maps are complex ACLs that allow conditions to be tested against a packet or route using the **match** command. If the conditions match, then actions, specified by the **set** command, can be taken to modify the attributes of the packet or route.
- When used for redistribution filtering, a route map is applied to the route redistribution process by adding the **route-map** command to the end of the **redistribute** routing protocol command.
- When route maps are used to filter routing update traffic, they can be applied to incoming or outgoing traffic. The routes that are permitted can have their attributes set or changed by the **set** command in the route map.
- A passive interface can be configured to suppress updates on a specific interface or on all interfaces.

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-4-34

Operating a Network Using Multiple IP Routing Protocols

Overview

Simple routing protocols work well for simple networks, but as networks grow and become more complex, it may be necessary to change routing protocols. Often, the transition between routing protocols takes place gradually, so there are multiple routing protocols that are operating in the network for variable lengths of time. This lesson examines several reasons for using more than one routing protocol.

It is important to understand how to exchange routing information between these routing protocols and how Cisco routers operate in an environment with multiple routing protocols. This lesson describes redistribution from one routing protocol to another and how Cisco routers make route selections when multiple protocols are active in the network.

Objectives

Upon completing this lesson, you will be able to explain what route redistribution is and why it may be necessary. This ability includes being able to meet these objectives:

- Describe a complex routing network
- Define route redistribution
- Determine default metrics for redistributed routes
- Determine where to redistribute
- Identify caveats of redistribution

Describe a Complex Routing Network

This topic describes the need to use multiple IP routing protocols.

Complex Routing Scenarios

- Networks can be dispersed:
 - Politically
 - Geographically
 - As a result of acquisitions—company mergers
- Careful design and traffic optimization are required:
 - Redistribution
 - Routing traffic filtering
 - Summarization

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—4-2

There are many reasons for complex routing scenarios. When a network grows, it typically becomes more complex as network administrators adjust it based on factors like political borders, geographical borders, and mergers with other companies. In all such situations, network administrators face complex routing scenarios, in which the use of multiple IP routing protocols is not an exception but the rule. With large numbers of routers and routing protocols running, the number of routing updates increases. As the network grows larger, traffic from those updates can slow down the network, indicating that a change is required.

There are many possible solutions when the design of a network must be changed or optimized. A scalable routing protocol may be necessary, and the following careful design and traffic optimizations are suggested:

- Redistribution
- Routing traffic filtering
- Summarization

Routing traffic filtering and summarization are described in separate lessons. In this lesson, redistribution between different routing protocols is described.

Using Multiple Routing Protocols

- Temporarily during conversion or migration only
- Application-specific protocols:
 - One size does not always fit all
- Political boundaries:
 - Groups that do not work well with others
- Mismatch between devices:
 - Multivendor interoperability
 - Host-based routers

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—4-3

Multiple routing protocols may be necessary in the following situations:

- When you are migrating from an older interior gateway protocol (IGP) to a new IGP. Multiple redistribution boundaries may exist until the new protocol has completely displaced the old protocol. The same applies to company mergers between companies that are each using a different routing protocol.
- When the use of a new protocol is desired, but the old routing protocol is needed for host systems—for example, for UNIX host-based routers that are running Routing Information Protocol (RIP).
- When some departments do not want to upgrade their routers to support a new routing protocol.
- In mixed-router vendor environments. In these environments, you can use a routing protocol that is specific to Cisco, such as Enhanced Interior Gateway Routing Protocol (EIGRP), in the Cisco portion of the network and a common standards-based routing protocol, like Open Shortest Path First (OSPF), to communicate with devices from other vendors.

When multiple routing protocols are running in different parts of the network, there may be a need for hosts in one part of the network to reach hosts in the other part. One solution is to advertise a default route into each routing protocol, but this is not always the best policy. The network design may not allow default routes.

If there is more than one way to get to a destination network, routers may need information about routes in other parts of the network to determine the best path to that destination. Additionally, if there are multiple paths, a router must have sufficient information to determine a loop-free path to the remote networks.

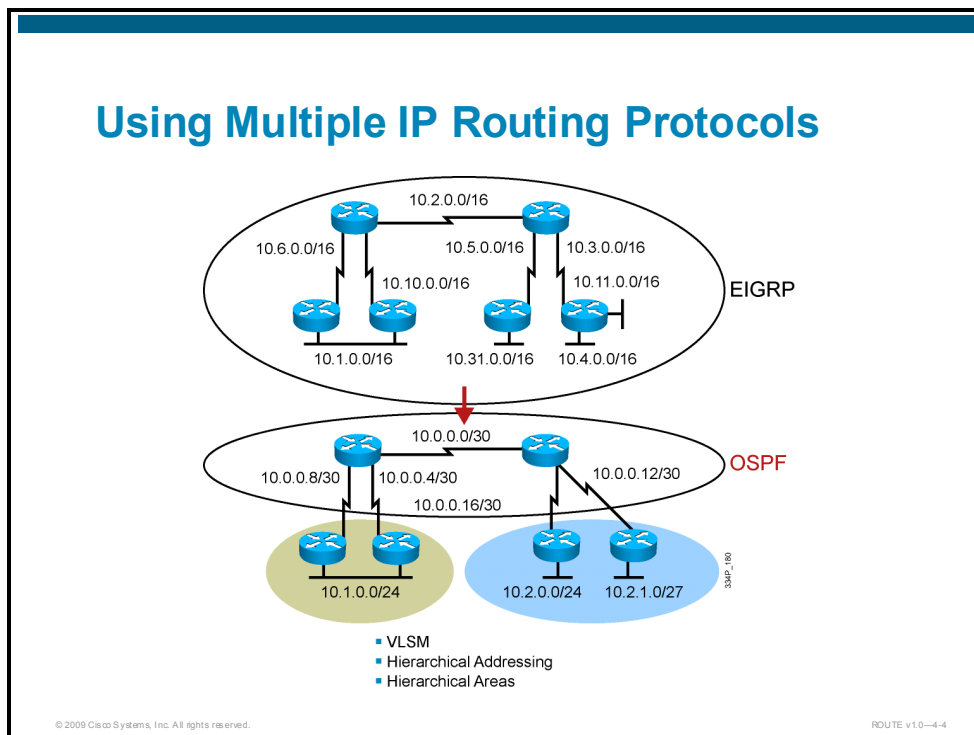
Cisco routers allow internetworks using different routing protocols, referred to as routing domains or autonomous systems, to exchange routing information through a feature called route redistribution.

Redistribution is how routers connect different routing domains so that they can exchange and advertise routing information between the different autonomous systems.

Note The term autonomous system (AS), as used here, refers to an internetwork that is using any one of many routing protocols. The specific routing protocol may be an IGP or exterior gateway protocol (EGP). This is a different use of the term “AS” than is used with Border Gateway Protocol (BGP).

Defining Route Redistribution

This topic describes the need to exchange routing information between different protocols to allow for complete network convergence.



Whatever the reason for running multiple protocols, network administrators must conduct migration from one routing protocol to another or redistribution of routing information between them carefully and thoughtfully.

It is important for network administrators to understand what must be changed and to create a detailed plan before making any changes. An accurate topology map of the network and an inventory of all network devices are also critical for success.

Network administrators must keep in mind the requirements and capabilities of routing protocols when planning redistribution. Link-state routing protocols, such as OSPF, require a hierarchical network structure. Network administrators need to decide which routers will reside in the backbone area and how to divide the other routers into areas. While EIGRP does not require a hierarchical structure, it operates much more effectively within one. Network administrators must carefully plan the redistribution strategy to avoid disrupting network traffic or causing outages.

Redistribution

- Routes learned by some other means are selectively redistributed into a routing protocol from one of three sources:
 - Another routing protocol
 - Static routes
 - Directly connected routes
- Routing loop prevention:
 - Only routes used by the router itself are redistributed.
- Double redistribution inside the same router is not allowed.

© 2009 Cisco Systems, Inc. All rights reserved.

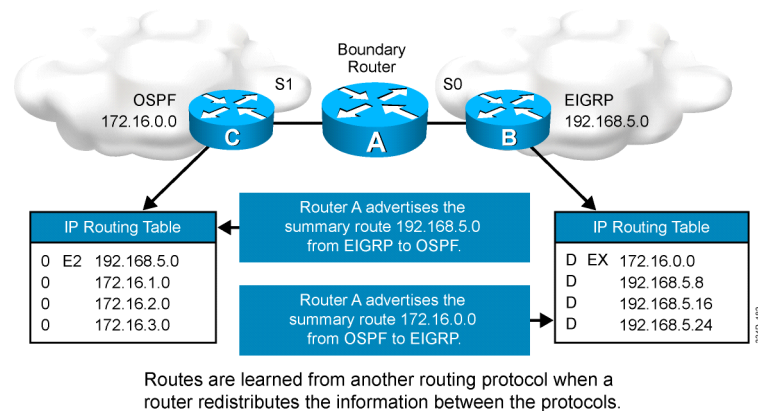
ROUTE v1.0—4-5

Redistribution is the process of using a routing protocol to advertise routes that are learned by usual means of learning routes, such as by another routing protocol, static routes, or directly connected routes.

While it is desirable that you run a single routing protocol throughout your entire IP internetwork, multiprotocol routing is common for many reasons. These reasons include company mergers, multiple departments that are managed by multiple network administrators, and multivendor environments. Running different routing protocols is often part of a network design. In any case, if you have a multiprotocol environment, redistribution is a necessity.

For you to be able to have a scalable solution and limit the amount of routing update traffic, the redistribution process must selectively insert the routes that are learned. Redistribution can lead to routing loops, which must be avoided. Only routes that are used by the router itself should be redistributed, and double redistribution inside the same router should not be allowed.

Redistributing Route Information



Within each AS, the internal routers have complete knowledge about their network. The router that interconnects the autonomous systems is called a boundary router. The boundary router must be running all the routing protocols that will be exchanging routes.

In most cases, route redistribution must be configured to redistribute routes from one routing protocol to another routing protocol. The only time that redistribution is automatic in IP routing protocols is between Interior Gateway Routing Protocol (IGRP) and EIGRP processes that are running on the same router and using the same AS number.

When a router redistributes routes, it allows a routing protocol to advertise routes that were not learned through that routing protocol. These redistributed routes could have been learned via a different routing protocol, such as when redistributing between EIGRP and OSPF, and they also could have been learned from static routes or by a direct connection to a network.

Routers can redistribute static and connected routes as well as routes from other routing protocols.

Redistribution is always performed outbound. The router that is doing redistribution does not change its routing table. When, for instance, redistribution between OSPF and EIGRP is configured, the OSPF process on the boundary router takes the EIGRP routes in the routing table and advertises them as OSPF routes to its OSPF neighbors.

Likewise, the EIGRP process on the boundary router takes the OSPF routes in the routing table and advertises them as EIGRP routes to its EIGRP neighbors. Then, both autonomous systems know about the routes of the other system, and each AS can make informed routing decisions for these networks.

EIGRP neighbors use the EIGRP external (D EX) listing to route traffic that is destined for the OSPF AS via the boundary router. The boundary router must have the OSPF routes for that destination network in its routing table to be able to forward the traffic.

For this reason, routes must be in the routing table for them to be redistributed. This requirement may seem self-evident, but it can also be a source of confusion. For instance, if a router learns about a network via EIGRP and OSPF, only the EIGRP route is put in the routing table, because it has a lower administrative distance. Suppose RIP is also running on this router, and you want to redistribute OSPF routes into RIP. That network will not be redistributed into RIP, because it is in the routing table as an EIGRP route, not as an OSPF route.

Default Metrics for Redistributed Routes

This topic describes default metrics that are used by different routing protocols when redistribution is implemented and identifies the seed metrics that are used by various routing protocols.

Using Seed Metrics

- The initial, or seed, metric for a route is derived from the directly connected router interface.
- Once a compatible metric is established, the metric will increase in increments just like any other route.

R1#

```
router eigrp 110
 network 10.0.0.0
 redistribute rip
 default-metric 1000 100 250 100 1500
```

- Use the **default-metric** command to establish the seed metric for the route or specify the metric when redistributing.

When a router advertises a link that is directly connected to one of its interfaces, the initial or seed metric—also referred to as the default metric—that is used is derived from the characteristics of that interface, and the metric increases in increments as the routing information is passed to other routers.

For OSPF, the seed metric is based on the bandwidth of the interface. For EIGRP, the seed metric is based on the interface bandwidth and delay. For RIP, the seed metric starts with a hop count of 0 and increases in increments from router to router.

Redistributed routes are not physically connected to a router; they are learned from other routing protocols. It is difficult to translate from one metric to another, from hops to bandwidth. If a boundary router is to redistribute information between routing protocols, it must be able to translate the metric of one routing protocol into the metric of the other routing protocol.

For example, if a boundary router receives a RIP route, the route will have a hop count as its metric. To redistribute the route into OSPF, the router must translate the hop count into a cost metric that the OSPF routers will understand. This seed metric is defined during redistribution configuration. Once the seed metric for a redistributed route is established, the metric will increase in increments normally within the AS.

Note The exception to the rule that metrics increase in normal increments within the AS is for OSPF external type 2 (E2) routes. These routes maintain their initial metric regardless of how far they are propagated across an AS.

The **default-metric** command, which is used in routing process configuration mode, establishes the seed metric for all redistributed routes. Cisco routers also allow the seed metric to be specified as part of the **redistribution** command, either with the *metric* option or by using a route map. Whichever way you set the initial seed metric, you should set it to a value that is larger than the largest metric within the receiving AS, to help prevent suboptimal routing and routing loops.

For more details about the **default-metric** (RIP) command, go to the Cisco IOS IP Routing: RIP Command Reference via the following link

http://www.cisco.com/en/US/docs/ios/iproute_rip/command/reference/irr_book.html

For more details about the **default-metric** (OSPF) command, go to the Cisco IOS IP Routing: OSPF Command Reference via the following link :

http://www.cisco.com/en/US/docs/ios/iproute_ospf/command/reference/iro_book.html

For more details about the **default-metric** (EIGRP) command, go to the Cisco IOS IP Routing: EIGRP Command Reference via the following link :

http://www.cisco.com/en/US/docs/ios/iproute_eigrp/command/reference/ire_book.html

Default Seed Metrics

Protocol	Default Seed Metric
RIP	Infinity
EIGRP	Infinity
OSPF	20 for all except BGP, which is 1
BGP	BGP metric set to IGP metric value

© 2009 Cisco Systems, Inc. All rights reserved.

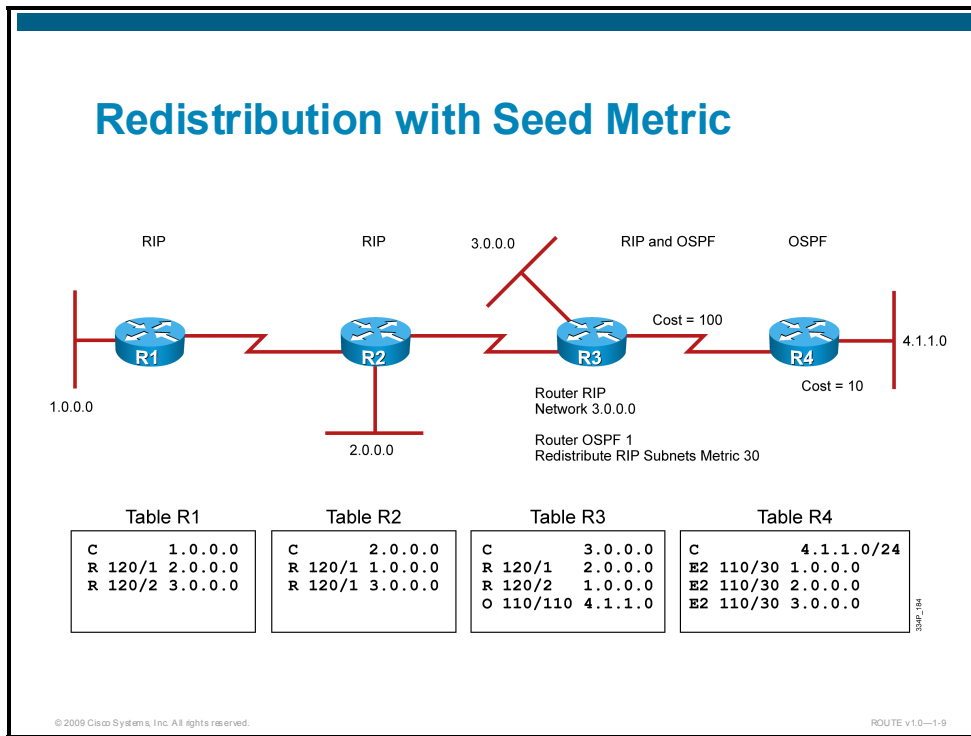
ROUTE v1.0-4-8

The figure shows the default seed metric value for redistributed routes for each IP routing protocol. RIP and EIGRP do not advertise a redistributed route unless a seed metric is configured.

These protocols interpret the seed metric of 0 as infinity by default. A metric of infinity tells the router that the route is unreachable; therefore, it should not be advertised. The default metric for redistributing routes into each of the protocols is as follows:

- The default for RIP and EIGRP is infinity; you must specify a default metric, or routes will not be redistributed.
- For OSPF, the redistributed routes have a default type 2 metric of 20, except for redistributed BGP routes, which have a default type 2 metric of 1.
- For BGP, the redistributed routes maintain the IGP routing metrics.

Redistribution with Seed Metric

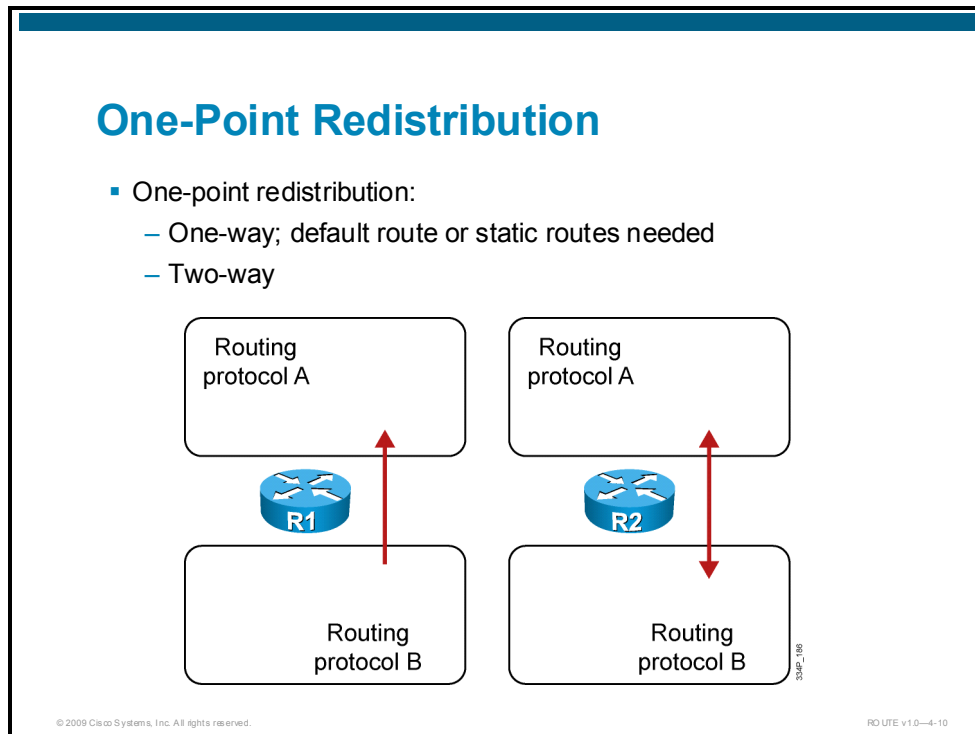


The figure illustrates a seed metric of 30 that is implemented by OSPF on the redistributed RIP routes. The link cost of the serial link to router R4 is 100. So, the cost for networks 1.0.0.0, 2.0.0.0, and 3.0.0.0 in R4 is the seed metric (30) plus the link cost (100)—130. Notice that the metrics of the three networks in the RIP cloud are irrelevant in the OSPF cloud, because the objective is to have each OSPF router forward traffic for the three networks to the border (redistributing) router.

A metric of infinity tells the router that the route is unreachable; therefore, it should not be advertised. When redistributing routes into RIP and EIGRP, you must specify a default metric. For OSPF, the redistributed routes have a default type 2 metric of 20, except for redistributed BGP routes, which have a default type 2 metric of 1. For BGP, the redistributed routes maintain the IGP routing metrics.

Determining Where to Redistribute

This topic describes the optimum redistribution process between two protocols, the process for points of distribution in a network, and ways of identifying possible routing loops.



One-point redistribution defines only one redistribution point between two routing protocols. Routes are redistributed on one router only. The redistribution can be:

- One-way
- Two-way

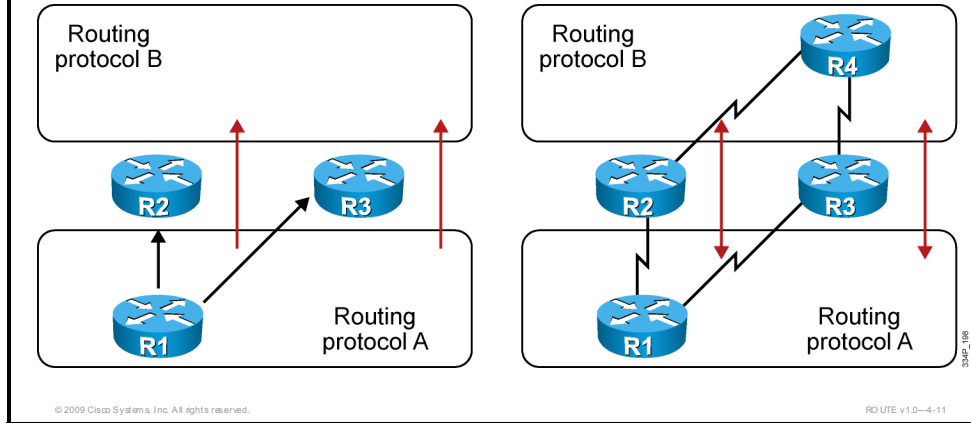
Redistribution is one-way if routes from routing protocol A are redistributed into routing protocol B but not vice versa. Redistribution is two-way if routes from routing protocol A are redistributed into routing protocol B, and routes from routing protocol B are also redistributed into routing protocol A.

One-way redistribution requires the use of a default route or static routes. If routes are redistributed from routing protocol A into routing protocol B, routing protocol B devices are aware of all the routing information. At the same time, devices in the routing protocol A autonomous system are aware of routing information for their own AS only, and reachability for destinations that are outside the routing protocol A autonomous system requires the use of a default route or one or more static routes.

One-way or two-way redistribution at one point is always safe, because one-point redistribution represents the only exit and entrance from one routing protocol to another. Routing loops cannot be inadvertently created.

One-Way and Two-Way Multipoint Redistribution

- Multipoint redistribution:
 - One-way
 - Two-way



Multipoint redistribution is redistribution between two routing protocols that takes place on two or more separate devices that are running both routing protocols. Two possibilities exist:

- Multipoint one-way redistribution
- Multipoint two-way redistribution

Multipoint redistribution is likely to introduce routing loops. Even one-way multipoint redistribution is dangerous, and generic multipoint two-way redistribution is highly problematic. Problems often result from differences in administrative distance between the two protocols, as well as incompatible metrics. Statically assigned metrics are used in redistribution points.

Multipoint one-way redistribution only works well under these circumstances:

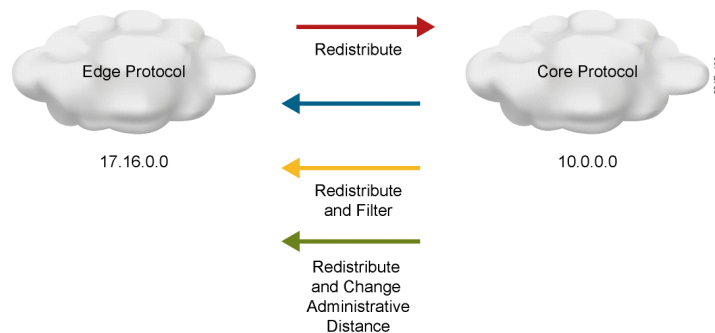
- The receiving routing protocol supports different administrative distances for internal and external routes. Routing protocols that support different administrative distances include EIGRP, BGP, and late maintenance releases of OSPF.
- The external administrative distance of the receiving routing protocol is higher than the administrative distance of the sending routing protocol.

Multipoint two-way redistribution difficulties include the following:

- Suboptimal routing (only part of the total cost is considered in routing decisions)
- Self-sustained routing loops on route loss

In multipoint redistribution scenarios, preventing routing loops is a main concern. The redistribution configuration should insert only internal routes from routing protocol A to B and vice versa. Routes at the redistribution points should be tagged and then filtered based on the tags that are used when doing redistribution in the other direction. Propagation of the metric from A to B and vice versa is recommended, even though it is not sufficient to prevent loops. The easiest way to avoid loops when using two-way redistribution is to use a default route.

Redistribution Techniques



The safest way to perform redistribution is to redistribute routes in only one direction, on only one boundary router within the network. To do this, you must first determine which routing protocol is the core routing protocol, and which ones are edge routing protocols.

The core routing protocol is the main routing protocol that is running in the network. During a transition between routing protocols, the core is the new routing protocol and the edge is the old routing protocol. In networks that run multiple routing protocols all the time, the core is usually the more advanced routing protocol.

If redistribution must be done in both directions, or on multiple boundary routers, the redistribution should be tuned to avoid problems like suboptimal routing and routing loops.

Depending on your network design, you may use any of the following redistribution techniques:

- Redistribute a default route about the core AS into the edge AS. In one-way redistribution, routes from the edge routing protocols are redistributed into the core routing protocol, and a default route is sent back to the edge routers. This technique helps prevent route feedback, suboptimal routing, and routing loops.
- Redistribute multiple static routes about the core AS into the edge AS. The edge routes are still redistributed into the core, but static routes for the core networks are redistributed into the edge protocol and sent to the edge routers. This method works if there is one redistribution point only, but it may cause route feedback if there are multiple points.
- Redistribute routes from the core AS into the edge AS with filtering to block out inappropriate routes. For example, routes from the edge should not be redistributed back into the edge routers from the core via another redistribution point (when there are multiple boundary routers).

- Redistribute all routes from the core AS into the edge AS and from the edge AS into the core AS. Then, modify the administrative distance that is associated with the external routes so that they are not the selected routes when multiple routes exist for the same destination. In some cases, the route that is learned by the native (local) routing protocol is better, but it may have a higher (less believable) administrative distance.

If two routing protocols advertise routes to the same destination, information from the routing protocol with the lowest administrative distance is placed in the routing table. By default, a route that is redistributed into a routing protocol inherits the default administrative distance of that routing protocol.

Caveats of Redistribution

This topic describes how to avoid routing loops in the network that are caused by redistribution and create a distribution and loop map for a given network.

Redistribution Implementation Considerations

- Problems:
 - Routing loop
 - Suboptimal path selection
 - Incompatible routing information
 - Inconsistent convergence time- Solutions:
 - Administrative distance
 - Route maps
 - Distribution lists
 - Manipulation of metrics

© 2009 Cisco Systems, Inc. All rights reserved. R0 UTE v1.0-4-13

Redistribution of routing information adds to the complexity of a network and increases the potential for routing confusion, so it should be used only when necessary. The key issues that arise when you are using redistribution are the following:

- **Routing feedback (routing loops):** Depending on how you employ redistribution, routers may send routing information that is received from one AS back into that same AS. The feedback is similar to the routing loop problem that occurs in distance vector topologies.
- **Incompatible routing information:** Because each routing protocol uses different metrics to determine the best path, path selection using the redistributed route information may be suboptimal. The metric information about a route cannot be translated exactly into a different protocol, so the path that a router chooses may not be the best.

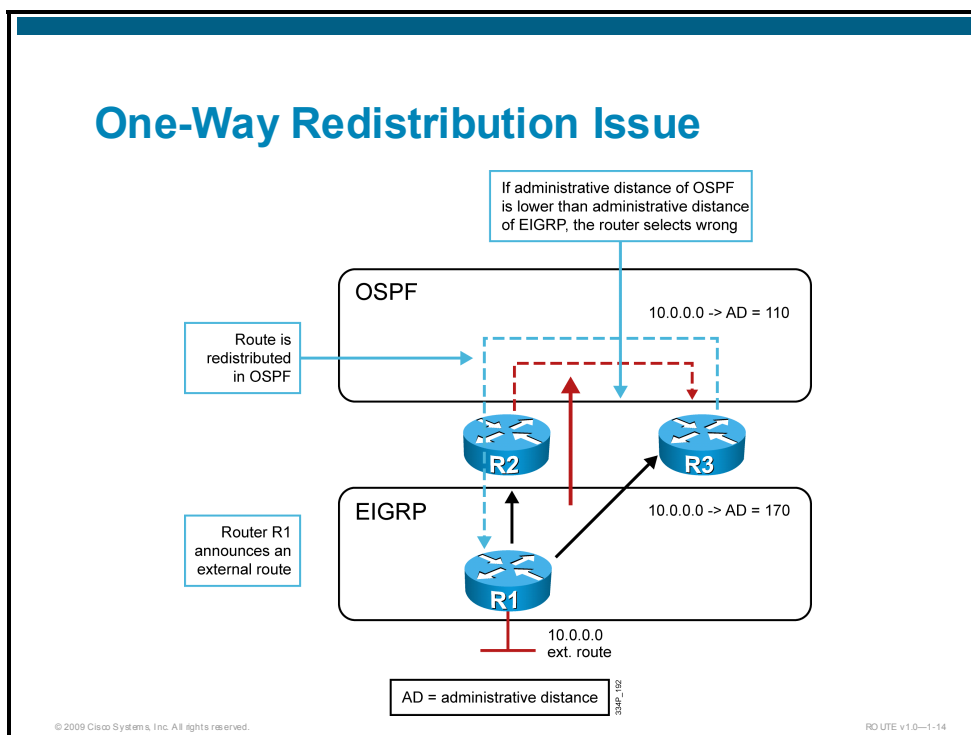
To prevent suboptimal routing, as a rule, you should assign to redistributed routes a seed metric that is higher than any routes that are native to the redistributing protocol. For instance, if RIP routes are being redistributed into OSPF and the highest OSPF metric is 50, the redistributed RIP routes should be assigned an OSPF metric that is higher than 50.

- **Inconsistent convergence time:** Different routing protocols converge at different rates. For example, RIP converges more slowly than EIGRP, so if a link goes down, the EIGRP network will learn about it before the RIP network.

Note Good planning will ensure that these issues do not cause problems in your network.

Good planning can eliminate the majority of issues, but additional configuration might be required. Some issues may be solved by changing the administrative distance, manipulating the metrics, and filtering using route maps and distribute lists.

One-Way Redistribution Issue



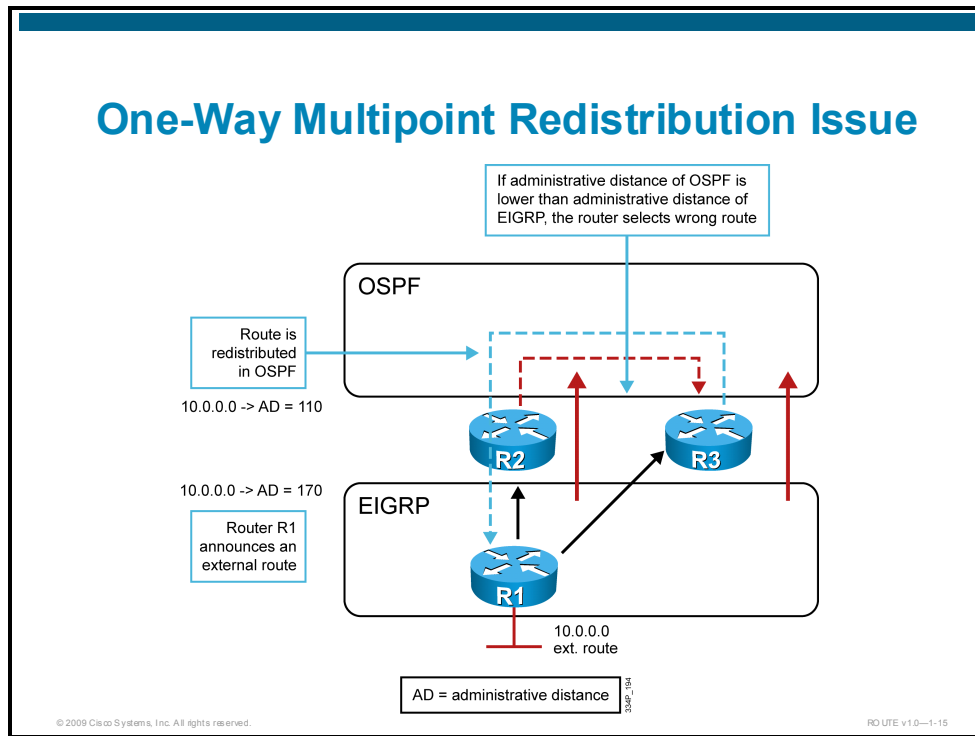
Multipoint one-way redistribution can cause routing confusion if routing protocols are using different administrative distances. It works properly only in the following circumstances:

- If the receiving routing protocol supports different administrative distances for internal and external routes
- If the external administrative distance of one routing protocol (OSPF in the figure) is higher than the administrative distance of the second routing protocol (EIGRP in the figure)

Routing protocols that support different administrative distances include EIGRP, BGP, and OSPF.

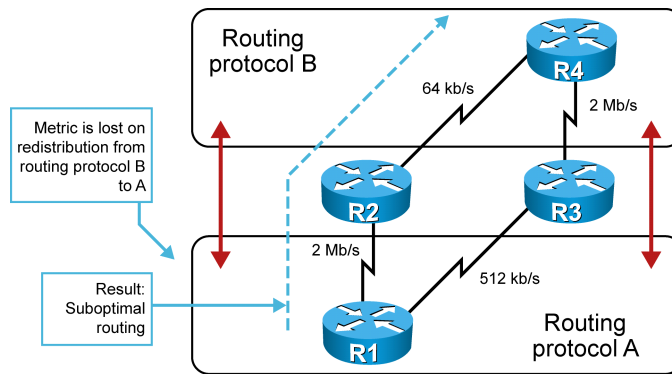
In the figure, R1 from EIGRP is announcing an external route by sending routing updates to R2 and R3. Both neighboring routers are running two routing protocols, and the redistribution between EIGRP and OSPF takes place on R2. R3 receives routing update information for the same route directly from R1 and via R2, which is sending a redistributed route through OSPF updates. The administrative distance of OSPF (110) is lower than the administrative distance of external EIGRP routes (170), so R3 selects the wrong route. Instead of sending packets directly to R1, R3 prefers the path via R2, and the result is suboptimal routing.

One-Way Multipoint Redistribution Issue



In the figure, R1 from EIGRP is announcing routes by sending routing updates to R2 and R3. Both neighboring routers are running two routing protocols, and redistribution between EIGRP and OSPF takes place. Because multipoint redistribution is configured, R2 and R3 receive routing update information for the same route directly from R1 and via OSPF (R2 from R3 and R3 from R2). The routing update information that they receive via OSPF comes in the form of redistributed routes that they receive in OSPF updates. If the administrative distance of OSPF (110) is lower than the administrative distance of EIGRP (170), then R2 also selects the wrong route. Instead of sending packets directly to R1, R2 prefers the path via R3, and the result is suboptimal routing. It can also result in route flapping, because R2 is also redistributing the same EIGRP route in OSPF. R3 also prefers the OSPF path and sends packets via R2. The packet may be bouncing back and forth between R2 and R3 in the OSPF network until its Time to Live (TTL) value reaches 0, before being discarded. Network 10.0.0.0 is then unreachable.

Two-Way Multipoint Redistribution Issue



Generic multipoint two-way redistribution is highly problematic and requires careful design and configuration.

Routing protocols have incompatible metrics, and during redistribution, the metric information can be lost. To avoid problems that are associated with incompatible metrics, statically assigned metrics can be used in redistribution points.

A number of problems can occur during multipoint two-way redistribution:

- Suboptimal routing (only part of the total cost is considered in routing decisions)
- Self-sustained routing loops upon route loss

To prevent routing loops in multipoint redistribution scenarios, you should take into account the following recommendations during configuration:

- Insert only internal routes from routing protocol A to B and vice versa.
- Tag routes in redistribution points, and filter based on tags when doing redistribution.
- Propagate metrics from routing protocol A to routing protocol B properly (though this is not sufficient to prevent loops).
- Use default routes to avoid two-way redistribution.

The figure shows a two-way multipoint redistribution issue in which the cost of the internal links in routing protocol A is completely different from the cost of the links in routing protocol B. It is obvious that the best path between R1 and R4 is via R3, but during redistribution from routing protocol B to routing protocol A, the metric is lost and R1 sends the packets toward R4 via R2. The result is suboptimal routing.

Summary

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

Summary

- Big, complex networks must be adjusted based on factors like political borders, geographical borders, and mergers with other companies. In such complex scenarios, running multiple IP routing protocols is common.
- Using a routing protocol to advertise routes that are learned by some other means, such as by another routing protocol, static routes, or directly connected routes, is called redistribution.
- When a router advertises a link that is directly connected to one of its interfaces, the initial, or seed, metric that is used is derived from the characteristics of that interface. During redistribution, the default metric is used; when redistributing routes into RIP and EIGRP, the default metric must be specified.

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—4-17

Summary (Cont.)

- Redistribution can be one-point (one-way or two-way) or multipoint (one-way or two-way). One-point redistribution, whether one-way or two-way, is always safe, because it represents the only exit from one routing protocol to another. In contrast, multipoint redistribution is likely to introduce routing loops, so careful design and configuration must be applied.
- The key issues that arise when you are using redistribution are routing feedback (routing loops), incompatible routing information, and inconsistent convergence time. Solutions to these issues include correct design in terms of the administrative distance, manipulation of metrics, and filtering using route maps and distribute lists.

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—4-18

Configuring and Verifying Route Redistribution

Overview

Configuring route redistribution can be simple or complex, depending on the mix of routing protocols that you want to redistribute. The commands that are used to enable redistribution and assign metrics vary slightly, depending on the routing protocols that are being redistributed. Before configuring the exchange of routing information between routing protocols, you must understand the procedures for and requirements of each routing protocol.

Redistribution must be configured correctly for each routing protocol to obtain proper results. This lesson describes how to configure route redistribution between various interior gateway protocols (IGPs). The commands for each protocol are covered. In addition, the impact of route redistribution is analyzed.

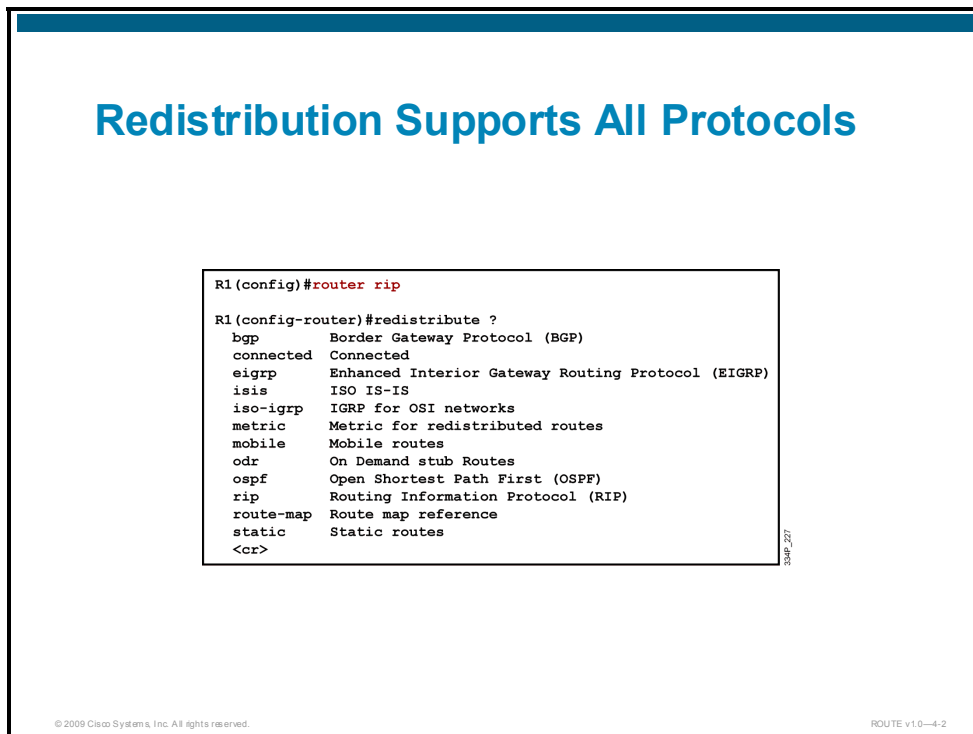
Objectives

Upon completing this lesson, you will be able to configure route redistribution between multiple IP routing protocols. This ability includes being able to meet these objectives:

- Define examples of redistribution
- Define the administrative distance attribute
- Define route redistribution using administrative distance
- Identify the impact of using administrative distance for route manipulation
- Describe route redistribution using route maps

Examples of Redistribution

This topic describes how to redistribute between Routing Information Protocol (RIP), Open Shortest Path First (OSPF), and Enhanced Interior Gateway Routing Protocol (EIGRP), as well as the procedures that are necessary to configure route redistribution.



As shown in the figure, redistribution supports all routing protocols. Additionally, static and connected routes can be redistributed to allow the routing protocol to advertise the routes without using a network statement for them.

Routes are redistributed into a routing protocol, so the **redistribute** command is entered under the routing process that is to receive the routes. Before implementing redistribution, consider these points:

- Only protocols that support the same protocol stack are redistributed. For example, you can redistribute between IP RIP and OSPF, because they both support the TCP/IP stack.
- The method that is used to configure redistribution varies slightly among different routing protocols and combinations of routing protocols. Some routing protocols require a metric to be configured during redistribution, but others do not.

Note As of Cisco IOS Release 12.3, Interior Gateway Routing Protocol (IGRP) is no longer supported.

The following generic steps apply to all routing protocol combinations; however, the commands that are used to implement these steps may vary. For configuration commands, it is important that you review the Cisco IOS Software documentation for the specific routing protocols that need to be redistributed.

Note In this topic, the terms “core” and “edge” are generic terms that are used to simplify the discussion about redistribution.

1. Locate the boundary router that requires configuration of redistribution. Selecting a single router for redistribution minimizes the likelihood of creating routing loops that are caused by feedback.
2. Determine which routing protocol is the core, or backbone, protocol. Typically, this protocol is OSPF or EIGRP.
3. Determine which routing protocol is the edge, or short-term (in the case of migration), protocol. Determine whether all routes from the edge protocol need to be propagated into the core. Consider methods that reduce the number of routes.
4. Select a method for injecting the required edge protocol routes into the core. Simple redistribution using summaries at network boundaries minimizes the number of new entries in the routing table of the core routers.

When you have planned the edge-to-core redistribution, consider how to inject the core routing information into the edge protocol. Your choice depends on your network.

For more details about the **redistribute** (IP) command, go to the Cisco IOS IP Routing: Protocol-Independent Command Reference via the following link:

http://www.cisco.com/en/US/docs/ios/iproute_pi/command/reference/iri_book.html

Steps to Configure Redistribution into RIP

Enter router RIP configuration mode

```
R1(config-router)#
```

```
redistribute ospf 1
```

- Configure redistribution from another routing protocol; OSPF is used here.

```
R1(config-router)# redistribute ospf 1 ?
match      Redistribution of OSPF routes
metric     Metric for redistributed routes
route-map  Route map reference
...
<cr>
```

- These optional parameters influence redistribution into the RIP routing protocol.

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—4-3

The figure shows how to configure for redistribution from OSPF process 1 into RIP.

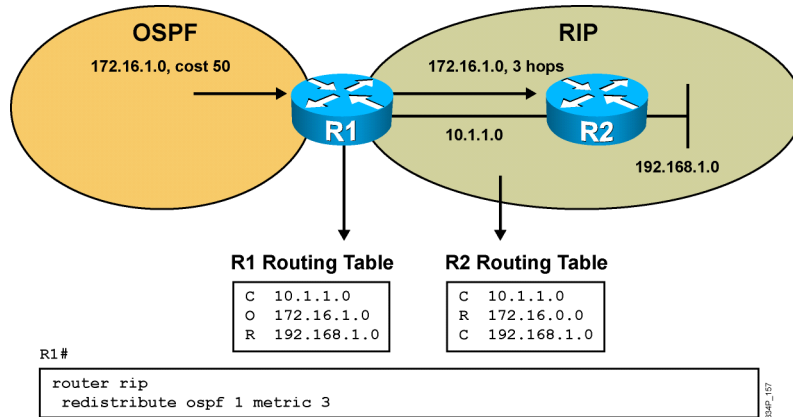
The first step is to enter router RIP configuration mode. The **router rip** command is used to access the routing process into which routes need to be redistributed. In this case, it is the RIP routing process.

The next step is to use the **redistribute (IP)** command inside router configuration mode to specify the routing protocol that is to be redistributed into RIP. In this case, it is the OSPF routing process number 1. You can use optional keywords to change how distribution is performed. For example, you might modify the default metric or route filtering using route maps.

Note The default metric is infinity, except when you are redistributing a static or connected route. In that case, the default metric is 1.

Redistributing into RIP

- Redistribute into RIP and assign the metric 3 (3 hops) to all OSPF routes.



In the figure, routes from OSPF process 1 are being redistributed into RIP and given a seed metric of 3. The default metric is infinity, and a network administrator was required to change it. Because no route type is specified, both internal and external OSPF routes are redistributed into RIP. In the routing table on router R2, you can see that OSPF route 172.16.1.0, which is redistributed into RIP on R1, is announced as an RIP route in the RIP autonomous system (AS).

Steps to Configure Redistribution into OSPF

Enter router OSPF configuration mode

```
R1(config-router)#
```

```
redistribute eigrp 100 metric 100 subnets
```

- Configure redistribution from another routing protocol; EIGRP with AS 100 is used here.

```
R1(config-router)# redistribute eigrp 100 ?
metric          Metric for redistributed routes
metric-type     OSPF/IS-IS exterior metric type for
                redistributed routes
route-map       Route map reference
subnets       Consider subnets for redistribution into OSPF
tag            Set tag for routes redistributed into OSPF
...
<cr>
```

- These optional parameters influence redistribution into OSPF routing protocol.

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-4-5

The figure shows how to configure for redistribution from EIGRP AS 100 into an OSPF routing process.

The first step is to enter router OSPF configuration mode. The **router ospf 1** command is used to access the OSPF routing process into which the routes need to be redistributed. In this case, it is OSPF process 1.

The next step is to use the **redistribute (IP)** command in router configuration mode to specify the routing protocol that is to be redistributed into OSPF. In this case, it is the EIGRP routing process for AS 100. You can use optional keywords to change the way distribution is performed. For example, you might modify the default metric or route filtering using route maps. Redistribution into OSPF can also be limited to a defined number of prefixes using the **redistribute maximum-prefix** router configuration command. The threshold parameter will default to logging a warning at 75 percent of the defined maximum value that is configured. After the router reaches the defined maximum number of prefixes, it will not redistribute any more routes. Additional parameters can be used to avoid the limitation that is placed on redistribution; the maximum value number simply becomes a second point at which a warning message is logged. This command was introduced in Cisco IOS Release 12.0(25)S. It was integrated into Cisco IOS Release 12.2(18)S and Cisco IOS Release 12.3(4)T and later.

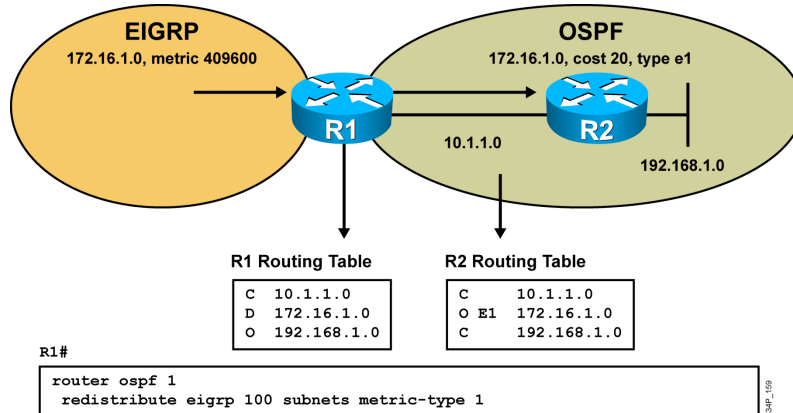
Note The default metric is 20, and the default metric type is 2. Subnets do not redistribute by default.

For more details about the **redistribute maximum-prefix** command, go to the Cisco IOS IP Command Reference, Volume 2 of 4: Routing Protocols via the following link:

http://www.cisco.com/en/US/docs/ios/12_3t/ip_route/command/reference/ip2_o1gt.html

Redistributing into OSPF

- Redistribute from EIGRP AS 100 into OSPF and change the metric type from type 2 to type 1.



In the figure, routes from EIGRP AS 100 that are being redistributed into OSPF process 1 are given a default metric of 20. The default metric type is external type 2 (E2); in the figure, the default metric type was changed to external type 1 (E1). This setting means that the metric increases in increments whenever updates are passed through the network. The command contains the **subnets** option, so subnets are redistributed as well.

In the routing table on R2, you can see that EIGRP route 172.16.1.0, which is redistributed into OSPF on R1, is announced as an OSPF route inside the OSPF AS. It is also apparent that the route is an E1 type.

Steps to Configure Redistribution into EIGRP

- Enter router EIGRP configuration mode

```
R1(config-router)#
```

```
redistribute ospf 1
```

- Configure redistribution from another routing protocol; OSPF with process number 1 is used here.

```
R1(config-router)#redistribute ospf 1 ?
match      Redistribution of OSPF routes
metric     Metric for redistributed routes
route-map  Route map reference
...
<CR>
```

- These optional parameters influence redistribution into EIGRP routing protocol.

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—4-7

The figure shows how to configure redistribution from an OSPF into an EIGRP AS 100 routing process.

The first step is to enter router EIGRP configuration mode. The **router eigrp 100** command is used to access the EIGRP routing process into which routes need to be redistributed. In this case, it is EIGRP AS 100.

The next step is to use the **redistribute (IP)** command inside router configuration mode to specify the routing protocol that is to be redistributed into EIGRP AS 100. In this case, it is OSPF process 1. You can use optional keywords to change the way that distribution is performed. For example, you might modify default metric or route filtering using route maps.

Note The default metric is infinity.

When you are redistributing a static or connected route into EIGRP, the default metric is equal to the metric of the associated interface.

Default Metric

```
R1 (config-router) #
```

```
default-metric 10000 100 255 1 1500
```

- The default metric is set to: Bandwidth in kb/s = 10000, Delay in tens of microseconds = 100, Reliability = 255 (maximum), Load = 1 (minimum), MTU = 1500 bytes
- You need to use this command when redistributing from another protocol with incompatible metric into EIGRP.
- You do not need this command when redistributing:
 - Static routes to interface or connected networks
 - Between EIGRP processes

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v10-4-8

You do not need default metrics to redistribute EIGRP into itself. To set metrics for EIGRP, use this form of the **default-metric** (EIGRP) router configuration command. EIGRP needs five metrics when redistributing other protocols: bandwidth, delay, reliability, load, and maximum transmission unit (MTU).

- **Bandwidth:** The minimum bandwidth of the route in kilobits per second. It can be any positive integer.
- **Delay:** The route delay in tens of microseconds. It can be 0 or any positive number that is a multiple of 39.1 ns.
- **Reliability:** The likelihood of successful packet transmission, expressed as a number between 0 and 255. The value 255 means 100 percent reliability; 0 means no reliability.
- **Loading:** The effective bandwidth of the route, expressed as a number from 1 to 255 (255 is 100 percent loading).
- **MTU:** The minimum MTU size of the route in bytes. It can be any positive integer.

You must use the **default-metric** (EIGRP) router configuration command when redistributing from another protocol with an incompatible metric into EIGRP, unless you use the **metric** parameter of the **redistribute** command. The **redistribute** command is able to set the metric that is used during the redistribution.

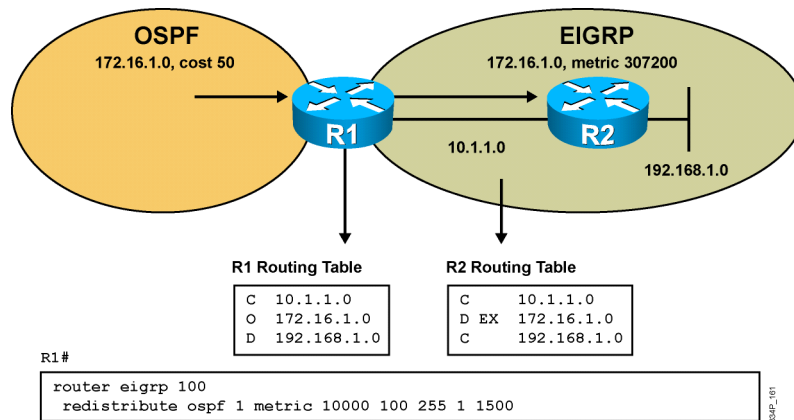
You do not need to use the **default-metric** (EIGRP) router configuration command when redistributing static routes to interface or connected networks between EIGRP processes.

For more details about the **default-metric** (EIGRP) command, go to the Cisco IOS IP Routing: EIGRP Command Reference via the following link:

https://www.cisco.com/en/US/docs/ios/iproute_eigrp/command/reference/ire_book.html

Redistributing into EIGRP

- Redistribute OSPF 1 routing process updates into EIGRP AS 100 and overwrite the default metric.

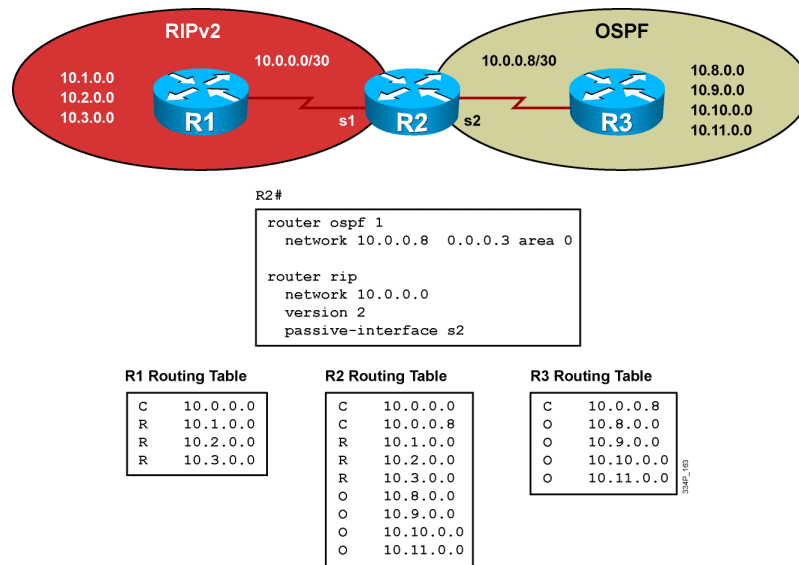


In the figure, routes from OSPF process 1 are redistributed into EIGRP AS 100. The default metric is infinity, but to reflect the normal metric, it must be changed. You can use the **default-metric** (EIGRP) command or change the metric and use the **redistribution** (IP) command. In the figure, the default metric has been changed with the **redistribution** command as follows: bandwidth in kilobits per second = 10000, delay in tens of microseconds = 100, reliability = 255 (maximum), load = 1 (minimum), MTU = 1500 bytes.

You can see from the routing table on R2 that OSPF route 172.16.1.0, which is redistributed into EIGRP AS 100 on R1, is announced as an EIGRP route inside EIGRP AS 100. The redistributed route appears as an external EIGRP (D EX) route.

External EIGRP (D EX) routes have a higher administrative distance than internal EIGRP (D) routes, so internal EIGRP routes are preferred over external EIGRP routes.

Example: Redistributing into OSPF— Before Redistribution



The figure shows the network of a hypothetical company. The network begins with two routing domains, or autonomous systems—one using OSPF and one using RIP version 2 (RIPv2). R2 is the boundary router; it connects directly to one router within each routing domain and runs both protocols.

R1 is in the RIP domain and is advertising subnets 10.1.0.0, 10.2.0.0, and 10.3.0.0 to R2. R3 is in the OSPF domain and is advertising subnets 10.8.0.0, 10.9.0.0, 10.10.0.0, and 10.11.0.0 to R2.

The configuration of R2 is shown in the figure. RIP is required to run on the serial 1 interface only; therefore, the **passive-interface** command is given for the serial 2 interface. The **passive-interface** command prevents RIP from sending route advertisements out that interface. OSPF is configured on serial 2.


The figure also shows the routing tables of R1, R2, and R3. Each routing domain is separate, and routers within them only recognize routes that are communicated from their own routing protocols.

The only router with information on all the routes is R2.

The goal of redistribution in this network is for all routers to recognize all routes within the company. To accomplish this goal, redistribution is planned:

- Redistribute RIP routes into OSPF.
- Redistribute OSPF routes into the RIP domain.

Example: Configuring Redistribution at Router R2



```
R2#
router ospf 1
 network 10.0.0.8 0.0.0.3 area 0
 redistribute rip subnets metric 300

router rip
 network 10.0.0.0
 version 2
 passive-interface s2
 redistribute ospf 1 metric 5
```

R1 Routing Table

C	10.0.0.0
R	10.0.0.8
R	10.1.0.0
R	10.2.0.0
R	10.3.0.0
R	10.8.0.0
R	10.9.0.0
R	10.10.0.0
R	10.11.0.0

R2 Routing Table

C	10.0.0.0
C	10.0.0.8
R	10.1.0.0
R	10.2.0.0
R	10.3.0.0
O	10.8.0.0
O	10.9.0.0
O	10.10.0.0
O	10.11.0.0

R3 Routing Table

C	10.0.0.8
O E2	10.1.0.0
O E2	10.2.0.0
O E2	10.3.0.0
O	10.8.0.0
O	10.9.0.0
O	10.10.0.0
O	10.11.0.0

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-1-11

R2 is the boundary router, so redistribution is configured on it. The figure shows how R2 is configured to accomplish the required redistribution.

RIP is redistributed under the OSPF process. In this example, the metric is set under the **redistribute** command. Other options include specifying a default metric or accepting the OSPF default metric of 20.

The **default-metric** command assigns a seed metric to all routes that are redistributed into OSPF from any origins. If a metric value is configured under a specific **redistribute** command, this value overrides the default metric value. A value of 300 is selected, because it is a worse metric than any of the native OSPF routes.

Under the RIP process, routes are redistributed in from OSPF process 1. These routes are redistributed into RIP with a metric of 5. A value of 5 is selected, because it is higher than any metric in the RIP network.

The figure also shows the routing tables of all three routers after redistribution is completed. The goal is accomplished. All routers now have routes to all remote subnets. There is complete reachability within the entire network.

R1 and R3 now have many more routes to keep track of than before. Each router is also affected by topology changes in the routing domain of the other router.

Depending on network requirements, you may be able to increase efficiency by summarizing the routes before redistributing them. Remember that route summarization hides information.

If routers in the other autonomous systems are required to track topology changes within the network, route summarization should not be performed, because it hides information that the routers need.

A more typical case is that the routers need to recognize topology changes only within their own routing domains. In this case, performing route summarization is appropriate.

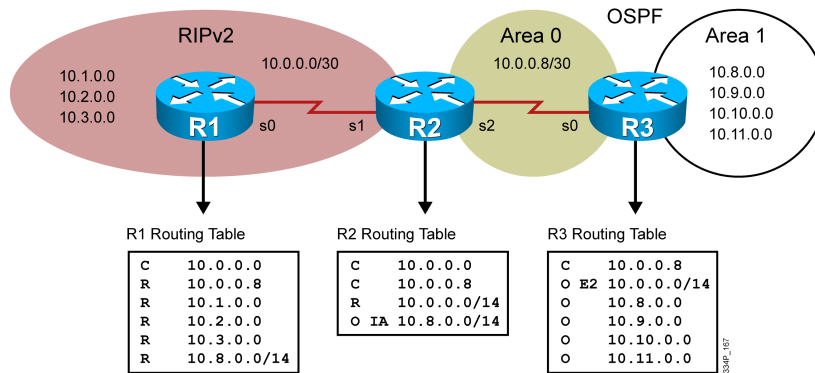
Example: Routing Tables After Summarizing Routes and Redistributions

R1#

```
interface s0
ip summary-address rip 10.0.0.0 255.252.0.0
```

R3#

```
router ospf 1
area 1 range 10.8.0.0 255.252.0.0
```



If routes are summarized before redistribution, the routing tables of each router are significantly smaller. R2 benefits the most; it now has only four routes to keep track of instead of nine. R1 has five routes instead of eight, and R3 has six routes to keep track of instead of eight.

The following list describes how summarization should be configured in this network:

- **R1, RIP:** For RIPv2, the summarization command is given at the interface connecting R2 with R1. This summary address is advertised out of that interface instead of the individual subnets. One limitation of RIP is that the subnet mask of the summary address must be greater than or equal to the default mask for the major classful network. The **ip summary-address rip** interface command is used for summarization.

Note This summary includes 10.0.0.0, which is acceptable in this case because it is directly connected with a longer mask.

- **R3, OSPF:** You must perform summarization in OSPF at an Area Border Router (ABR) or an Autonomous System Boundary Router (ASBR). Create another OSPF area that includes the four subnets that are to be summarized. Give the command for summarization under the OSPF process at R3, which becomes an ABR. The **area range** router configuration command is used for summarization, as shown in the figure.

The Administrative Distance Attribute

Cisco routers use administrative distance when using more than one routing protocol. This topic describes the features of administrative distance in terms of routing protocols.

Administrative Distance

- Administrative distance is a way of ranking the trustworthiness of routing information.
- A lower administrative distance means a route is more trustworthy.

Route Source	Default Distance
Connected interface	0
Static route	1
EIGRP summary route	5
External BGP	20
Internal EIGRP	90
IGRP	100
OSPF	110
IS-IS	115
RIPv1, RIPv2	120
ODR	160
External EIGRP	170
Internal BGP	200
Unknown	255

© 2009 Cisco Systems, Inc. All rights reserved. PD UTE v1.0-4-13

Most routing protocols have metric structures and algorithms that are not compatible with other protocols. It is critical for a network using multiple routing protocols to have seamless exchange of route information and the ability to select the best path across multiple protocols.

Cisco routers use a value called administrative distance to select the best path when they learn two or more routes to the same destination from different routing protocols. Administrative distance is a way of ranking the trustworthiness of a routing protocol. Cisco has assigned a default administrative distance value to each routing protocol that is supported on its routers.

Each routing protocol is prioritized from most believable to least believable. Some examples of prioritization are as follows:

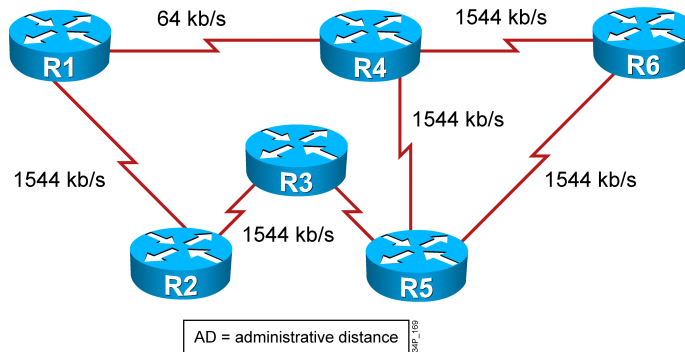
- Prefer manually configured routes (static routes) to dynamically learned routes
- Prefer protocols with sophisticated metrics to protocols with more deterministic metrics
- Prefer External Border Gateway Protocol (EBGP) to most other dynamic protocols

The table in the figure lists the default administrative distances of the protocols that are supported by Cisco routers. The administrative distance is a value between 0 and 255. The lower the administrative distance value, the higher the reliability of the protocol.

Note As of Cisco IOS Release 12.3, IGRP is no longer supported.

Administrative Distance Example

- From R1 to R6:
 - RIP (AD 120) will choose R1-R4-R6
 - OSPF (AD 110) will choose R1-R2-R3-R5-R6
 - EIGRP (AD 90) will choose R1-R2-R3-R5-R6



The figure has multiple sources of routing information that are active at the same time. A lower administrative distance value means that the source of routing information is more believable.

For the best path from R1 to R6, the following route choices are made if all four routing protocols are active on all the routers:

- RIP (administrative distance 120) chooses R1 to R4 to R6 based on hop count (two hops versus four hops the other way).
- OSPF (administrative distance 110) calculates the default metric as 100 Mb/s divided by the interface bandwidth, where bandwidth is the speed of each link.

The path R1 to R4 to R6 default metric is $(100 \text{ Mb/s} / 64 \text{ kb/s}) + (100 \text{ Mb/s} / 1.544 \text{ Mb/s}) = 1562 + 64$, or 1626. The R1 to R2 to R3 to R5 to R6 path default metric is $1544 + 1544 + 1544 + 1544 = 6176$; therefore, OSPF chooses the R1 to R2 to R3 to R5 to R6 path.

Although OSPF and Intermediate System-to-Intermediate System (IS-IS) are both link-state routing protocols that converge quickly, OSPF is more trustworthy than IS-IS. This is because OSPF bases its default metric on bandwidth and is therefore more likely to pick the best path.

- EIGRP (administrative distance 90) calculates the default metric as $(100,000,000 / \text{bandwidth}) + \text{delay}$, where bandwidth is the speed of the slowest link in the path and delay is cumulative across the path.

Assuming a uniform link delay of 100, the R1 to R4 to R6 default metric is $(10^8 / 64) + 200 = 1,562,700$, and the R1 to R2 to R3 to R5 to R6 default metric is $(10^8 / 1544) + 400 = 65,166$.

EIGRP chooses R1 to R2 to R3 to R5 to R6. Although EIGRP and OSPF routing protocols both converge quickly and consider bandwidth, EIGRP is more trustworthy than OSPF. This is because EIGRP takes more information into account in its calculation.

Because EIGRP has the lowest administrative distance of the four protocols, only the EIGRP path is put into the routing table.

Route Redistribution Using Administrative Distance

This topic describes how to manipulate redistributed routes to avoid network loops and instability by using the administrative distance attribute.

Steps to Configure Redistribution Using Administrative Distance

- Enter router configuration mode
- Configure the redistribution from another routing protocol
- Modify the administrative distance
 - Remember that the configuration is different for different routing protocols

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-4-15

Redistribution using administrative distance requires the following steps:

- Step 1** Enter router configuration mode.
- Step 2** Configure redistribution from another routing protocol.
- Step 3** Modify the administrative distance.

Note Keep in mind that the configuration varies for different routing protocols.

These steps need information about the protocols that are used and their administrative distances.

Modifying Administrative Distance

```
R1(config-router)#
```

```
distance ospf external 100 inter-area 100 intra-area 100
```

- Used for OSPF to define the external, interarea, and intra-area distance

```
R1(config-router)#
```

```
distance eigrp 90 170
```

- Used for EIGRP to define the internal and external distance

```
R1(config-router)#
```

```
distance bgp 20 200 200
```

- Used for BGP to define external, internal, and local distance

© 2009 Cisco Systems, Inc. All rights reserved.

RO UTE v1.0-4-16

The **distance** (IP) command configures the administrative distance per specific network that is inserted into routing table, regardless of the IP routing protocol that is used. The specific, protocol-dependent **distance** commands are used to define the administrative distance for each routing protocol.

For OSPF, use the **distance ospf** command. OSPF assigns different administrative distance values to external, interarea, and intra-area routes. The default value is 110 for each of them.

For EIGRP, use the **distance eigrp** command. EIGRP assigns different administrative distance values to routes that are learned natively through EIGRP and to routes that are redistributed from other sources. By default, natively learned routes have an administrative distance of 90, but external routes have an administrative distance of 170.

For BGP, use the **distance bgp** command. BGP assigns different administrative distance values to routes that are learned through EBGP, routes that are learned through IBGP, and local routes. The default values that are used are 20, 200, and 200, respectively.

For more details about the **distance** (IP) command, go to the Cisco IOS IP Routing: Protocol-Independent Command Reference via the following link:

http://www.cisco.com/en/US/docs/ios/iproute_pi/command/reference/iri_book.html

Modifying Administrative Distance (Cont.)

- The **distance eigrp** command is used for EIGRP protocol.
- The **distance (IP)** command is protocol-independent.
 - Used to define an administrative distance for routes that are inserted into the routing table
 - Used for all protocols, except for EIGRP, OSPF, and BGP when the dedicated, protocol-specific **distance** commands are used.

R1#

```
router eigrp 100
network 192.168.7.0
network 172.16.0.0
distance eigrp 80 130
distance 90 192.168.7.0 0.0.0.255
distance 120 172.16.1.3 0.0.0.255
```

304P_208

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-4-17

In some cases, a router selects a suboptimal path if it assigns a better administrative distance to a routing protocol with a worse route.

Assigning a higher administrative distance to an undesired routing protocol ensures that routers select routes from the desired routing protocol. The figure illustrates the commands for changing the default administrative distance.

The **distance** command changes the default administrative distance for all protocols. The **distance** commands that are dedicated to EIGRP, OSPF, and BGP redistribution can also be used.

The **distance** router configuration command is used to specify the administrative distance for particular routes that are inserted into a routing table.

In the figure, the **router eigrp** global configuration command sets up EIGRP routing in AS 100. The **network** router configuration commands specify EIGRP routing on networks 192.168.7.0 and 172.16.0.0. The **distance eigrp 80 130** command sets the administrative distance for internal EIGRP routes to 80 and for external EIGRP routes to 130. Additionally, the **distance 90 192.168.7.0 0.0.0.255** command sets the administrative distance to 90 for all routers on the Class C network 192.168.7.0, and the **distance 120 172.16.1.3 0.0.0.255** command sets the administrative distance to 120 for the router with the address 172.16.1.3.

For more details about the **distance (IP)** command, go to the Cisco IOS IP Routing: Protocol-Independent Command Reference via the following link:

http://www.cisco.com/en/US/docs/ios/iproute_pi/command/reference/iri_book.html

For more details about the **distance eigrp** command, go to the Cisco IOS IP Routing: EIGRP Command Reference via the following link:

http://www.cisco.com/en/US/docs/ios/iproute_eigrp/command/reference/ire_book.html

Modifying Administrative Distance (Cont.)

- The **distance ospf** command is used for OSPF protocol.
- The **distance (IP)** command is protocol-independent.
 - Used to define an administrative distance for routes that are inserted into the routing table.
 - Used for all protocols, except for EIGRP, OSPF, and BGP when the dedicated, protocol-specific **distance** commands are used.

R1#

```
router ospf 10
network 192.168.7.0 0.0.0.255 area 0
network 172.16.0.0 0.0.0.255 area 3
distance ospf external 100 inter-area 100 intra-area 100
distance 90 10.0.0.0 0.0.0.255
distance 110 10.11.0.0 0.0.0.255
distance 130 10.11.12.0 0.0.0.255
```

334P_201

© 2009 Cisco Systems, Inc. All rights reserved.

RO UTE v1.0—1-18

In the figure, the **router ospf** global configuration command sets up OSPF routing with process 10. The **network** router configuration commands specify OSPF routing on networks 192.168.7.0 and 172.16.0.0. The **distance ospf external 100 inter-area 100 intra-area 100** command sets the administrative distance for internal, interarea, and intra-area OSPF routes to 100 (default values are 110). Additionally, the **distance 90 10.0.0.0 0.0.0.255**, **distance 110 10.11.0.0 0.0.0.255**, and **distance 130 10.11.12.0 0.0.0.255** commands set the administrative distance to 90, 110, and 130 for the least to the most specific networks, respectively.

For more details about the **distance (IP)** command, go to the Cisco IOS IP Routing: Protocol-Independent Command Reference via the following link:

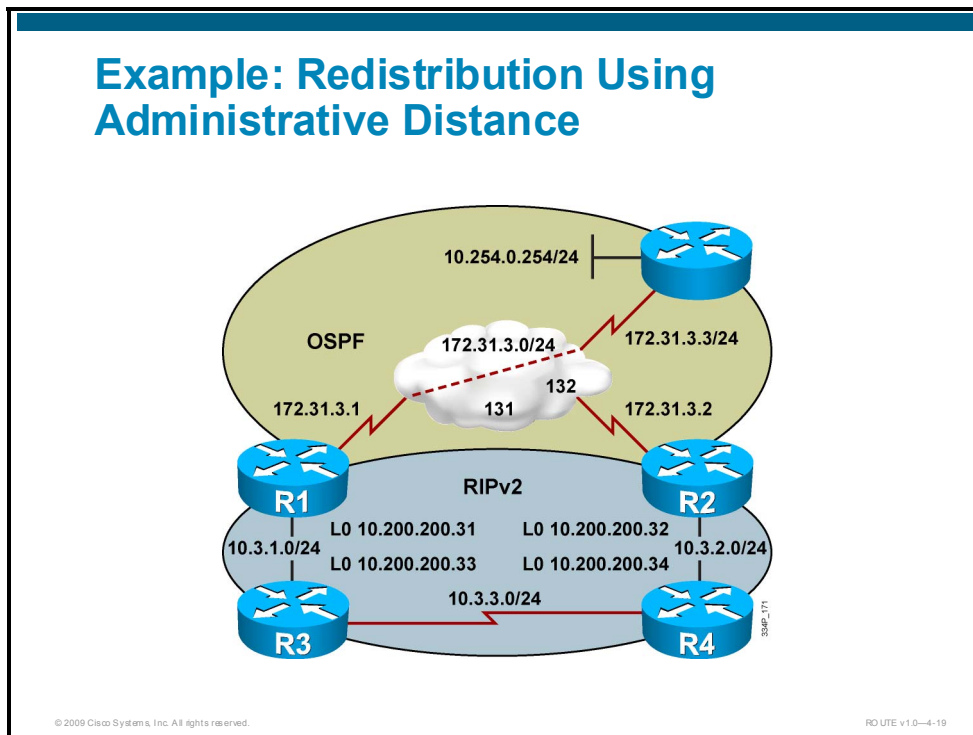
http://www.cisco.com/en/US/docs/ios/iproute_pi/command/reference/iri_book.html

For more details about the **distance ospf** command, go to the Cisco IOS IP Routing: OSPF Command Reference via the following link:

http://www.cisco.com/en/US/docs/ios/iproute_ospf/command/reference/iro_book.html

Impact of Using Administrative Distance for Route Manipulation

This topic describes the before and after results of using the administrative distance to modify route redistribution. Assess the impact of administrative distance changes on routing tables.



The figure shows a network that is using multiple routing protocols. There are several ways to correct path selection problems in a redistribution environment. These examples show how problems occur, where they occur, and one possible way to resolve them. Redistribution using administrative distance is presented.

The figure illustrates a network with RIP and OSPF routing domains. Recall that OSPF is more believable than RIP, because OSPF has an administrative distance of 110 and RIP has an administrative distance of 120.

If, for example, a boundary router (R1 or R2) learns about network 10.3.3.0 via RIPv2 and also via OSPF, the OSPF route will be used and inserted into the routing table, because OSPF has a lower administrative distance than RIPv2, even though the path via OSPF might be the longer (worse) path.

Example: Redistribution Using Administrative Distance (Cont.)

R1#

```
router ospf 1
redistribute rip metric 10000 metric-type 1 subnets
network 172.31.0.0 0.0.255.255 area 0
!
router rip
version 2
redistribute ospf 1 metric 5
network 10.0.0.0
no auto-summary
```

R2#

```
router ospf 1
redistribute rip metric 10000 metric-type 1 subnets
network 172.31.3.2 0.0.0.0 area 0
!
router rip
version 2
redistribute ospf 1 metric 5
network 10.0.0.0
no auto-summary
```

3849_173

© 2009 Cisco Systems, Inc. All rights reserved.

RO UTE v1.0—4-20


The figure illustrates the configurations for R1 and R2. These configurations redistribute RIP into OSPF and OSPF into RIP on both routers.

The redistribution into OSPF sets a default OSPF metric of 10000 to make these routes less preferred than native OSPF routes and protect against route feedback. The redistribute statement also sets the metric type to E1, so that the route metrics continue to accrue, and the router redistributes subnet information.

The redistribution into RIP sets a default RIP metric of 5 to also protect against route feedback.

Example: Redistribution Using Administrative Distance (Cont.)

With OSPF and RIP running



```
R2#show ip route
<Output Omitted>

Gateway of last resort is not set

  172.31.0.0/24 is subnetted, 7 subnets
O    172.31.55.0 [110/2343] via 172.31.3.3, 00:09:46, Serial0/0/0
C    172.31.3.0 is directly connected, Serial0/0/0
O    172.31.2.0 [110/1562] via 172.31.3.3, 00:09:46, Serial0/0/0
10.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
O E1 10.3.1.0/24 [110/10781] via 172.31.3.1, 00:09:47, Serial0/0/0
O E1 10.3.3.0/24 [110/10781] via 172.31.3.1, 00:04:51, Serial0/0/0
C    10.3.2.0/24 is directly connected, fastethernet0/0
O E1 10.200.200.31/32 [110/10781] via 172.31.3.1, 00:09:48, Serial0/0/0
O E1 10.200.200.34/32 [110/10781] via 172.31.3.1, 00:04:52, Serial0/0/0
C    10.200.200.32/32 is directly connected, Loopback0
O E1 10.200.200.33/32 [110/10781] via 172.31.3.1, 00:04:52, Serial0/0/0
O E2 10.254.0.0/24 [110/50] via 172.31.3.3, 00:09:48, Serial0/0/0
```

■ R2 includes suboptimal paths and loops.

The figure displays the routing table on R2 after redistribution has occurred. R2 learned RIP and OSPF routes, but it lists only OSPF routes in the routing table.

The first edge router to set up redistribution has a normal routing table and retains the RIP routes. The second edge router chooses the OSPF routes over its RIP routes. The paths to the internal RIP routes are shown as going through the core because of the dual mutual redistribution points.

OSPF is informed about the RIP routes via redistribution. OSPF then advertises the RIP routes via OSPF routes to its neighboring router. The neighbor router is also informed about the same routes via RIP. However, OSPF has a better administrative distance than RIP, so the RIP routes are not put into the routing table.

OSPF was configured on R1 first, and R2 then received information about the internal (native RIP) routes from both OSPF and RIP. It prefers the OSPF routes, because OSPF has a lower administrative distance. Therefore, none of the RIP routes appear in the table.

Refer back to the topology diagram to trace some of the routes. The redistribution has resulted in suboptimal paths to many of the networks.

For instance, 10.200.200.34 is a loopback interface on R4. R4 is directly attached to R2. However, the OSPF path to that loopback interface goes through R1, then R3, and then R4 before it reaches its destination. The OSPF path that is taken is actually a longer (worse) path than the more direct RIP path.

Example: Redistribution Using Administrative Distance (Cont.)

```
R1#
router ospf 1
 redistribute rip metric 10000 metric-type 1
 subnets
 network 172.31.0.0 0.0.255.255 area 0
 distance 125 0.0.0.0 255.255.255.255 64
 !
router rip
 version 2
 redistribute ospf 1 metric 5
 network 10.0.0.0
 no auto-summary
 !
access-list 64 permit 10.3.1.0 0.0.0.255
access-list 64 permit 10.3.3.0 0.0.0.255
access-list 64 permit 10.3.2.0 0.0.0.255
access-list 64 permit 10.200.200.31
access-list 64 permit 10.200.200.34
access-list 64 permit 10.200.200.32
access-list 64 permit 10.200.200.33

R2#
router ospf 1
 redistribute rip metric 10000 metric-type 1
 subnets
 network 172.31.3.2 0.0.0.0 area 0
 distance 125 0.0.0.0 255.255.255.255 64
 !
router rip
 version 2
 redistribute ospf 1 metric 5
 network 10.0.0.0
 no auto-summary
 !
access-list 64 permit 10.3.1.0 0.0.0.255
access-list 64 permit 10.3.3.0 0.0.0.255
access-list 64 permit 10.3.2.0 0.0.0.255
access-list 64 permit 10.200.200.31
access-list 64 permit 10.200.200.34
access-list 64 permit 10.200.200.32
access-list 64 permit 10.200.200.33
```

One of the boundary routers (R2 in this example) selected the poor paths, because OSPF has a better administrative distance than RIP. You can change the administrative distance of the redistributed RIP routes to ensure that the boundary routers select the native RIP routes, as shown in the figure.

The **distance** command modifies the administrative distance of the OSPF routes to the networks that match ACL 64.

ACL 64 is used to match all the native RIP routes. The **access-list 64 permit 10.3.1.0** command configures a standard access control list (ACL) to permit the 10.3.1.0 network. Other similar access list statements permit the other internal native RIP networks.

In the figure, both of the redistributing routers are configured to assign an administrative distance of 125 to OSPF routes that are advertised for the networks that are listed in ACL 64.

ACL 64 has permit statements for the internal native RIP networks of 10.3.1.0, 10.3.2.0, and 10.3.3.0, as well as the loopback networks of 10.200.200.31, 10.200.200.32, 10.200.200.33, and 10.200.200.34.


When either one of the redistributing routers learns about these networks from RIP, it selects the routes that are learned from RIP (with a lower administrative distance of 120) over the same routes that are learned from OSPF (with an administrative distance of 125) and puts only the RIP routes in the routing table.

Note that the **distance** command is part of the OSPF routing process configuration, because the administrative distance should be changed for these routes when they are advertised by OSPF, not by RIP.

You need to configure the **distance** command on both redistributing routers, because either one can have suboptimal routes, depending on which redistributing router sends the OSPF updates about the RIP networks to the other redistributing router first.

Example: Redistribution Using Administrative Distance (Cont.)

With OSPF changing administrative distance



```
Gateway of last resort is not set

  172.31.0.0/16 is variably subnetted, 8 subnets, 2 masks
O   172.31.55.4/32 [110/781] via 172.31.33.4, 00:00:01, Serial0/0/0
C   172.31.33.0/24 is directly connected, Serial0/0/0
O   172.31.33.1/32 [110/1562] via 172.31.33.4, 00:00:01, Serial0/0/0
O   172.31.33.4/32 [110/781] via 172.31.33.4, 00:00:01, Serial0/0/0
O   172.31.44.4/32 [110/781] via 172.31.33.4, 00:00:01, Serial0/0/0
O   172.31.22.4/32 [110/781] via 172.31.33.4, 00:00:01, Serial0/0/0
O   172.31.11.4/32 [110/781] via 172.31.33.4, 00:00:03, Serial0/0/0
O   172.31.66.4/32 [110/781] via 172.31.33.4, 00:00:03, Serial0/0/0
  10.0.0.0/8 is variably subnetted, 8 subnets, 2 masks
R   10.3.1.0/24 [120/2] via 10.3.2.4, 00:00:03, FastEthernet0/0
R   10.3.3.0/24 [120/1] via 10.3.2.4, 00:00:03, FastEthernet0/0
C   10.3.2.0/24 is directly connected, FastEthernet0/0
R   10.200.200.31/32 [120/3] via 10.3.2.4, 00:00:04, FastEthernet0/0
R   10.200.200.34/32 [120/1] via 10.3.2.4, 00:00:04, FastEthernet0/0
C   10.200.200.32/32 is directly connected, Loopback0
R   10.200.200.33/32 [120/2] via 10.3.2.4, 00:00:04, FastEthernet0/0
O E2 10.254.0.0/24 [110/50] via 172.31.33.4, 00:00:04, Serial0/0/0
```

▪ Router R2 prefers RIP routes.

The figure shows that R2 now retains the more direct paths to the internal networks by learning them from RIP. However, some routing information is lost with this configuration. For example, depending on the actual bandwidths, the OSPF path may have been better for the 10.3.1.0 network. It may have made sense not to include 10.3.1.0 in the ACL.

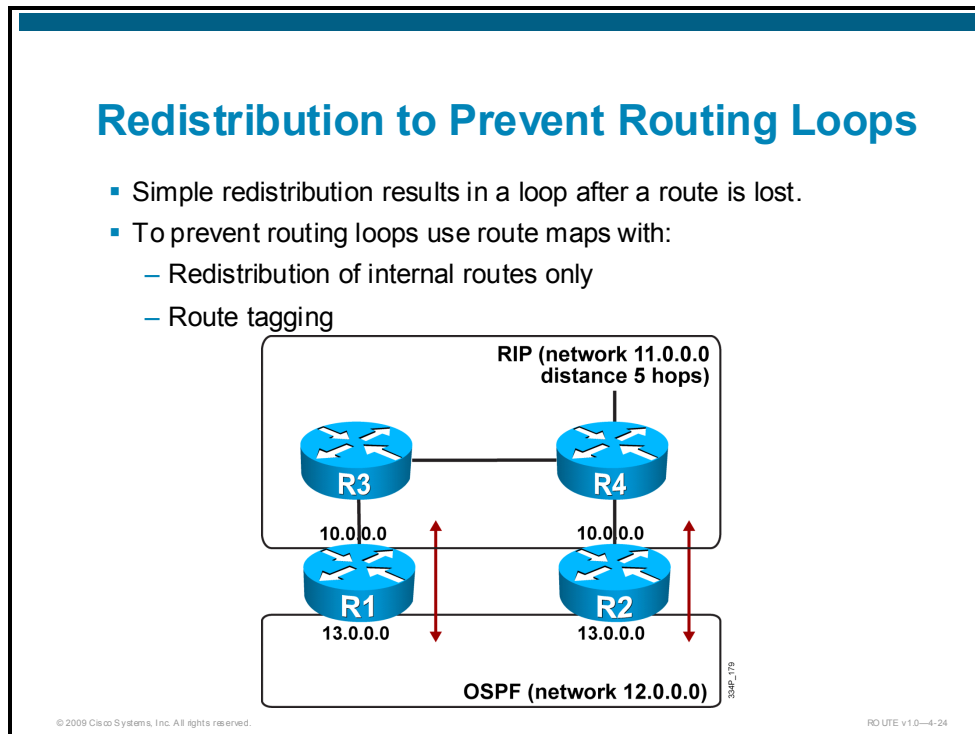
This example illustrates the importance of knowing your network before you implement redistribution, and of closely examining which routes that the routers are selecting after redistribution is enabled.

Pay particular attention to routers that can select from a number of possible redundant paths to a network, because they are more likely to select suboptimal paths.

The most important benefit of using administrative distance to control route preference is that no path information is lost; the OSPF information is still in the OSPF database. If the primary path is lost, the OSPF path can reassert itself, and the router will maintain connectivity with those networks.

Redistribution Using Route Maps

This topic describes how to implement route maps with route redistribution to prevent routing loops.



The figure represents multipoint two-way redistribution. With this kind of redistribution, it is very likely that routing loops exist. Suppose that R1 and R2 are both connected to a RIP network 10.0.0.0, and that they both learn from RIP routers about the network 11.0.0.0. Also suppose that R1 and R2 are both connected to an OSPF network 13.0.0.0, and that they both learn from other OSPF routers about the network 12.0.0.0.

Simple redistribution is configured as follows:

```
router rip
network 10.0.0.0
redistribute ospf 1 metric 3
!
router ospf 1
network 13.0.0.0
redistribute rip subnets
```

As network 11.0.0.0 is learned from the other RIP routers, it will also be redistributed into OSPF. Similarly, as network 12.0.0.0 is learned from other OSPF routers, it will also be redistributed into RIP.

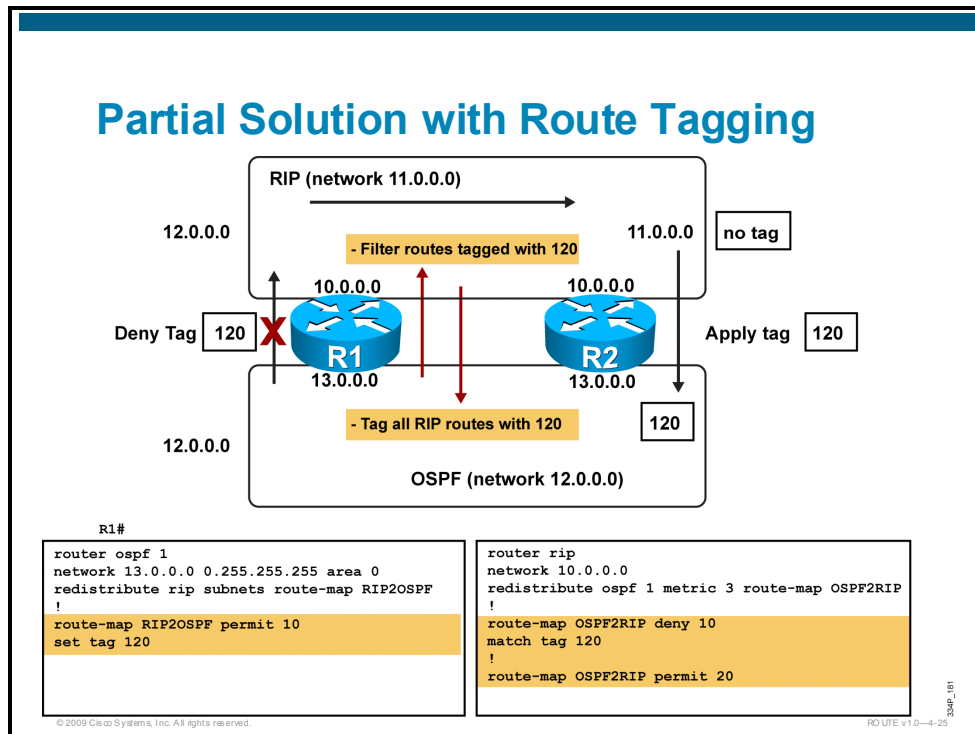
With this type of multipoint two-way redistribution, routing loops, or suboptimal routing, may occur. As an example, look at the 11.0.0.0 network. As this network is learned natively within the RIP section of the network, R4 first sees it with a hop count of 5. R4 then propagates this route to R3 and R2 with a hop count of 6. R3 propagates the route to R1 with a hop count of 7, while R2 redistributes into OSPF. Now R1 has a choice to make. It has a route to the 11.0.0.0 network from RIP with an administrative distance of 120 (RIP) and the same network with an administrative distance of 110 (OSPF). Because OSPF has a better (lower) administrative distance, R1 redistributes the network back to RIP with the metric that is set in the **redistribute**

command (**redistribute ospf 1 metric 3**). R3 then starts preferring the path R1-R2-R4 to reach 11.0.0.0, because the hop count advertised by R1 is 3 while the hop count advertised by R4 is 6. This results in suboptimal routing. Worse, because R3 now prefers the path to R1, it will advertise this to R4 with a hop count of 4. R4 now has the choice of the route from R3 with a hop count of 4 or the true path to the 11.0.0.0 network with a hop count of 5. R4 will select the path to R3 and advertise this to R2. There is now a routing loop (R4, R2, R1, R3, R4). Packets destined for the 11.0.0.0 network that enter this loop will bounce around the loop and never reach the destination. Network 11.0.0.0 becomes unreachable.

To prevent routing loops, use route maps for the following:

- Redistributing internal routes only
- Route tagging

Partial Solution with Route Tagging



To avoid routing loops, several solutions exist, which often have to be combined. Among these solutions you can either use route maps to redistribute internal routes or use route tagging. You can also manipulate the metric during the redistribution process.

The solution using route maps and tagging is shown in the figure. RIP routes that are redistributed into OSPF are tagged with the value 120 during the redistribution process. In the opposite direction, routes that are redistributed into RIP are filtering tagged updates. Updates with tag 120 are not allowed to go back to the RIP routing protocol. This will effectively break the routing loop.

As shown in the figure, the **route-map** command has two statements: One is for tagging, and the other one is for filtering updates.

Only one part of the configuration (for R1) is shown in the figure. The same configuration is needed in R2 as well.

Redistribution with Tagging Verification

```
R1#show ip ospf database external 11.0.0.0
OSPF Router with ID(13.0.0.1) (Process ID 1)
Type-5 AS External Link States
Routing Bit Set on this LSA
LS age: 313
Options: (No TOS-capability, DC)
LS Type: AS External Link
Link State ID: 11.0.0.0 (External Network Number )
Advertising Router: 13.0.0.2
LS Seq Number: 8000002A
Checksum: 0x8D02
Length: 36
Network Mask: /8
Metric Type: 2
TOS: 0
Metric: 20
Forward Address: 0.0.0.0
External Route Tag: 120

R1#show ip rip database 11.0.0.0 255.0.0.0
%Route not in database

R1#show ip route 11.0.0.0
Routing entry for 11.0.0.0/8
Known via "ospf 1", distance 110, metric 20, tag 120, type extern 2
advertised by rip metric 7
Last update from 13.0.0.2 on Serial0, 00:00:19 ago
Routing Descriptor Blocks:
* 13.0.0.2, from 13.0.0.2, 00:00:17 ago, via Serial0
Route metric is 20, traffic share count is 1
```

- Routing tables after tagging

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-4-28

The solution as implemented on the previous page will solve the routing issue for network 11.0.0.0. Network 11.0.0.0 will be redistributed into OSPF, but will not be redistributed back to RIP, and R3 will not use suboptimal routing to reach 11.0.0.0 anymore.

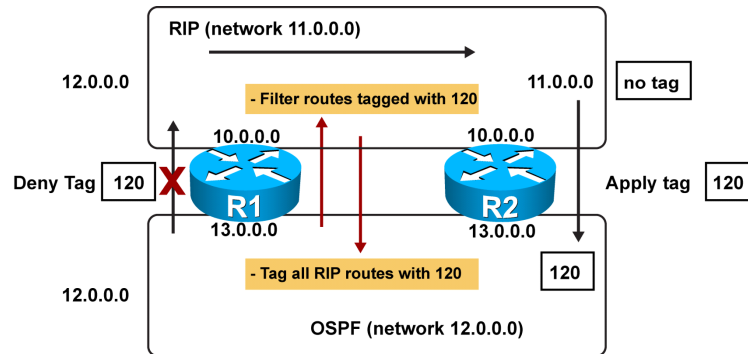
Nevertheless, this partial solution does not solve the issue for R1 and R2. Because OSPF is preferred to RIP, R1 will still learn network 11.0.0.0 from OSPF (via the redistribution through R2), and will still prefer the OSPF path to send packets to 11.0.0.0 subnets. An easy way to verify this point is to use the **show ip routing protocol database** command.

In the figure, the **show ip ospf database 11.0.0.0** command displays that network 11.0.0.0 is indeed known via OSPF, and has a tag of 120. This command is run on R1. Both R1 and R2 are redistributing RIP routes into OSPF. The same command on R2 would show that 11.0.0.0 is also known via OSPF, with R1 as the best next hop.

The **show ip rip database 11.0.0.0 255.0.0.0** command shows that the network is not known. This does not mean that R1 does not receive routes to 11.0.0.0 from RIP neighbors—it means that R1 prefers OSPF to reach 11.0.0.0, and therefore does not keep a route to 11.0.0.0 in its RIP database. The same command run on R2 would display the same result.

The **show ip route 11.0.0.0** command provides more information, showing that R1 prefers OSPF to reach network 11.0.0.0, but also learns about 11.0.0.0 from RIP, with a metric of 7. Running the same command on R2 would display the same result. Running the same command on any of the RIP routers (R3 or R4) would show that the RIP routers do learn properly about 11.0.0.0, and that the loop issue is solved on the RIP side.

Solution with Route Tagging and Distribute List



R1#

```
router ospf 1
network 13.0.0.0 0.255.255.255 area 0
redistribute rip subnets route-map RIP2OSPF
distribute-list 1 in serial0
!
route-map RIP2OSPF permit 10
set tag 120
!
access-list 1 deny 11.0.0.0
access-list 1 permit any
```

```
router rip
network 10.0.0.0
redistribute ospf 1 metric 3 route-map OSPF2RIP
!
route-map OSPF2RIP deny 10
match tag 120
!
route-map OSPF2RIP permit 20
```

An additional step needs to be performed to ensure that R1 will prefer RIP to reach a RIP network and, similarly, that R2 will prefer RIP to reach a RIP network. This can be done using several techniques, including distance or administrative distance manipulation. An additional possibility, presented here, is to ensure that R1 or R2 will not learn the route to 11.0.0.0 through OSPF. An easy way to achieve this goal is to use distribute lists. Distribute lists allow you to decide which routes will be sent to a router, or received by a router via a given routing protocol.

To configure a distribute list, first configure an access list defining the routes to be allowed or forbidden. In the example here, access list 1 denies the network 11.0.0.0 but allows all other networks.

Then, within the routing process, you can call the access list to apply its filter to the routes to be received or sent. In the example here, the distribute list is applied to the serial0 interface, inbound. The serial0 interface is the interface to the OSPF network, as seen on the previous page. The effect of this distribute list is that R1 will not receive any route to 11.0.0.0 via its serial0 interface in the OSPF process. Therefore, R1 will ignore the route to 11.0.0.0 redistributed by R2 into the OSPF process, and will use RIP to reach network 11.0.0.0.

Redistribution with Tagging and Distribute List Verification

```
R1#show ip route 12.0.0.0
Routing entry for 12.0.0.0/8
  Known via "ospf 1", distance 110, metric 20,
  tag 20, type extern 2
  Redistributing via ospf 1
  Last update from 13.0.0.2 on Serial0, 00:07:22 ago
  Routing Descriptor Blocks:
  * 13.0.0.2, from 13.8.0.1, 00:07:22 ago, via Serial0
    Route metric is 20, traffic share count is 1

R1#show ip route 11.0.0.0
Routing entry for 11.0.0.0/8
  Known via "rip", distance 120, metric 2
  tag 40
  Redistributing via rip
  Last update from 10.0.0.2 on Serial1, 00:00:15 ago
  Routing Descriptor Blocks:
  * 10.0.0.2, from 10.0.0.2, 00:00:15 ago, via Serial1
    Route metric is 2, traffic share count is 1
```

- Routing table after tagging and distribute list are applied
- R2 configuration and result would be similar

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-4-28

A similar configuration should be done on R2. This way, both R1 and R2 will use RIP to reach RIP networks, but 11.0.0.0 is still redistributed into OSPF, so that the other routers within the OSPF network will still learn the route to 11.0.0.0. This distribute list configuration is typically performed on ASBRs.

Summary

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

Summary

- Redistribution supports all routing protocols. Additionally, static and connected routes can be redistributed to allow the routing protocol to advertise the routes without using a network statement for them.
- Administrative distance is a way of ranking the trustworthiness of routing information. A lower administrative distance indicates that a route is more trustworthy.
- The **distance** (IP) command configures the administrative distance per specific network inserted into routing table, regardless of the IP routing protocol used. The specific protocol-dependent distance commands are used to define administrative distance for a specific routing protocol.

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—4-29

Summary (Cont.)

- Redistribution using administrative distance can correct path selection problems in a redistribution environment.
- It is very likely that routing loops are introduced when you use multipoint two-way redistribution. Two solutions are to redistribute only internal routes, to use route tagging, distribute lists, distance manipulation, or a combination of these techniques.

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—4-30

Lab 4-1 Debrief

Overview

In Lab 4-1, you configured route redistribution between multiple IP routing protocols. First, you had to configure RIPv2, OSPF, and EIGRP to implement redistribution between them. During the redistribution task, you then redistributed static routes and routes between routing protocols. In the last task, you implemented filtering for selective redistribution and to avoid routing loops for two-way redistribution.

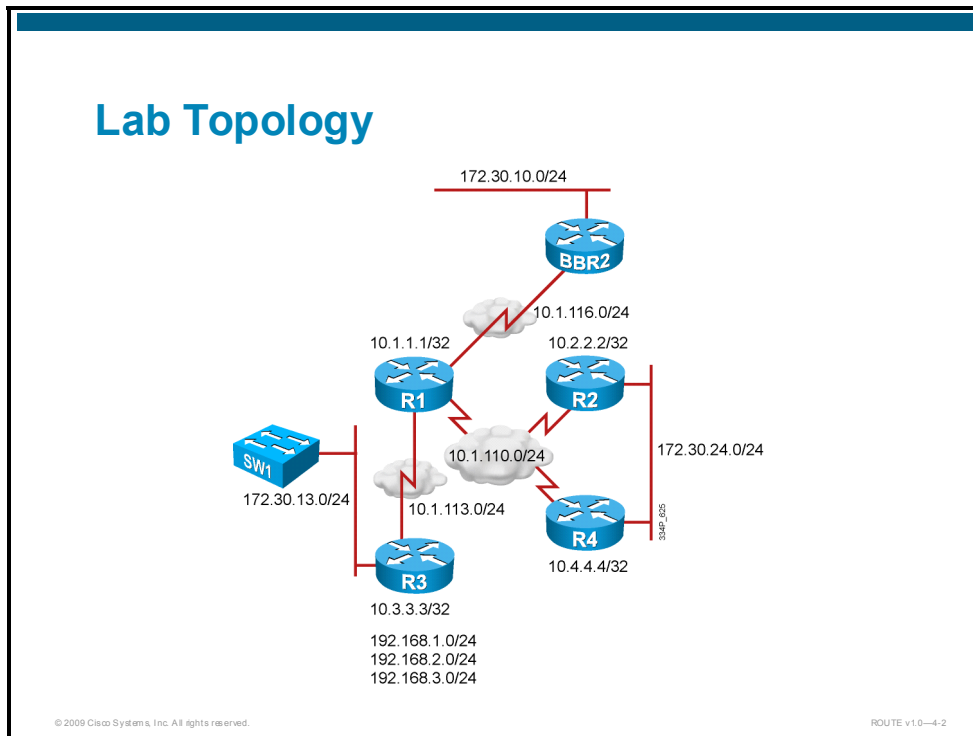
Objectives

Upon completing this lesson, you will be able to configure redistribution between multiple routing protocols, as well as implement route filtering. This ability includes being able to meet these objectives:

- Identify the implementation and verification tasks for redistribution between the OSPF and EIGRP and the OSPF and RIPv2 protocols
- Present a sample solution and identify possible alternative solutions

Lab Overview and Verification

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



The figure presents the physical lab topology that is used for Lab 4-1: “Configure Route Redistribution Between Multiple IP Routing Protocols.” The topology uses four pod routers and one backbone router. All routers participate in the multiple routing protocol configurations. Redistribution is configured between them on various routers that are based on the routing protocol configuration.

Based on the topology, you need to identify the required parameters and route redistribution between multiple IP routing protocols. Because two-way redistribution is configured in the network, loop-avoidance mechanisms must be used—filtering or another path-control tool.

Lab Review: What Did You Accomplish?

- **Task 1:** Implement RIP, EIGRP, and OSPF.
 - Which steps did you take to configure RIPv2, OSPF, and EIGRP?
- **Task 2:** Implement RIP redistribution.
 - Which steps did you take to redistribute loopback interfaces into RIP and to redistribute RIP routes into EIGRP?
 - Have you also configured a static route on R3 to establish connectivity in the opposite direction?
- **Task 3:** Implement OSPF two-way redistribution.
 - Which steps did you take to redistribute from OSPF to EIGRP?
 - Which steps did you take to redistribute from OSPF to RIP?

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-4-3

In the first task, you implemented the following routing protocols:

- RIP on routers R1 and R3
- EIGRP on router R1 to exchange the routing protocol information with BBR2
- OSPF on routers R1, R2, and R4

For RIP and OSPF configuration, all the parameters must be defined by the network administrator. For EIGRP, the requirements involve gathering the existing EIGRP routing configuration on BBR2, because EIGRP requires that the same AS number be used among all EIGRP routers in one EIGRP domain.

In the second task, redistribution for RIP involved redistribution of loopback interfaces into RIP and redistribution of the RIP routes into an EIGRP. For bidirectional connectivity, the way back from EIGRP to RIP needed to be defined; static routes are used for this.

In the third task, two-way redistribution between OSPF and EIGRP, as well as between OSPF and RIP, required careful implementation of filters, because routing loops should be avoided.

Verification

- Did you have enough information to create an implementation plan?
- Did you successfully configure RIPv2, EIGRP, and OSPF so that all the routes are present in the IP routing table?
- Did you successfully configure redistribution into RIP and verify the presence of the requested route?
- Did you successfully configure redistribution of RIP into EIGRP and verify the presence of the requested route?
- Did you successfully configure redistribution from OSPF into RIP and from OSPF into EIGRP?

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—4-4

A common approach to verifying the implementation process for a routing protocol is as follows:

- Evaluate if enough information was gathered to create a good implementation plan.
- Check the routes in the IP routing table belonging to each routing protocol. Successful configuration of RIPv2, EIGRP, and OSPF is important, because the base configuration must be prepared to successfully configure route redistribution between the protocols.
- Configure redistribution into RIP, and verify the presence of the routes inside the IP routing table.
- Configure redistribution from RIPv2 into EIGRP, and verify the presence of the routes inside the IP routing table.
- Configure redistribution from OSPF into RIPv2 and from OSPF into EIGRP, and verify the presence of the routes inside the IP routing table.

Checkpoints

- Configure RIPv2, EIGRP, and OSPF.
- Examine and verify that the routing protocols are up and that routes exist inside the IP routing table.
- Configure redistribution into RIP.
- Configure redistribution from RIP into EIGRP.
- Verify the RIP database and IP routing table.
- Configure two-way redistribution between OSPF and EIGRP.
- Verify an EIGRP topology and OSPF database.
- Configure two-way redistribution between OSPF and RIP.
- Verify the OSPF and RIP databases and IP routing table.

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0-4-5

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


- Configure RIP, EIGRP, and OSPF.
- Examine and verify that the routing protocols are up and that routes exist inside the IP routing table.
- Configure redistribution into RIPv2.
- Configure redistribution from RIPv2 into EIGRP.
- Verify the RIP database and IP routing table.
- Configure two-way redistribution between OSPF and EIGRP.
- Verify the EIGRP topology and OSPF database.
- Configure two-way redistribution between OSPF and RIP.
- Verify the OSPF and RIP databases and IP routing table.

Sample Solution and Alternatives

This topic describes a sample solution and possible alternatives.

Sample Solution

- RIPv2, OSPF, and EIGRP are configured on the routers.
- Redistribution takes place into RIP and from RIP to EIGRP.
- Loopback interfaces are redistributed into RIP and the RIP routes into EIGRP.
- Static routes are configured.
- Two-way redistribution from OSPF into EIGRP and from OSPF into RIP is configured.



© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—4-6


A sample solution includes the implementation details and the details for each task of the implementation plan. Different solutions are possible; the figure shows a few details of a successful configuration.

The proper implementation of route redistribution between multiple IP routing protocols includes the following details:

- RIP, OSPF, and EIGRP are implemented.
- The existing configuration is taken into account in terms of EIGRP AS number, IP addressing, and subnets that are used. Keep in mind the OSPF network type for the serial multipoint subinterface.
- Loopback interfaces are redistributed into RIP and the RIP routes into EIGRP.
- Static routes are configured if two-way redistribution is not implemented, in order to provide reachability back to the source routing protocol as well.
- Two-way redistribution from OSPF into EIGRP and from OSPF to RIP is configured. Route filtering is applied to avoid redistribution of the wrong routing information and the creation of a routing loop.

Alternative Solutions

- Different routing protocols can be used.
- Static routes can be used.
- Different path control tools can be used to manipulate the redistribution.
- Redistribution can be configured by changing the administrative distance and metric.



To provide reachability in the network, several routing protocols can be used. It is possible to migrate from one routing protocol to another to avoid redistribution or because the second routing protocol provides more optimized path selection and manipulation of routes.

Static routes can be used to manipulate the path or to influence the destination. Static routes provide more control over the reachability in the network, but using them is not a scalable solution; for every new destination network, a new static route must be defined. When the network grows, static routes are rarely used. In the alternative solution, static routes can be used instead of redistribution to provide reachability for a few destination networks. All other destinations that are not specified by the static routes are not reachable if there is no entry in the routing table.

Two-way redistribution requires either the use of route filtering or another path manipulation mechanism to avoid routing loops, or the use of filtering to hide destinations from networks.

For redistribution, changing the metric and administrative distance is another way to manipulate paths and the reachability of destinations.

Q and A

1. Why is routing protocol selection important?
2. Why is changing the administrative distance and metric important?
3. What must be used when configuring two-way redistribution?
4. Why are path control tools important?

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—4-8

1. A routing protocol exchanges routing updates and populates the IP routing table, which is used for destination-based forwarding. Different routing protocols process the routing updates in different ways.
2. The administrative distance defines the trustworthiness of the routes that are provided by a routing protocol. The metric provides the importance, or the quality, of routes within a routing protocol. By manipulating the administrative distance and metric value, you can implement path manipulation.
3. Because two-way redistribution is highly problematic, you must use route filters with it to avoid routing loops.
4. Path control tools are important for filtering routes that are not desired. They can also be used to manipulate the paths of packets to avoid suboptimal paths.

Summary

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

Summary

- Configure RIPv2, EIGRP, and OSPF to establish reachability in the network and enable destination-based forwarding.
- Configure redistribution into RIP and from RIP to EIGRP to maintain the optimum path across the network.
- Configure two-way redistribution of OSPF into EIGRP and OSPF into RIP, where the last example path control or route filtering must be used.

Module Summary

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

Module Summary

- Common network performance issues, including excessive routing updates, incorrectly configured route filters, and the use of different protocols in different areas within the same AS, can be controlled using prefix lists, distribution lists, and route maps. The **passive-interface** command can be used to suppress routing update traffic.
- Companies with complex networks face the challenges of political borders, geographic borders, and, sometimes, mergers with other companies.
- Routes are redistributed into a routing protocol by using the **redistribute** command under the routing process that is to receive the routes.

© 2009 Cisco Systems, Inc. All rights reserved. ROUTE v1.0—4-1

This module covered IP route redistribution and the control of redistributed routing updates. It also covered the use of passive interfaces and route maps for this control. Finally, it covered the use of route maps for policy-based routing (PBR).

Any two IP routing protocols can be redistributed. However, many types of incorrect information may be propagated. Passive interfaces, distribute lists, and route maps are some of the tools that are used to control these updates. Route maps may also be used to implement PBR for cost savings, quality of service (QoS), and other purposes that are driven by enterprise policy.

Module Self-Check

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

- Q1) Which two of these are common network performance issues? (Choose two.) (Source: Assessing Network Routing Performance and Security Issues)
- A) excessive routing updates
 - B) low-speed CPU in the router
 - C) use of different protocols in different areas within the same autonomous system
 - D) automatic summarization of routing protocols
- Q2) Which four tools can be used for route filtering? (Choose four.) (Source: Assessing Network Routing Performance and Security Issues)
- A) access lists
 - B) prefix lists
 - C) summarization
 - D) distribute lists
 - E) QoS
 - F) route maps
- Q3) Which matching method cannot be used with a prefix list filter? (Source: Assessing Network Routing Performance and Security Issues)
- A) mask filtering
 - B) using **ge**
 - C) using **le**
 - D) using **ge** and **le**
 - E) packet length
- Q4) Which two prefixes are matched by the following command: **ip prefix-list List2 permit 192.168.0.0/16 ge 18**? (Choose two.) (Source: Assessing Network Routing Performance and Security Issues)
- A) 192.168.0.0/16
 - B) 192.168.0.0/20
 - C) 192.168.2.0/24
 - D) 192.128.0.0/24
- Q5) Which prefix list matches all host routes? (Source: Assessing Network Routing Performance and Security Issues)
- A) **ip prefix-list A permit 0.0.0.0/0 ge 32**
 - B) **ip prefix-list B permit 128.0.0.0/2 ge 17**
 - C) **ip prefix-list C permit 0.0.0.0/0 le 32**
 - D) **ip prefix-list D permit 0.0.0.0/0**
- Q6) Which two tools can be used to match traffic that is controlled by a distribute list? (Choose two.) (Source: Assessing Network Routing Performance and Security Issues)
- A) access lists
 - B) prefix lists
 - C) route maps
 - D) match lists

- Q7) Which three of these are common route map applications? (Choose three.) (Source: Assessing Network Routing Performance and Security Issues)
- A) redistribution route filtering
 - B) policy-based routing
 - C) BGP policy implementation
 - D) network performance testing
- Q8) Route maps cannot match networks based on access lists and change the metric value and metric type during the redistribution process. (Source: Assessing Network Routing Performance and Security Issues)
- A) true
 - B) false
- Q9) The **passive-interface** command is used to suppress the sending of routing updates from an interface. In RIP, it suppresses both the sending and the receiving of routing updates. (Source: Assessing Network Routing Performance and Security Issues)
- A) true
 - B) false
- Q10) Which command sets all interfaces as passive by default? (Source: Assessing Network Routing Performance and Security Issues)
- A) **passive-interface all**
 - B) **passive-interface default**
 - C) **passive-interface default all**
 - D) **passive-interface all default**
- Q11) You can use PBR to bypass the default destination-based forwarding of routers and route maps that are applied to interfaces for the processing of inbound packets. (Source: Assessing Network Routing Performance and Security Issues)
- A) true
 - B) false
- Q12) Which steps must you use to configure policy routing? (Source: Assessing Network Routing Performance and Security Issues)
-
-
-
- Q13) Which configuration command must be used to apply the route map that is named **ChangeIntf** to locally originated traffic? (Source: Assessing Network Routing Performance and Security Issues)
- A) **ip policy route-map ChangeIntf**
 - B) **ip local policy route-map ChangeIntf**
 - C) **ip policy local route-map ChangeIntf**
 - D) **ip policy route-map local ChangeIntf**
- Q14) Redistribution supports BGP, EIGRP, IS-IS, OSPF, RIP, static, and connected routes. (Source: Configuring and Verifying Route Redistribution)
- A) true
 - B) false

- Q15) What does the **redistribute eigrp 100 subnets metric-type 1** command do when used in router configuration mode under OSPF process 1? (Source: Configuring and Verifying Route Redistribution)
- A) redistributes from EIGRP AS 100 into OSPF and changes the metric type from type 2 to type 1
 - B) redistributes from EIGRP AS 100 into OSPF process 1 and changes the metric value of each subnet
 - C) redistributes only from EIGRP AS 100 subnets of metric type 1
 - D) redistributes from EIGRP and processes the first 100 subnets of metric type 1
- Q16) All these statements are true about the **default-metric 10000 100 255 1 1500** command that is used in the EIGRP routing process configuration mode except which one? (Source: Configuring and Verifying Route Redistribution)
- A) This command is needed when redistributing from another protocol with an incompatible metric into EIGRP.
 - B) This command is not needed or applied when redistributing static routes to interfaces or connected networks.
 - C) This command is not needed or applied when redistributing between the EIGRP and IGRP processes.
 - D) This command is not needed or applied when redistributing between the EIGRP and RIP processes.
- Q17) Administrative distance is a way of ranking the trustworthiness of routing information. A higher administrative distance means a route is more trustworthy. (Source: Configuring and Verifying Route Redistribution)
- A) true
 - B) false
- Q18) Which routing protocol has the lowest administrative distance? (Source: Configuring and Verifying Route Redistribution)
- A) OSPF
 - B) external EIGRP
 - C) RIP
 - D) internal EIGRP
- Q19) In which two ways can you avoid routing loops that might result from two-way redistribution? (Choose two.) (Source: Configuring and Verifying Route Redistribution)
- A) Use the **passive-interface** command to suppress routing updates.
 - B) Use route maps to redistribute internal routes only.
 - C) Use route maps to implement route tagging.
 - D) Use the **no-loop** keyword when redistributing routes.

Module Self-Check Answer Key

- Q1) A, C
- Q2) A, B, D, F
- Q3) E
- Q4) B, C
- Q5) A
- Q6) A, C
- Q7) A, B, C
- Q8) B
- Q9) B
- Q10) B
- Q11) A
- Q12) Create a route map to configure a new policy for routing, apply the route map to an incoming interface or locally originated traffic.
- Q13) B
- Q14) A
- Q15) A
- Q16) D
- Q17) B
- Q18) D
- Q19) B, C