

Building Scalable Cisco Internetworks

Volume 1

Version 3.0

Student Guide

Editorial, Production, and Graphic Services: 06.14.06

**Corporate Headquarters**

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters

Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: 31 0 20 357 1000
Fax: 31 0 20 357 1100

Americas Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters

Cisco Systems, Inc.
168 Robinson Road
#28-01 Capital Tower
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

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

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia • Cyprus • Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia • Ireland • Israel • Italy • Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland • Portugal • Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden • Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe



© 2006 Cisco Systems, Inc. All rights reserved. CCSP, CCVP, the Cisco Square Bridge logo, Follow Me Browsing, and StackWise are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, and iQuick Study are service marks of Cisco Systems, Inc.; and Access Registrar, Aironet, BPX, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Fast Step, FormShare, GigaDrive, GigaStack, HomeLink, Internet Quotient, IOS, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, LightStream, Linksys, MeetingPlace, MGX, the Networkers logo, Networking Academy, Network Registrar, *Packet*, PIX, Post-Routing, Pre-Routing, ProConnect, RateMUX, ScriptShare, SlideCast, SMARTnet, The Fastest Way to Increase Your Internet Quotient, and TransPath are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0601R)

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.



***Students, this letter describes important
course evaluation access information!***

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

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

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

Sincerely,

Cisco Systems Learning

Table of Contents

Volume 1

<i>Course Introduction</i>	<i>1</i>
Overview	1
Learner Skills and Knowledge	1
Course Goal and Objectives	2
Course Flow	3
Additional References	4
Cisco Glossary of Terms	4
Your Training Curriculum	5
<i>Network Requirements</i>	<i>1-1</i>
Overview	1-1
Module Objectives	1-1
<i>Describing Network Requirements</i>	<i>1-3</i>
Overview	1-3
Objectives	1-3
Cisco Network Models	1-4
Enterprise Composite Network Model	1-8
Traffic Conditions in a Converged Network	1-10
Cisco SONA Framework and IIN	1-11
Routing and Routing Protocols	1-15
Summary	1-17
References	1-17
<i>Configuring EIGRP</i>	<i>2-1</i>
Overview	2-1
Module Objectives	2-1
<i>Introducing EIGRP</i>	<i>2-3</i>
Overview	2-3
Objectives	2-3
EIGRP Capabilities and Attributes	2-4
Underlying Processes and Technologies	2-6
EIGRP Operation	2-8
Example: EIGRP Tables	2-12
EIGRP Metric	2-15
Calculating the EIGRP Metric	2-16
Example: EIGRP Metric Calculation	2-18
Integrating the EIGRP and IGRP Routes	2-20
Summary	2-21
<i>Implementing and Verifying EIGRP</i>	<i>2-23</i>
Overview	2-23
Objectives	2-23
Configuring Basic EIGRP	2-24
Example: Configuring EIGRP for IP	2-27
Using a Wildcard Mask in EIGRP	2-29
Example: Wildcard Mask in EIGRP	2-29
Configuring the ip default-network Command	2-30
Example: ip default-network Command	2-31
Verify EIGRP IP Routes	2-32
Example: EIGRP Configuration	2-32
Example: R2 EIGRP Configuration	2-33
Verify EIGRP IP Operations	2-37
Summary	2-43

<i>Configuring Advanced EIGRP Options</i>	2-45
Overview	2-45
Objectives	2-45
Route Summarization	2-46
Configuring Manual Route Summarization	2-49
Example: Summarizing EIGRP Routes	2-50
Example: Router C Routing Table	2-51
Load Balancing Across Equal Paths	2-52
Configuring Load Balancing Across Unequal-Cost Paths	2-53
Example: Variance	2-54
EIGRP Bandwidth Use Across WAN Links	2-56
Configuring EIGRP Bandwidth Use Across WAN Links	2-59
Example: WAN Configuration—Frame Relay Hub-and-Spoke Topology	2-59
Example: WAN Configuration—Hybrid Multipoint	2-61
Summary	2-62
<i>Configuring EIGRP Authentication</i>	2-63
Overview	2-63
Objectives	2-63
Router Authentication	2-64
MD5 Authentication	2-66
Configuring MD5 Authentication	2-68
Example: MD5 Authentication Configuration	2-73
Example: R1 Configuration for MD5 Authentication	2-74
Example: R2 Configuration for MD5 Authentication	2-75
Verifying MD5 Authentication	2-76
Troubleshooting MD5 Authentication	2-77
Example: Successful MD5 Authentication	2-77
Example: Troubleshooting MD5 Authentication Problems	2-78
Summary	2-79
<i>Using EIGRP in an Enterprise Network</i>	2-81
Overview	2-81
Objectives	2-81
Scalability in Large Networks	2-82
EIGRP Queries	2-83
EIGRP Stubs	2-84
Example: Limiting Updates and Queries: Using EIGRP Stub	2-88
Example: eigrp stub Parameters	2-89
SIA Connections	2-91
Preventing SIA Connections	2-92
Graceful Shutdown	2-94
Summary	2-96
Module Summary	2-97
Module Self-Check	2-99
Module Self-Check Answer Key	2-108
<i>Configuring OSPF</i>	3-1
Overview	3-1
Module Objectives	3-1

<i>Introducing the OSPF Protocol</i>	3-3
Overview	3-3
Objectives	3-3
Link-State Routing Protocols	3-4
OSPF Area Structure	3-7
OSPF Adjacency Databases	3-10
Calculating the OSPF Metric	3-13
Example: SPF Calculation	3-14
Link-State Data Structures	3-15
Summary	3-16
<i>OSPF Packet Types</i>	3-17
Overview	3-17
Objectives	3-17
OSPF Packet Types	3-18
Establishing OSPF Neighbor Adjacencies	3-20
Exchanging and Synchronizing LSDBs	3-22
Maintaining Network Routes	3-26
Maintaining Link-State Sequence Numbers	3-28
Example: LSA Sequence Numbers and Maximum Age	3-29
Verifying Packet Flow	3-30
Example: debug ip ospf packet	3-30
Summary	3-32
<i>Configuring OSPF Routing</i>	3-33
Overview	3-33
Objectives	3-33
Configuring Basic Single-Area and Multiarea OSPF	3-34
Example: Configuring OSPF on Internal Routers of a Single Area	3-36
Example: Configuring OSPF for Multiple Areas	3-37
Configuring a Router ID	3-38
Verifying the OSPF Router ID	3-41
Verifying OSPF Operation	3-43
Example: The show ip route ospf Command	3-45
Example: The show ip ospf interface Command	3-46
Example: The show ip ospf neighbor Command	3-47
Summary	3-49
<i>OSPF Network Types</i>	3-51
Overview	3-51
Objectives	3-51
Introducing OSPF Network Types	3-52
Adjacency Behavior for a Point-to-Point Link	3-53
Adjacency Behavior for a Broadcast Network Link	3-54
Selecting the DR and BDR	3-56
Adjacency Behavior for an NBMA Network	3-58
OSPF over Frame Relay Configuration Options	3-60
Example: Sample Configuration of a Router Using OSPF Broadcast Mode	3-63
OSPF over Frame Relay NBMA Configuration	3-64
Example: neighbor Command	3-66
Example: show ip ospf neighbor Command	3-67
OSPF over Frame Relay Point-to-Multipoint Configuration	3-68
Example: Point-to-Multipoint Configuration	3-69
Using Subinterfaces in OSPF over Frame Relay Configuration	3-72
Example: Point-to-Point Subinterface	3-74
Example: Multipoint Subinterface	3-76
Example: OSPF over NBMA Topology Summary	3-77

Tracking OSPF Adjacencies	3-78
Example: debug Output for Point-to-Point Mode	3-78
Example: debug ip ospf adj Output for Broadcast Mode	3-79
Summary	3-82

Link-State Advertisements **3-83**

Overview	3-83
Objectives	3-83
OSPF Router Types	3-84
Example: OSPF Hierarchical Routing	3-85
OSPF Virtual Links	3-88
Example: OSPF Virtual Link Configuration	3-91
Example: show ip ospf virtual-links Command	3-92
OSPF LSA Types	3-95
Type 1	3-95
Type 2	3-95
Types 3 and 4	3-96
Type 5	3-96
Type 6	3-96
Type 7	3-96
Type 8	3-96
Types 9, 10, and 11	3-96
Example: LSA Type 4—Summary LSA	3-100
Interpreting the OSPF LSDB and Routing Table	3-102
Example: Interpreting the OSPF Database	3-102
Configuring OSPF LSDB Overload Protection	3-108
Changing the Cost Metric	3-110
Summary	3-111

OSPF Route Summarization **3-113**

Overview	3-113
Objectives	3-113
OSPF Route Summarization	3-114
Example: Using Route Summarization	3-116
Configuring OSPF Route Summarization	3-117
Example: Route Summarization Configuration at ABR	3-119
Example: Route Summarization Configuration at ASBR	3-120
Benefits of a Default Route in OSPF	3-121
Example: Default Routes in OSPF	3-121
Configuring a Default Route in OSPF	3-122
Example: Default Route Configuration	3-124
Summary	3-125

Configuring OSPF Special Area Types **3-127**

Overview	3-127
Objectives	3-127
Configuring OSPF Area Types	3-128
Configuring Stub Areas	3-130
Example: OSPF Stub Area Configuration	3-133
Configuring Totally Stubby Areas	3-134
Example: Totally Stubby Configuration	3-136
Interpreting Routing Tables	3-137
Example: Routing Table in a Standard Area	3-137
Example: Routing Table in a Stub Area	3-138
Example: Routing Table in a Stub Area with Summarization	3-139
Example: Routing Table in a Totally Stubby Area	3-140
Configuring NSSAs	3-141
Example: NSSA Configuration	3-144
Example: NSSA Totally Stubby Configuration	3-145

Verifying All Stub Area Types	3-146
Summary	3-147
<i>Configuring OSPF Authentication</i>	<i>3-149</i>
Overview	3-149
Objectives	3-149
Types of Authentication	3-150
Configuring Simple Password Authentication	3-151
Example: Simple Password Authentication Configuration	3-153
Example: R2 Configuration for Simple Password Authentication	3-154
Verifying Simple Password Authentication	3-155
Configuring MD5 Authentication	3-156
Example: MD5 Authentication Configuration	3-159
Example: R2 Configuration for MD5 Authentication	3-160
Verifying MD5 Authentication	3-161
Troubleshooting Simple Password Authentication	3-162
Example: Successful Simple Password Authentication	3-162
Example: Troubleshooting Simple Password Authentication Problems	3-164
Troubleshooting MD5 Authentication	3-165
Example: Successful MD5 Authentication	3-165
Example: Troubleshooting MD5 Authentication Problems	3-167
Summary	3-168
Module Summary	3-169
Module Self-Check	3-171
Module Self-Check Answer Key	3-188
<i>The IS-IS Protocol</i>	<i>4-1</i>
Overview	4-1
Module Objectives	4-1
<i>Introducing IS-IS and Integrated IS-IS Routing</i>	<i>4-3</i>
Overview	4-3
Objectives	4-3
IS-IS Routing	4-4
Integrated IS-IS Routing	4-8
Principles and Issues of Integrated IS-IS Design	4-9
The ES-IS Protocol	4-11
OSI Routing Levels	4-13
Level 0 Routing	4-13
IS-IS Level 1 Routing	4-13
IS-IS Level 2 Routing	4-14
Level 3 Routing	4-14
Summary of Routing Levels	4-14
Comparing IS-IS to OSPF	4-15
Summary of Differences between OSPF and Integrated IS-IS	4-19
Summary	4-20

Performing IS-IS Routing Operations

4-21

Overview	4-21
Objectives	4-21
NSAP Addresses	4-22
NET Addresses	4-28
IS-IS Routing Levels	4-30
Intra-Area and Interarea Addressing and Routing	4-31
Example: Identifying Systems—OSI Addressing in Networks	4-33
Example: OSI Area Routing	4-34
IS-IS PDUs	4-37
Link-State Packets	4-39
Example: LSP TLV Examples	4-41
Implementing IS-IS in NBMA Networks	4-42
Implementing IS-IS in Broadcast Networks	4-44
LSP and IIH Levels	4-47
Level 1 and Level 2 LSP	4-47
Level 1 and Level 2 IIH	4-47
Example: Comparing Broadcast and Point-to-Point Topologies	4-49
LSDB Synchronization	4-50
Example: LSDB Synchronization—LAN	4-52
Example: LSDB Synchronization: Point-to-Point	4-53
Example: WAN Adjacencies	4-55
Summary	4-56

Configuring Basic Integrated IS-IS

4-57

Overview	4-57
Objectives	4-57
Integrated IS-IS in a CLNS Environment	4-58
Configuring Integrated IS-IS	4-61
Example: Simple Integrated IS-IS Configuration	4-66
Optimizing IS-IS	4-67
Example: Tuning IS-IS Configuration	4-70
Configuring Route Summarization in IS-IS	4-71
Verifying IS-IS Configuration	4-72
Example: Is Integrated IS-IS Running?	4-72
Verifying CLNS IS-IS Structures	4-74
Example: OSI Intra-Area and Interarea Routing	4-76
Summary	4-80
Module Summary	4-81
References	4-81
Module Self-Check	4-83
Module Self-Check Answer Key	4-92
Example: WAN Adjacencies	4-94

Course Introduction

Overview

Building Scalable Cisco Internetworks (BSCI) v3.0 is recommended training for individuals seeking Cisco CCNP® certification. The course instructs network administrators of medium-to-large network sites on the use of advanced routing in implementing scalability for Cisco routers that are connected to LANs and WANs. The goal is to train network administrators to dramatically increase the number of routers and sites using these techniques instead of redesigning the network when additional sites or wiring configurations are added.

Learner Skills and Knowledge

This topic lists the skills and knowledge that learners must possess to benefit fully from the course.

Learner Skills and Knowledge

Cisco CCNA® certification

Note: Practical experience with deploying and operating networks based on Cisco network devices and Cisco IOS software is strongly recommended.

Course Goal and Objectives

This topic describes the course goal and objectives.

Course Goal

“To train network administrators on the techniques to plan, implement, and monitor a scalable IP routing network.”

Building Scalable Cisco Internetworks



© 2006 Cisco Systems, Inc. All rights reserved. BSCI v3.0-1-3

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

- Describe the converged network requirements of various network and networked applications within the Cisco network architectures
- Implement and verify EIGRP operations
- Build a scalable multiarea network with OSPF
- Configure Integrated IS-IS in a single area
- Manipulate routing and packet flow
- Implement and verify BGP for enterprise ISP connectivity
- Implement and verify multicast forwarding using PIM and related protocols
- Describe how IPv6 functions to satisfy the increasingly complex requirements of hierarchical addressing

Course Flow

This topic presents the suggested flow of the course materials.


Course Flow					
	Day 1	Day 2	Day 3	Day 4	Day 5
A M	Course Introduction	Configuring OSPF	The IS-IS Protocol	Implementing BGP	Implementing Multicast
	Network Requirements				
	Configuring EIGRP				
Lunch					
P M	Configuring EIGRP	Configuring OSPF	Manipulating Routing Updates	Implementing BGP	Implementing IPv6

© 2006 Cisco Systems, Inc. All rights reserved. BSCI v3.0—1-4

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

Additional References

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



Cisco Glossary of Terms

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

Your Training Curriculum

This topic presents the training curriculum for this course.



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

Learner Introductions

- **Your name**
- **Your company**
- **Skills and knowledge**
- **Brief history**
- **Objective**



© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—1-7

Please introduce yourself to the class.

Network Requirements

Overview

The convergence of voice, video, and data has not only changed the conceptual network models but has also affected the way that networks support services and applications.

This module describes Cisco conceptual models and architectures for converged networks.

Module Objectives

Upon completing this module, you will be able to describe the converged network requirements of various network and networked applications within the Cisco network architectures.

Describing Network Requirements

Overview

This lesson starts by introducing Cisco Enterprise Architectures and describing how they align with the traditional three-layer hierarchical network model. The Cisco Enterprise Composite Network Model is examined, and the traffic patterns in converged networks are discussed. The Cisco vision of the future Intelligent Information Network (IIN) and the Service-Oriented Network Architecture (SONA) are introduced. The lesson concludes with a discussion of where routing protocols fit into these models.


Objectives

Upon completing this lesson, you will be able to describe the converged network requirements of various network and networked applications within the Cisco network architectures. This ability includes being able to meet these objectives:

- Explain the Cisco conceptual network models, such as Cisco Enterprise Architectures and the Cisco hierarchical network model
- Describe the Cisco Enterprise Composite Network Model
- Describe the traffic conditions in a converged network
- Describe the IIN and the Cisco SONA framework
- Describe routing and routing protocols

Cisco Network Models

This topic describes Cisco network models, starting with the Cisco Enterprise Architectures and their mapping to traditional three-layer hierarchical network model.



Cisco provides an enterprise-wide systems architecture that helps companies to protect, optimize, and grow the infrastructure that supports their business processes. The architecture provides for integration of the entire network—campus, data center, WAN, branches, and teleworkers—offering staff secure access to tools, processes, and services.

The Cisco Enterprise *Campus* Architecture combines a core infrastructure of intelligent switching and routing with tightly integrated productivity-enhancing technologies, including IP communications, mobility, and advanced security. The architecture provides the enterprise with high availability through a resilient multilayer design, redundant hardware and software features, and automatic procedures for reconfiguring network paths when failures occur. Multicast provides optimized bandwidth consumption, and quality of service (QoS) prevents oversubscription to ensure that real-time traffic, such as voice and video, or critical data is not dropped or delayed. Integrated security protects against and mitigates the impact of worms, viruses, and other attacks on the network—even at the port level. Cisco enterprise-wide architecture extends support for standards, such as 802.1x and Extensible Authentication Protocol (EAP). It also provides the flexibility to add IPsec and Multiprotocol Label Switching (MPLS) virtual private networks (VPNs), identity and access management, and VLANs to compartmentalize access. These features help improve performance and security and decrease costs.


The Cisco Enterprise *Data Center* Architecture is a cohesive, adaptive network architecture that supports the requirements for consolidation, business continuance, and security while enabling emerging service-oriented architectures, virtualization, and on-demand computing. IT staff can easily provide departmental staff, suppliers, or customers with secure access to applications and resources, which simplifies and streamlines management, significantly reducing overhead. Redundant data centers provide backup using synchronous and asynchronous data and application replication. The network and devices offer server and application load balancing to maximize performance. This solution allows the enterprise to scale without major changes to the infrastructure.

The Cisco Enterprise *Branch* Architecture allows enterprises to extend head-office applications and services, such as security, IP communications, and advanced application performance to thousands of remote locations and users or to a small group of branches. Cisco integrates security, switching, network analysis, caching, and converged voice and video services into a series of integrated services routers in the branch so that the enterprises can deploy new services when they are ready without buying new equipment. This solution provides secure access to voice, mission-critical data, and video applications—anywhere, anytime. Advanced network routing, VPNs, redundant WAN links, application content caching, and local IP telephony call processing provide a robust architecture with high levels of resilience for all the branch offices. An optimized network leverages the WAN and LAN to reduce traffic and save bandwidth and operational expenses. The enterprise can easily support branch offices with the ability to centrally configure, monitor, and manage devices located at remote sites, including tools such as AutoQoS that proactively resolve congestion and bandwidth issues before they affect network performance.

The Cisco Enterprise *Teleworker* Architecture allows enterprises to securely deliver voice and data services to remote small or home offices over a standard broadband access service, providing a business resiliency solution for the enterprise and a flexible work environment for employees. Centralized management minimizes the IT support costs, and robust integrated security mitigates the unique security challenges of this environment. Integrated security and identity-based networking services enable the enterprise to help extend campus security policies to the teleworker. Staff can securely log in to the network over an “always-on” VPN and gain access to authorized applications and services from a single cost-effective platform. Productivity can further be enhanced by adding an IP phone, providing cost-effective access to a centralized IP communications system with voice and unified messaging services.

The Cisco Enterprise *WAN* Architecture offers the convergence of voice, video, and data services over a single IP communications network, which enables the enterprise to cost-effectively span large geographic areas. QoS, granular service levels, and comprehensive encryption options help ensure the secure delivery of high-quality corporate voice, video, and data resources to all corporate sites, enabling staff to work productively and efficiently wherever they are located. Security is provided with multiservice VPNs (IPsec and MPLS) over Layer 2 or Layer 3 WANs or hub-and-spoke or full-mesh topologies.

Cisco Hierarchical Network Model




Traditionally, the three-layer hierarchical model has been used in network design. The model provides a modular framework that allows flexibility in network design and facilitates implementation and troubleshooting. The hierarchical model divides networks or their modular blocks into the access, distribution, and core layers, with these features:


- **Access layer:** This layer is used to grant user access to network devices. In a network campus, the access layer generally incorporates switched LAN devices with ports that provide connectivity to workstations and servers. In the WAN environment, the access layer at remote sites or teleworkers may provide access to the corporate network across WAN technology.
- **Distribution layer:** This layer aggregates the wiring closets and uses switches to segment workgroups and isolate network problems in a campus environment. Similarly, the distribution layer aggregates WAN connection at the edge of the campus and provides policy-based connectivity.
- **Core layer (also referred to as the backbone):** This layer is a high-speed backbone and is designed to switch packets as fast as possible. Because the core is critical for connectivity, it must provide a high level of availability and adapt to changes very quickly.

Note The hierarchical model can be applied to any network type, such as LANs, WANs, wireless LANs (WLANs), metropolitan-area networks (MANs), and VPNs, and to any modular block of the Cisco networking model.

Hierarchical Campus Model



Hierarchical Network Model WAN




For example, the hierarchical model can be applied specifically to the enterprise campus.

It can also be applied to the enterprise WAN. Obviously, another model is required to break down and analyze an existing modern enterprise network or to plan a new one.

Enterprise Composite Network Model

This topic describes the Enterprise Composite Network Model.




Since the intelligent network service *security* has become of critical importance to all network planning and implementation, Cisco has developed a set of best practices for security. These best practices constitute a blueprint for network designers and administrators for the proper deployment of security solutions to support network solutions and the existing network infrastructure. This blueprint is called “SAFE”.

SAFE includes the Enterprise Composite Network Model, which can be used by network professionals to describe and analyze any modern enterprise network.

Three functional areas are defined by the model:

- **Enterprise Campus:** This functional area contains the modules required to build a hierarchical, highly robust campus network. Access, distribution, and core principles are applied to these modules.
- **Enterprise Edge:** This functional area aggregates connectivity from the various elements at the edge of the enterprise network. It provides a description of connectivity to remote locations, the Internet, and remote users.
- **Service Provider Edge:** This area provides a description of connectivity to service providers such as Internet service providers (ISPs), WAN providers, and the public switched telephone network (PSTN).



Various modules form an integrated converged network that supports business processes.

As shown in the figure, the campus comprises six modules:

- Building, with access switches and end devices (PCs and IP phones)
- Building distribution, with distribution multilayer switches
- Core, sometimes called the backbone
- Edge distribution, which concentrates all branches and teleworkers accessing the campus via WAN or Internet
- Server farm, which represents the data center
- Management, which represents the network management functionality

Additional modules in the other functional areas represent e-commerce functionality, corporate Internet connections, remote access and VPN connections, and traditional WAN (Frame Relay, ATM, and leased lines with PPP) connections.

Traffic Conditions in a Converged Network

This topic describes the traffic types and requirements in converged networks.

Network Traffic Mix and Requirements

- **Converged network traffic mix:**
 - **Voice and video traffic**
 - **Voice applications traffic**
 - **Mission-critical applications traffic**
 - **Transactional traffic**
 - **Routing update traffic**
 - **Network management traffic**
- **Key requirements:**
 - **Performance (bandwidth, delay, jitter)**
 - **Security (access, transmission)**

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0–1.4

Converged networks with integrated voice, video, and data contain various traffic patterns:

- Voice and video traffic, for example, IP telephony, and video broadcast and conferencing
- Voice applications traffic, generated by voice-related applications (such as contact centers)
- Mission-critical traffic, generated, for example, by stock exchange applications
- Transactional traffic, generated by e-commerce applications
- Routing update traffic, from routing protocols like Routing Information Protocol (RIP), Open Shortest Path First Protocol (OSPF), Enhanced Interior Gateway Routing Protocol (EIGRP), Intermediate System-to-Intermediate System Protocol (IS-IS), and Border Gateway Protocol (BGP)
- Network management traffic

The diversity of the traffic mix poses stringent requirements on the network in terms of performance and security. The requirements significantly differ, depending on the traffic type. For example, voice and video require constant bandwidth and low delay and jitter, while the transactional traffic requires high reliability and security with relatively low bandwidth. Video traffic is frequently carried as IP multicast traffic. Also, voice applications, such as IP telephony, require high reliability and availability because the user expectations for “dial tone” in the IP network are exactly the same as in traditional phone network. To meet the traffic requirements in the network, for example, voice and video traffic must be treated differently from other traffic, such as web-based traffic. QoS mechanisms are mandatory in converged networks.

Security is a key issue not only in fixed networks but also in wireless mobility, where access to the network is possible virtually anywhere. Several security strategies, such as device hardening with strict access control and authentication, intrusion protection, intrusion detection, traffic protection with encryption, and others, can minimize or even totally remove network security threats.

Cisco SONA Framework and IIN

This topic describes Cisco SONA, which guides an evolution of enterprise networks toward IIN; the IIN and its features are also described.

Cisco SONA Framework


- **Cisco Service-Oriented Network Architecture (SONA) is an architectural framework.**
- **Cisco SONA brings several advantages to enterprises:**
 - **Outlines how enterprises can evolve toward the Intelligent Information Network (IIN)**
 - **Illustrates how to build integrated systems across a fully converged intelligent network**
 - **Improves flexibility and increases efficiency**
 - **Optimizes applications, processes, and resources**

© 2006 Cisco Systems, Inc. All rights reserved. BSCI v3.0—1.8

Cisco is helping organizations to address new IT challenges, such as the deployment of service-oriented architectures, web services, and virtualization. Cisco SONA is an architectural framework that guides the evolution of enterprise networks to an IIN. The Cisco SONA framework provides several advantages to enterprises:

- Outlines the path toward the IIN
- Illustrates how to build integrated systems across a fully converged IIN
- Improves flexibility and increases efficiency, which results in optimized applications, processes, and resources

Cisco SONA uses the extensive product line services, proven architectures, and experience of Cisco and its partners to help enterprises achieve their business goals.



The Cisco SONA framework shows how integrated systems can both allow a dynamic, flexible architecture and provide for operational efficiency through standardization and virtualization. It centers on the concept that the network is the common element that connects and enables all components of the IT infrastructure. Cisco SONA outlines these three layers of the IIN:

- **Networked infrastructure layer:** This layer is where all of the IT resources are interconnected across a converged network foundation. The IT resources include servers, storage, and clients. The network infrastructure layer represents how these resources exist in different places in the network, including the campus, branch, data center, WAN and MAN, and teleworker. The objective for customers in this layer is to have “anywhere and anytime” connectivity.
- **Interactive services layer:** This layer enables efficient allocation of resources to applications and business processes delivered through the networked infrastructure. This layer comprises these services:
 - Voice and collaboration services
 - Mobility services
 - Security and identity services
 - Storage services
 - Computer services
 - Application networking services
 - Network infrastructure virtualization
 - Services management
 - Adaptive management services

Application layer: This layer includes business applications and collaboration applications. The objective for customers in this layer is to meet business requirements and achieve efficiencies by leveraging the interactive services layer.

Intelligent Information Network

- IIN integrates networked resources and information assets.
- IIN extends intelligence across multiple products and infrastructure layers.
- IIN actively participates in the delivery of services and applications.
- Three phases in building an IIN are:
 - Integrated transport
 - Integrated services
 - Integrated applications

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—1-11

The Cisco vision of the future IIN encompasses these features:

- Integration of networked resources and information assets that have been largely unlinked. The modern converged networks with integrated voice, video, and data require that IT departments more closely link the IT infrastructure with the network.
- Intelligence across multiple products and infrastructure layers. The intelligence built into each component of the network is extended network-wide and applies end to end.
- Active participation of the network in the delivery of services and applications. With added intelligence, the IIN makes it possible for the network to actively manage, monitor, and optimize service and application delivery across the entire IT environment.

With the listed features, the IIN offers much more than basic connectivity, bandwidth for users, and access to applications. The IIN offers an end-to-end functionality and centralized, unified control that promotes true business transparency and agility.


The IIN technology vision offers an evolutionary approach that consists of three phases in which functionality can be added to the infrastructure as required:

- **Integrated transport:** Everything—data, voice, and video—consolidates onto an IP network for secure network convergence. By integrating data, voice, and video transport into a single, standards-based, modular network, organizations can simplify network management and generate enterprise-wide efficiencies. Network convergence also lays the foundation for a new class of IP-enabled applications delivered through Cisco IP Communications solutions.
- **Integrated services:** Once the network infrastructure has been converged, IT resources can be pooled and shared or “virtualized” to flexibly address the changing needs of the organization. Integrated services help to unify common elements, such as storage and data center server capacity. By extending virtualization capabilities to encompass server, storage, and network elements, an organization can transparently use all of its resources more efficiently. Business continuity is also enhanced because shared resources across the IIN provide services in the event of a local systems failure.

- **Integrated applications:** With Cisco Application-Oriented Networking (AON) technology, Cisco has entered the third phase of building the IIN. This phase focuses on making the network “application aware” so that it can optimize application performance and more efficiently deliver networked applications to users. In addition to capabilities such as content caching, load balancing, and application-level security, Cisco AON makes it possible for the network to simplify the application infrastructure by integrating intelligent application message handling, optimization, and security into the existing network.

Routing and Routing Protocols

This topic describes routing and routing protocols.



To review, the focus of this course is on selecting, planning, implementing, tuning, and troubleshooting IP advanced routing protocols. It is a Cisco CCNP®-level technical course.

All of the models and tools described previously are important in the initial part of this process—selecting and planning.

The best practice is to use one IP routing protocol throughout the enterprise if possible. In many cases, this practice is not possible, which will be discussed in detail in another module. For example, Border Gateway Protocol (BGP) will be a factor in the Corporate Internet and E-Commerce modules if multihoming to ISPs is implemented. For remote access and VPN users, static routes are almost always used. Therefore, dealing with multiple routing protocols is likely.

The Enterprise Composite Network Model can assist in determining where each routing protocol is implemented, where the boundaries are, and how traffic flows are managed.

It is obvious that advanced IP routing protocols must be implemented in all core networks to support high availability requirements. Less advanced routing protocols (such as RIP and Interior Gateway Routing Protocol [IGRP]) and static routes may exist at the access and distribution levels within modules.

Routing Protocol Comparison

Parameters	EIGRP	OSPF	IS-IS
Size of Network (Small-Medium-Large-Very Large)	Large	Large	Very Large
Speed of Convergence (Very High-High-Medium-Low)	Very High	High	High
Use of VLSM (Yes-No)	Yes	Yes	Yes
Mixed-Vendor Devices (Yes-No)	No	Yes	Yes
Network Support Staff Knowledge (Good-Poor)	Good	Good	Fair

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—1-13

The figure represents a simple comparison of three IP routing protocols. The remainder of this course consists of technical detail on each of these, as well as BGP, IP multicast, and IP version 6 (IPv6).

Summary

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

Summary

- **Cisco Enterprise Architectures with a hierarchical network model facilitate the deployment of converged networks.**
- **The Cisco Enterprise Composite Network Model defines three functional areas: Enterprise Campus, Enterprise Edge, and Service Provider Edge.**
- **Converged networks with their traffic mix have higher demands on the network and its resources.**
- **The SONA framework guides the evolution of the enterprise network toward the IIN.**
- **The network models can be important tools for selecting and implementing an advanced IP routing protocol.**

© 2006 Cisco Systems, Inc. All rights reserved. BSCI v3.0—1-14

References

For additional information, refer to these resources:

- Cisco Systems, Inc. The Intelligent Information Network.
<http://www.cisco.com/go/iin>
- Cisco Systems, Inc. Service-Oriented Network Architecture.
<http://www.cisco.com/go/sona>

Configuring EIGRP

Overview

In routing environments, Enhanced Interior Gateway Routing Protocol (EIGRP) offers benefits and features over historical distance-vector routing protocols such as Routing Information Protocol version 1 (RIPv1) and Interior Gateway Routing Protocol (IGRP). These benefits include rapid convergence, lower bandwidth utilization, and multiple routed protocol support (for IP, as well as for Internetwork Packet Exchange [IPX] and AppleTalk).

This module describes how EIGRP works and how to implement and verify EIGRP operations. Advanced topics including route summarization, load balancing, EIGRP bandwidth usage, and authentication are also explored. The module concludes with a discussion of EIGRP issues and problems and how to correct them.

Module Objectives

Upon completing this module, you will be able to implement and verify EIGRP operations. This ability includes being able to meet these objectives:

- Explain how EIGRP selects routes between routers in diverse, large-scale internetworks
- Describe how to implement EIGRP routing
- Configure advanced EIGRP features for scalable networks
- Implement authentication in an EIGRP network
- Describe, recognize, and correct common EIGRP issues and problems

Lesson 1

Introducing EIGRP

Overview

To select the appropriate routing protocols for an internetwork, you must understand the key features and terminology that are necessary to evaluate a given protocol against other choices. Routing protocols are distinguished by the way that they select the best pathway and the way that they calculate the routing protocol metric.

This lesson reviews the benefits of Enhanced Interior Gateway Routing Protocol (EIGRP) and discusses the four underlying technologies within EIGRP. The three tables that EIGRP uses in the path selection process are also described, and the EIGRP metric calculation is explored in detail.


Objectives

Upon completing this lesson, you will be able to explain how EIGRP selects routes between routers in diverse, large-scale internetworks. This ability includes being able to meet these objectives:

- Describe the key capabilities that distinguish EIRGP from other routing protocols
- Identify the four key technologies employed by EIGRP
- Describe how EIGRP operates
- Describe the five components of the metric used by EIGRP
- Calculate the EIGRP metric for a range of pathways between routers
- Explain how IGRP routes are integrated into EIGRP routes and vice-versa

EIGRP Capabilities and Attributes

Key capabilities that distinguish EIGRP from other routing protocols include fast convergence, support for variable-length subnet masking (VLSM), support for partial updates, and support for multiple network layer protocols. This topic describes these capabilities.



EIGRP is a Cisco proprietary protocol that combines the advantages of link-state and distance vector routing protocols. EIGRP has its roots as a distance vector routing protocol and is predictable in its behavior. Like its predecessor, Interior Gateway Routing Protocol (IGRP), EIGRP is easy to configure and is adaptable to a wide variety of network topologies. The addition of several link-state features, such as dynamic neighbor discovery, makes EIGRP an advanced distance vector protocol. While EIGRP is compatible with existing IGRP networks, EIGRP is an *enhanced* IGRP because of its rapid convergence and the guarantee of a loop-free topology at all times. A hybrid protocol, EIGRP uses the Diffusing Update Algorithm (DUAL) and includes the following key features:

- **Fast convergence:** A router running EIGRP stores all its neighbors' routing tables so that it can quickly adapt to alternate routes. If no appropriate route exists, EIGRP queries its neighbors to discover an alternate route. These queries propagate until an alternate route is found.
- **VLSM support:** EIGRP is a classless routing protocol, which means that it advertises a subnet mask for each destination network; this structure enables EIGRP to support discontinuous subnetworks and VLSM. With EIGRP, routes are automatically summarized at the major network number boundary, but EIGRP can be configured to summarize on any bit boundary on any router interface.

- **Partial updates:** EIGRP does not send periodic updates. Instead, it sends partial triggered updates; these are sent only when the path or the metric changes for a route, and they contain information about the changed routes only. Propagation of partial updates is automatically bounded so that only those routers that need the information are updated. As a result of these two capabilities, EIGRP consumes significantly less bandwidth than IGRP. This behavior is different from that of link-state protocols, in which an update is transmitted to *all* link-state routers within an area.
- **Multiple network-layer protocol support:** EIGRP supports IP, AppleTalk, and Novell NetWare Internet Packet Exchange (IPX) through the use of protocol-dependent modules. These modules are responsible for protocol requirements specific to the network layer. The rapid convergence and sophisticated metric structure of EIGRP offers superior performance and stability when implemented in IPX and AppleTalk networks.

Note This course covers only the TCP/IP implementation of EIGRP.

Other EIGRP features include the following:

- **Seamless connectivity across all data link layer protocols and topologies:** EIGRP does not require special configuration to work across any Layer 2 protocols. Other routing protocols, such as Open Shortest Path First (OSPF), use different configurations for different Layer 2 protocols, such as Ethernet and Frame Relay. EIGRP operates effectively in both LAN and WAN environments. WAN support for dedicated point-to-point links and nonbroadcast multiaccess (NBMA) topologies is standard for EIGRP. EIGRP accommodates differences in media types and speeds when neighbor adjacencies form across WAN links and can be configured to limit the amount of bandwidth that the protocol uses on WAN links.
- **Sophisticated metric:** EIGRP uses the same algorithm for metric calculation as IGRP but represents values in 32-bit format to give additional granularity. EIGRP supports unequal metric load balancing, which allows administrators to better distribute traffic flow in their networks.
- **Multicast and unicast:** EIGRP uses multicast and unicast, rather than broadcast. The multicast address used for EIGRP is 224.0.0.10.

Underlying Processes and Technologies

EIGRP employs four key technologies that combine to differentiate it from other routing technologies: neighbor discovery/recovery, reliable transport protocol (RTP), DUAL finite-state machine, and protocol-dependent modules. This topic describes these technologies.

EIGRP Key Technologies

- **Neighbor discovery/recovery**
 - Uses hello packets between neighbors
- **Reliable Transport Protocol (RTP)**
 - Guaranteed, ordered delivery of EIGRP packets to all neighbors
- **DUAL finite-state machine**
 - Selects lowest-cost, loop free, paths to each destination
- **Protocol-dependent modules (PDMs)**
 - EIGRP supports IP, AppleTalk, and Novell NetWare.
 - Each protocol has its own EIGRP module and operates independently of any of the others that may be running.

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—2-3

The four key technologies are described as follows:

- **Neighbor discovery/recovery mechanism:** Enables routers to dynamically learn about other routers on their directly attached networks. Routers also must discover when their neighbors become unreachable or inoperative. This process is achieved with low overhead by periodically sending small hello packets. As long as a router receives hello packets from a neighboring router, it assumes that the neighbor is functioning and the two can exchange routing information.
- **RTP:** Responsible for guaranteed, ordered delivery of EIGRP packets to all neighbors. It supports intermixed transmission of multicast or unicast packets. For efficiency, only certain EIGRP packets are transmitted reliably.

For example, on a multiaccess network that has multicast capabilities, such as Ethernet, it is not necessary to send hello packets reliably to all neighbors individually, so EIGRP sends a single multicast hello packet containing an indicator that informs the receivers that the packet need not be acknowledged. Other types of packets, such as updates, indicate in the packet that acknowledgment is required. RTP contains a provision for sending multicast packets quickly even when unacknowledged packets are pending, which helps ensure that convergence time remains low in the presence of links of varying speeds.

- **DUAL finite state machine:** Embodies the decision process for all route computations. DUAL tracks all routes advertised by all neighbors and uses distance information, known as a metric or cost, to select efficient, loop-free paths to all destinations.

- **Protocol-dependent modules (PDMs):** Responsible for network layer protocol-specific requirements. EIGRP supports IP, AppleTalk, and Novell NetWare; each protocol has its own EIGRP module and operates independently from any of the others that may be running. The IP-EIGRP module, for example, is responsible for sending and receiving EIGRP packets that are encapsulated in IP. IP-EIGRP is also responsible for parsing EIGRP packets and informing DUAL of the new information that has been received. IP-EIGRP asks DUAL to make routing decisions, the results of which are stored in the IP routing table. IP-EIGRP is responsible for redistributing routes learned by other IP routing protocols.

EIGRP Operation

EIGRP uses the neighbor table to list adjacent routers. The topology table lists all the learned routes to each destination, while the routing table contains the best route to each destination; this best route is called the successor route. A feasible successor route is a backup route to a destination, which is kept in the topology table. This topic describes how EIGRP uses these tables and routes in its operation.

EIGRP Neighbor Table

IP EIGRP Neighbor Table	
Next-Hop Router	Interface

List of directly connected routers running EIGRP with which this router has an adjacency

3145, 123

© 2006 Cisco Systems, Inc. All rights reserved. BSCI v3.0--2-4

When a router discovers and forms an adjacency with a new neighbor, it records the neighbor's address and the interface through which it can be reached as an entry in the neighbor table. One neighbor table exists for each PDM. The EIGRP neighbor table is comparable to the adjacencies database that link-state routing protocols use and serves the same purpose: to ensure bidirectional communication between each of the directly connected neighbors.

When a neighbor sends a hello packet, it advertises a hold time, which is the amount of time that a router treats a neighbor as reachable and operational. If a hello packet is not received within the hold time, the hold time expires, and DUAL is informed of the topology change.

The neighbor-table entry also includes information required by RTP. Sequence numbers are employed to match acknowledgments with data packets, and the last sequence number received from the neighbor is recorded so that out-of-order packets can be detected. A transmission list is used to queue packets for possible retransmission on a per-neighbor basis. Round-trip timers are kept in the neighbor-table entry to estimate an optimal retransmission interval.

DUAL Terminology

- **Selects lowest-cost, loop-free paths to each destination**
- **AD = cost between the next-hop router and the destination**
- **FD = cost from local router = AD of next-hop router + cost between the local router and the next-hop router**
- **Lowest-cost = lowest FD**
- **(Current) successor = next-hop router with lowest-cost, loop free path**
- **Feasible successor = backup router with loop-free path (AD of feasible successor must be less than FD of current successor route)**

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—2-9

DUAL uses distance information, known as a metric or cost, to select efficient, loop-free paths.


The lowest-cost route is calculated by adding the cost between the next-hop router and the destination—referred to as the advertised distance (AD)—to the cost between the local router and the next-hop router. The sum of these costs is referred to as the feasible distance (FD).

A successor, also called a current successor, is a neighboring router that has a least-cost path to a destination (the lowest FD) that is guaranteed not to be part of a routing loop; successors are used for forwarding packets. Multiple successors can exist if they have the same FD. By default, up to four successors can be added to the routing table (the router can be configured to accept up to six per destination).

As well as keeping least-cost paths, DUAL keeps backup paths to each destination. The next-hop router for a backup path is called the feasible successor. To qualify as a feasible successor, a next-hop router must have an AD less than the FD of the current successor route.

If the route via the successor becomes invalid (because of a topology change) or if a neighbor changes the metric, DUAL checks for feasible successors to the destination route. If one is found, DUAL uses it, avoiding the need to recompute the route. If no suitable feasible successor exists, a recomputation must occur to determine the new successor. Although recomputation is not processor-intensive, it does affect convergence time, so it is advantageous to avoid unnecessary recomputations.

EIGRP Topology Table



When the router dynamically discovers a new neighbor, it sends an update about the routes that it knows to its new neighbor and receives the same information from the new neighbor. These updates populate the topology table. The topology table contains all destinations advertised by neighboring routers. It is important to note that if a neighbor is advertising a destination, it must be using that route to forward packets; this rule must be strictly followed by all distance vector protocols.


The topology table also maintains the metric that each neighbor advertises for each destination (the AD) and the metric that this router would use to reach the destination via that neighbor (the FD). The FD is the cost of this router to reach the neighbor for this destination plus the neighbor's metric to reach the destination.

The topology table is updated when a directly connected route or interface changes or when a neighboring router reports a change to a route.

A topology-table entry for a destination can exist in one of two states: active or passive. A destination is in the *passive* state when the router is not performing a recomputation; it is in the *active* state when the router is performing a recomputation. If feasible successors are always available, a destination never has to go into the active state and avoids a recomputation. The desired state is passive state.

A recomputation occurs when a destination has no feasible successors. The router initiates the recomputation by sending a query packet to each of its neighboring routers. If the neighboring router has a route for the destination, it sends a reply packet; if it does not have a route, it sends a query packet to its neighbors. In this case, the route is also in the active state in the neighboring router. While a destination is in the active state, a router cannot change the destination's routing table information. After a router has received a reply from each neighboring router, the topology-table entry for the destination returns to the passive state, and the router can select a successor.

EIGRP IP Routing Table



A router compares all FDs to reach a specific network and then selects the route with the lowest FD and places it in the IP routing table; this is the successor route. The FD for the chosen route becomes the EIGRP routing metric to reach that network in the routing table.


Example: EIGRP Tables

Router C Tables:

IP EIGRP Neighbor Table	
Next-Hop Router	Interface
Router A	Ethernet 0
Router B	Ethernet 1

IP EIGRP Topology Table			
Network	Feasible Distance (EIGRP Metric)	Advertised Distance	EIGRP Neighbor
10.1.1.0 /24	2000	1000	Router A (E0)
10.1.1.0 /24	2500	1500	Router B (E1)

IP Routing Table			
Network	Metric (Feasible Distance)	Outbound Interface	Next Hop (EIGRP Neighbor)
10.1.1.0 /24	2000	Ethernet 0	Router A



© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0-2.8

Example: EIGRP Tables

The network shown in the figure is used to illustrate the EIGRP tables; the router C tables are displayed. Routers A and B have established a neighbor relationship with router C and have sent their routing tables to router C. Both routers A and B have paths to network 10.1.1.0/24, among many others that are not shown.

The routing table on router A has an EIGRP metric of 1000 for 10.1.1.0/24, so router A advertises 10.1.1.0/24 to router C with a metric of 1000. Router C installs the route to 10.1.1.0/24 via router A in its EIGRP topology table with an advertised distance of 1000.

Router B has network 10.1.1.0/24 with a metric of 1500 in its IP routing table, so router B advertises 10.1.1.0/24 to router C with an advertised distance of 1500. Router C places the route to 10.1.1.0/24 network via router B in the EIGRP topology table with an advertised distance of 1500.

Therefore, router C has two entries to reach 10.1.1.0/24 in its topology table. The EIGRP metric for router C to reach both routers A and B is 1000. This cost (1000) is added to the respective advertised distance from each router, resulting in the FDs from router C to reach network 10.1.1.0/24 shown in the figure.

Router C chooses the least-cost FD, which is 2000, via router A, and installs it in the IP routing table as the best route to reach 10.1.1.0/24. The EIGRP metric in the routing table is equal to the FD from the EIGRP topology table. Router A is the successor for the route to 10.1.1.0/24.

EIGRP Packets


- **Hello:** Establish neighbor relationships.
- **Update:** Send routing updates.
- **Query:** Ask neighbors about routing information.
- **Reply:** Respond to query about routing information.
- **ACK:** Acknowledge a reliable packet.

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—2.9

EIGRP uses the following five generic packet types:

- **Hello:** Routers use hello packets for neighbor discovery. The packets are sent as multicasts and do not require an acknowledgment.
- **Update:** Update packets contain route change information. They are sent reliably to the affected routers only. These updates can be unicast to a specific router or multicast to multiple attached routers.
- **Query:** When a router performs route computation and does not have a feasible successor, it sends a reliable query packet to its neighbors to determine if they have a feasible successor for the destination. Queries are normally multicast but can be retransmitted as unicast packets in certain cases.
- **Reply:** A router sends a reply packet in response to a query packet. Replies are unicast reliably to the originator of the query.
- **ACK:** The acknowledgment (ACK) packet acknowledges update, query, and reply packets. ACK packets are unicast hello packets and contain a nonzero acknowledgment number.



The process to establish and discover neighbor routes occurs simultaneously in EIGRP. A high-level description of the process is as follows, using the topology in the figure as an example:

1. A new router (router A) comes up on the link and sends a hello packet through all of its EIGRP-configured interfaces.
2. Routers receiving the hello packet (router B) on one interface reply with update packets that contain all the routes they have in their routing tables, except those learned through that interface (split horizon). Router B sends an update packet to router A, but a neighbor relationship is not established until router B sends a hello packet to router A. The update packet from router B has the initialization bit set, indicating that this is the initialization process. The update packet includes information about the routes that the neighbor (router B) is aware of, including the metric that the neighbor is advertising for each destination.
3. After both routers have exchanged hellos and the neighbor adjacency is established, router A replies to router B with an ACK packet, indicating that it received the update information.
4. Router A assimilates all update packets in its topology table. The topology table includes all destinations advertised by neighboring (adjacent) routers. It lists each destination, all the neighbors that can reach the destination, and their associated metric.
5. Router A then sends an update packet to router B.
6. Upon receiving the update packet, router B sends an ACK packet to router A.

After router A and router B successfully receive the update packets from each other, they are ready to update their routing tables with the successor routes from the topology table.

EIGRP Metric

EIGRP uses the same metric components as IGRP: delay, bandwidth, reliability, load, and maximum transmission unit (MTU), as described in this topic.

EIGRP Metric

- **Same metric components as IGRP:**
 - **Bandwidth**
 - **Delay**
 - **Reliability**
 - **Loading**
 - **MTU**
- **EIGRP metric is IGRP metric multiplied by 256.**

© 2006 Cisco Systems, Inc. All rights reserved. BSCI v3.0--2.11

EIGRP uses the same composite metric as IGRP to determine the best path, except that the EIGRP metric is multiplied by 256. The metric can be based on five criteria, but EIGRP uses only two of these criteria by default:

- **Bandwidth:** The smallest bandwidth between source and destination
- **Delay:** The cumulative interface delay along the path

The following criteria can be used, but are not recommended, because they typically result in frequent recalculation of the topology table:

- **Reliability:** This value represents the worst reliability between source and destination, based on keepalives.
- **Loading:** This value represents the worst load on a link between source and destination, computed based on the packet rate and the configured bandwidth of the interface.
- **MTU:** This criterion represents the smallest MTU in the path. MTU is included in the EIGRP routing update but is not actually used in the metric calculation.

Calculating the EIGRP Metric

This topic describes how the EIGRP metric is calculated.

EIGRP Metric Calculation

- **By default, EIGRP metric:**
 $\text{Metric} = \text{bandwidth (slowest link)} + \text{delay (sum of delays)}$
- **Delay = sum of the delays in the path, in tens of microseconds, multiplied by 256**
- **Bandwidth = $[10^7 / (\text{minimum bandwidth link along the path, in kilobits per second})] * 256$**
- **Formula with default K values (K1 = 1, K2 = 0, K3 = 1, K4 = 0, K5 = 0):**
 $\text{Metric} = [K1 * BW + ((K2 * BW) / (256 - \text{load})) + K3 * \text{delay}]$
- **If K5 not equal to 0:**
 $\text{Metric} = \text{metric} * [K5 / (\text{reliability} + K4)]$

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—2-12

EIGRP calculates the metric by adding together weighted values of different variables of the link to the network in question. The default constant weight values are K1 = K3 = 1, and K2 = K4 = K5 = 0.

In EIGRP metric calculations, when K5 is 0 (the default), variables (bandwidth, bandwidth divided by load, and delay) are weighted with the constants K1, K2, and K3. The following is the formula used:

$$\blacksquare \text{ Metric} = (K1 * \text{bandwidth}) + [(K2 * \text{bandwidth}) / (256 - \text{load})] + (K3 * \text{delay})$$

If these K values are equal to their defaults, the formula becomes the following:

$$\blacksquare \text{ Metric} = (1 * \text{bandwidth}) + [(0 * \text{bandwidth}) / (256 - \text{load})] + (1 * \text{delay})$$
$$\blacksquare \text{ Metric} = \text{bandwidth} + \text{delay}$$


If K5 is not equal to 0, the following additional operation is performed:

$$\blacksquare \text{ Metric} = \text{metric} * [K5 / (\text{reliability} + K4)]$$

K values are carried in EIGRP hello packets. Mismatched K values can cause a neighbor to be reset. (Only K1 and K3 are used, by default, in metric compilation.) These K values should be modified only after careful planning; changing these values can prevent your network from converging and is generally not recommended.

The format of the delay and bandwidth values used for EIGRP metric calculations is different from those displayed by the **show interface** command. The EIGRP delay value is the sum of the delays in the path, in tens of microseconds, multiplied by 256. The **show interface** command displays delay in microseconds. The EIGRP bandwidth is calculated using the minimum bandwidth link along the path, in kilobits per second. The value 10^7 is divided by this value, and then the result is multiplied by 256.

EIGRP Metrics Calculation Example



A → B → C → D	Least bandwidth 64 kbps	Total delay 6,000
A → X → Y → Z → D	Least bandwidth 256 kbps	Total delay 8,000

- **Delay is the sum of all the delays of the links along the paths:**
 $\text{Delay} = [\text{delay in tens of microseconds}] \times 256$
- **Bandwidth is the lowest bandwidth of the links along the paths:**
 $\text{Bandwidth} = [10,000,000 / (\text{bandwidth in kbps})] \times 256$

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—2-13

Example: EIGRP Metric Calculation

The figure illustrates an example network in which router A has two paths to reach networks behind router D. The bandwidths (in kilobits per second) and the delays (in tens of microseconds) of the various links are also shown in the figure.

The least bandwidth along the top path (A → B → C → D) is 64 kbps. The EIGRP bandwidth calculation for this path is as follows:

- $\text{Bandwidth} = (10^7 / \text{least bandwidth in kbps}) \times 256$
- $\text{Bandwidth} = (10,000,000 / 64) \times 256 = 156,250 \times 256 = 40,000,000$

The delay through the top path is as follows:

- $\text{Delay} = [(\text{delay A} \rightarrow \text{B}) + (\text{delay B} \rightarrow \text{C}) + (\text{delay C} \rightarrow \text{D})] \times 256$
- $\text{Delay} = [2000 + 2000 + 2000] \times 256$
- $\text{Delay} = 1,536,000$

Therefore, the EIGRP metric calculation for the top path is as follows:

- $\text{Metric} = \text{bandwidth} + \text{delay}$
- $\text{Metric} = 40,000,000 + 1,536,000$
- $\text{Metric} = 41,536,000$

The least bandwidth along the lower path (A → X → Y → Z → D) is 256 kbps. The EIGRP bandwidth calculation for this path is as follows:

- $\text{Bandwidth} = (10^7 / \text{least bandwidth in kbps}) \times 256$
- $\text{Bandwidth} = (10,000,000 / 256) \times 256 = 10,000,000$

The delay through the lower path is as follows:

- Delay = [(delay A → X) + (delay X → Y) + (delay Y → Z) + (delay Z → D)] * 256
- Delay = [2000 + 2000 + 2000 + 2000] * 256
- Delay = 2,048,000

Therefore, the EIGRP metric calculation for the lower path is as follows:


- Metric = bandwidth + delay
- Metric = 10,000,000 + 2,048,000
- Metric = 12,048,000

Router A therefore chooses the lower path, with a metric of 12,048,000, over the top path, with a metric of 41,536,000. Router A installs the lower path, with a next-hop router of X and a metric of 12,048,000, in the IP routing table.

The bottleneck along the top path, the 64-kbps link, can explain why the router takes the lower path. This slow link means that the rate of transfer to router D would be at a maximum of 64 kbps. Along the lower path, the lowest speed is 256 kbps, making the throughput rate up to that speed. Therefore, the lower path represents a better choice, for example, to move large files quickly.

Integrating the EIGRP and IGRP Routes

EIGRP represents its metrics in a 32-bit format, versus the 24-bit representation used by IGRP. This topic explains how IGRP routes are integrated into EIGRP routes and vice-versa.



The EIGRP 32-bit metric representation allows for a more granular decision than the IGRP 24-bit representation to calculate the best routes. The EIGRP metric value ranges from 1 to 4,294,967,296. The IGRP metric value ranges from 1 to 16,777,216.

As shown in the figure, EIGRP metrics are backward-compatible with IGRP. When integrating IGRP routes into an EIGRP domain using redistribution, the router multiplies the IGRP metric by 256 to compute the EIGRP-equivalent metric. When sending EIGRP routes to an IGRP routing domain, the router divides each EIGRP metric by 256 to achieve the equivalent IGRP metric.

Note IGRP is no longer supported, as of Cisco IOS Software Release 12.3.

Summary

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

Summary

- EIGRP capabilities include fast convergence and support for VLSM, partial updates, and multiple network layer protocols.
- EIGRP key technologies are neighbor discovery/recovery, RTP, DUAL finite-state machine, and PDMs.
- EIGRP uses three tables: neighbor table, topology table, and routing table. The routing table contains the best route to each destination, called the successor route. A feasible successor route is a backup route to a destination; it is kept in the topology table.
- EIGRP uses the same metric components as IGRP: delay, bandwidth, reliability, load, and MTU.
- By default, EIGRP metric equals bandwidth (slowest link) plus delay (sum of delays).
- EIGRP metrics are backward-compatible with IGRP; the EIGRP-equivalent metric is the IGRP metric multiplied by 256.

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—2-15

Lesson 2

Implementing and Verifying EIGRP

Overview

This lesson explains and demonstrates how to implement and verify basic Enhanced Interior Gateway Routing Protocol (EIGRP). This lesson also explains how to use the wildcard network mask option and how to use a default network with EIGRP. To assist in troubleshooting, this lesson introduces various Cisco IOS software **show** commands and defines the key fields in each.

Knowing the correct commands to use when you configure EIGRP helps to ensure that migration to this routing protocol is smooth and quick. Understanding which **show** command to use when troubleshooting the EIGRP configuration saves valuable time.

Objectives

Upon completing this lesson, you will be able to describe how to implement EIGRP routing. This ability includes being able to meet these objectives:

- Describe the commands used in a basic EIGRP configuration task
- Explain how to configure a router to use wildcard masks to select the interfaces and networks that will participate in EIGRP routing
- Configure the last-resort gateway or default route
- Verify that the router recognizes EIGRP neighbors and routes
- Verify EIGRP operations

Configuring Basic EIGRP

This topic describes the commands used in a basic EIGRP configuration task.

Configuring EIGRP

```
Router(config)#  
router eigrp autonomous-system-number
```

- Defines EIGRP as the IP routing protocol.
- All routers in the internetwork that must exchange EIGRP routing updates must have the same autonomous system number.

```
Router(config-router)#  
network network-number [wildcard-mask]
```

- Identifies attached networks participating in EIGRP.
- The *wildcard-mask* is an inverse mask used to determine how to interpret the address. The mask has wildcard bits, where 0 is a match and 1 is “don’t care.”

© 2006 Cisco Systems, Inc. All rights reserved. BSCI v3.0—2.2

To configure basic EIGRP for IP, perform the following steps:

- Step 1** Enable EIGRP and define the autonomous system (AS) using the **router eigrp** *autonomous-system-number* command. The AS number value must match on all routers within the AS.
- Step 2** Indicate which networks are part of the EIGRP AS using the **network** command. This command determines which interfaces of the router are participating in EIGRP and which networks the router advertises. The table lists the parameters for the **network** command.

network Command Parameters

Command	Description
<i>network-number</i>	The network, subnet, or the address of a directly connected interface. This parameter instructs the router to recognize which links to advertise to, which links to listen to advertisements on, and which networks to advertise.
<i>wildcard-mask</i>	(Optional) An inverse mask used to determine how to interpret the network number. The mask has wildcard bits, where 0 is a match and 1 is don't care. For example, 0.0.255.255 indicates a match in the first two octets.

If you do not use the optional wildcard mask, the EIGRP process assumes that all directly connected networks that are part of the overall major network will participate in the EIGRP routing process, and EIGRP will attempt to establish EIGRP neighbor relationships from each interface that is part of that Class A, B, or C major network.

Use the optional wildcard mask to identify a specific IP address, subnet, or network. The router interprets the network number using the wildcard mask to determine which connected networks will participate in the EIGRP routing process. If specifying an interface address, use the mask 0.0.0.0 to match all four octets of the address. An address and wildcard mask combination of 0.0.0.0 255.255.255.255 matches all interfaces on the router.

Configuring EIGRP (Cont.)

Router(config-if) #

bandwidth kilobits

- **Defines the interface's bandwidth for the purposes of sending routing update traffic.**

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—2-3


- Step 3** If you are using serial links, define the bandwidth of the link for the purposes of sending routing update traffic, using the **bandwidth kilobits** command. In this command, *kilobits* indicates the intended bandwidth in kilobits per second. For example, for a 64-kbps link, use the following command:

```
router(config-if)#bandwidth 64
```

If you do not change the bandwidth for these interfaces, EIGRP assumes that the bandwidth on the link is the default of T1 speed. If the link is actually slower, the router might not be able to converge, or routing updates might be lost.

For generic serial interfaces such as PPP or High-Level Data Link Control (HDLC), set the bandwidth to the line speed. For Frame Relay on point-to-point interfaces, set the bandwidth to the committed information rate (CIR). For Frame Relay multipoint connections, set the bandwidth to the sum of all CIRs, or if the permanent virtual circuits (PVCs) have different CIRs, then set the bandwidth to the lowest CIR multiplied by the number of PVCs on the multipoint connection.

Configuring EIGRP for IP



Network 192.168.1.0 is not configured on router A, because it is not directly connected to router A.

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—2-4

Example: Configuring EIGRP for IP

The figure illustrates the configuration of router A for EIGRP. Router A, along with all routers in the figure, is part of EIGRP AS 109. For EIGRP to establish a neighbor relationship, all neighbors must be in the same AS.

Because the wildcard mask is not used in the router A configuration, all interfaces on router A that are part of network 10.0.0.0/8 and network 172.16.0.0/16 participate in the EIGRP routing process. In this case, this includes all four interfaces. Note that network 192.168.1.0 is not configured in the EIGRP configuration on router A, because router A does not have any interfaces in that network.

In this example, consider what would happen if the following configuration were entered on router A:

```
routerA(config)#router eigrp 109
routerA(config-router)#network 10.1.0.0
routerA(config-router)#network 10.4.0.0
routerA(config-router)#network 172.16.7.0
routerA(config-router)#network 172.16.2.0
```

Router A would change the **network** commands to have classful networks, and the resulting configuration would be the following:

```
router eigrp 109
network 10.0.0.0
network 172.16.0.0
```


Alternatively, consider what would happen if the following configuration were entered on router A:

```
routerA(config)#router eigrp 109
routerA(config-router)#network 10.1.0.0 0.0.255.255
routerA(config-router)#network 10.4.0.0 0.0.255.255
routerA(config-router)#network 172.16.2.0 0.0.0.255
routerA(config-router)#network 172.16.7.0 0.0.0.255
```

In this case, router A uses the wildcard mask to determine which directly connected interfaces participate in the EIGRP routing process for AS 109. All interfaces that are part of networks 10.1.0.0/16, 10.4.0.0/16, 172.16.2.0/24, and 172.16.7.0/24 participate in the EIGRP routing process for AS 109; in other words, all four interfaces participate in EIGRP.

Using a Wildcard Mask in EIGRP

This topic examines how to configure a router running EIGRP with a wildcard mask to select the networks or interfaces that will participate in the EIGRP.



Example: Wildcard Mask in EIGRP

An example of when wildcard masks are useful is when a router in an AS connects to a router external to its AS. In this case, the router can be configured with a wildcard mask so the router does not try to form an adjacency with the router in the other AS.


For example, router C in the figure includes subnets of the Class B network 172.16.0.0 on all interfaces. The router C configuration in the figure uses the wildcard mask, because router C connects to a router external to AS 100 on its serial interface, and EIGRP with AS 100 should not be run there.

Without the wildcard mask, router C would send EIGRP packets to the external network, which would waste bandwidth and CPU cycles and would provide unnecessary information to the external network.

The wildcard mask tells EIGRP to establish a relationship with EIGRP routers from interfaces that are part of subnets 172.16.3.0/24 or 172.16.4.0/24, not 172.16.5.0/24.

Configuring the ip default-network Command

This topic describes the **ip default-network *network-number*** command that is used to configure the last-resort gateway or default route.



The EIGRP default route can be created with the **ip default-network *network-number*** command. A router configured with this command considers the network listed in the command as the last-resort gateway that it will announce to other routers. The table describes the parameter of the **ip default-network** command.

ip default-network Command Parameter

Command	Description
<i>network-number</i>	The classful destination network.

The network specified by this command must be reachable by the router that uses this command before it announces it as a candidate default route to other EIGRP routers. The network specified by this command must also be passed to other EIGRP routers so that those routers can use this network as their default network and set the gateway of last resort to this default network. This requirement means that the network must either be an EIGRP-derived network in the routing table or be generated using a static route, which has been redistributed into EIGRP.

Multiple default networks can be configured; downstream routers use the EIGRP metric to determine the best default route.

Example: ip default-network Command

For example, in the figure, router A is directly attached to external network 172.31.0.0/16. Router A is configured with the 172.31.0.0 network as a candidate default network using the **ip default-network 172.31.0.0** command.

This network is passed to router B because router A has it listed in a **network** command under the EIGRP process. Notice that the routing table for router A does not set the gateway of last resort; the **ip default-network** command does not benefit router A directly. On router B, the EIGRP-learned 172.31.0.0 network is flagged as a candidate default network (as indicated by the asterisk (*) in the routing table). Router B also sets the gateway of last resort to 10.64.0.2 (router A) to reach the default network of 172.31.0.0.


Note	When you configure the ip default-network command, a static route (the ip route command) is generated in the router configuration; however, the Cisco IOS software does not display a message to indicate this. The entry appears as a static route in the routing table of the router in which the command is configured, as can be seen in the router A configuration and routing table in the figure, which can be confusing if you want to remove the default network. The configuration must be removed with the no ip route command.

EIGRP and Interior Gateway Routing Protocol (IGRP) behave differently from Routing Information Protocol (RIP) when you are using the **ip route 0.0.0.0 0.0.0.0** command. For example, EIGRP does not redistribute the 0.0.0.0 0.0.0.0 default route by default. However, if the **network 0.0.0.0** command is added to the EIGRP configuration, it redistributes a default route as a result of the **ip route 0.0.0.0 0.0.0.0 interface** command (but not as a result of the **ip route 0.0.0.0 0.0.0.0 address** or **ip default-network** command). For example, the following configuration results in the 0.0.0.0 route being passed to the EIGRP neighbors of the router:

```
interface serial 0/0/0
  ip address 10.1.1.1 255.255.255.0
!
ip route 0.0.0.0 0.0.0.0 serial 0/0/0
!
router eigrp 100
  network 0.0.0.0
```

Verify EIGRP IP Routes

Use the **show ip eigrp neighbors** to verify that the router recognizes its neighbors. Use the **show ip route eigrp** command to verify that the router recognizes routes from its neighbors. This topic describes these commands.



Example: EIGRP Configuration

The figure shows a network used to illustrate the configuration, verification, and troubleshooting of EIGRP in this topic and the following topics. The configuration of the R1 router is also shown in this figure.

EIGRP is enabled in AS 100. The **network 172.16.1.0 0.0.0.255** command starts EIGRP on the Fast Ethernet 0/0 interface and allows router R1 to advertise this network. With the wildcard mask used, this command specifies that only interfaces on the 172.16.1.0/24 subnet will participate in EIGRP. Note, however, that the full Class B network 172.16.0.0 will be advertised, because EIGRP automatically summarizes routes on the major network boundary by default. The **network 192.168.1.0** command starts EIGRP on the Serial 0/0/1 interface, and allows router R1 to advertise this network.

R2 EIGRP Configuration

```
<output omitted>
interface FastEthernet0/0
  ip address 172.17.2.2 255.255.255.0

<output omitted>
interface Serial0/0/1
  bandwidth 64
  ip address 192.168.1.102 255.255.255.224

<output omitted>
router eigrp 100
 network 172.17.2.0 0.0.0.255
 network 192.168.1.0
```

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—2-8

Example: R2 EIGRP Configuration

The configuration of the R2 router is shown in this figure.

EIGRP is enabled in AS 100. The **network 172.17.2.0 0.0.0.255** command starts EIGRP on the Fast Ethernet 0/0 interface and allows router R2 to advertise this network. With the wildcard mask used, this command specifies that only interfaces on the 172.17.2.0/24 subnet will participate in EIGRP. Note, however, that the full Class B network 172.17.0.0 will be advertised, because EIGRP automatically summarizes routes on the major network boundary by default. The **network 192.168.1.0** command starts EIGRP on the serial 0/0/1 interface and allows router R2 to advertise this network.

Verifying EIGRP: show ip eigrp neighbors

```
R1#show ip eigrp neighbors
IP-EIGRP neighbors for process 100
H   Address          Interface    Hold   Uptime    SRTT    RTO   Q   Seq
                               (sec)                (ms)    Cnt  Num
0   192.168.1.102     Se0/0/1     10     00:07:22   10      2280  0   5
R1#
```

The EIGRP IP neighbor table can be displayed with the **show ip eigrp neighbors** command, as shown in the figure. This output table includes the following key elements:

- **H (handle):** A number used internally by the Cisco IOS software to track a neighbor.
- **Address:** The network-layer address of the neighbor.
- **Interface:** The interface on this router through which the neighbor can be reached.
- **Hold Time:** The maximum time, in seconds, that the router waits to hear from the neighbor without receiving anything from a neighbor before considering the link unavailable. Originally, the expected packet was a hello packet, but in current Cisco IOS software releases, any EIGRP packets received after the first hello from that neighbor reset the timer.
- **Uptime:** The elapsed time, in hours, minutes, and seconds, since the local router first heard from this neighbor.
- **Smoothed round trip time (SRTT):** The average number of milliseconds it takes for an EIGRP packet to be sent to this neighbor and for the local router to receive an acknowledgment of that packet. This timer is used to determine the retransmit interval, also known as the retransmission timeout (RTO).
- **RTO:** The amount of time, in milliseconds, that the router waits for an acknowledgment before retransmitting a reliable packet from the retransmission queue to a neighbor.
- **Queue count:** The number of packets waiting in the queue to be sent out. If this value is constantly higher than 0, a congestion problem might exist. A value of 0 indicates that no EIGRP packets are in the queue.
- **Seq Num:** The sequence number of the last update, query, or reply packet that was received from this neighbor.

Verifying EIGRP: show ip route eigrp

```
R1#show ip route eigrp
D    172.17.0.0/16 [90/40514560] via 192.168.1.102, 00:07:01, Serial0/0/1
    172.16.0.0/16 is variably subnetted, 2 subnets, 2 masks
D    172.16.0.0/16 is a summary, 00:05:13, Null0
    192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks
D    192.168.1.0/24 is a summary, 00:05:13, Null0

R1#show ip route
<output omitted>
Gateway of last resort is not set
D    172.17.0.0/16 [90/40514560] via 192.168.1.102, 00:06:55, Serial0/0/1
    172.16.0.0/16 is variably subnetted, 2 subnets, 2 masks
D    172.16.0.0/16 is a summary, 00:05:07, Null0
C    172.16.1.0/24 is directly connected, FastEthernet0/0
    192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks
C    192.168.1.96/27 is directly connected, Serial0/0/1
D    192.168.1.0/24 is a summary, 00:05:07, Null0
```

To verify that the router recognizes EIGRP routes for any neighbors, use the **show ip route eigrp** command, as shown in the figure. The figure also exhibits the **show ip route** command, which displays the full IP routing table, including the EIGRP routes.

EIGRP supports several route types: internal, external, and summary. EIGRP routes are identified with a D in the left column; any external EIGRP routes (from outside of this autonomous system) would be identified with a D EX.

After the network number, there is a field that looks similar to [90/40514560]. (The numbers may be different from the one in the example.) The second number in the brackets is the EIGRP metric; recall that the default EIGRP metric is the least-cost bandwidth plus the accumulated delays. The EIGRP metric for a network is the same as its feasible distance (FD) in the EIGRP topology table. The first number, 90 in this case, is the administrative distance; it is used to select the best path when a router learns two or more routes from different routing sources. For example, consider that this router also uses RIP, and RIP has a route to network 172.17.0.0 that is three hops away. The router, without administrative distance, cannot compare the three hops of RIP to an EIGRP metric of 40514560. The router does not know the bandwidth associated with hops, and EIGRP does not use hop count as a metric.

To correct this problem, Cisco Systems established an administrative distance value for each routing protocol; the lower the value, the more strongly preferred the route is. By default, EIGRP internal routes have an administrative distance of 90, and RIP has an administrative distance of 120. Because EIGRP has a metric based upon bandwidth and delays, it is preferred over the RIP hop count. As a result, in this example, the EIGRP route is installed in the routing table.

The next field, “via 192.168.1.102” in this example, identifies the address of the next-hop router to which this router passes the packets for the destination network 172.17.0.0/16. The next-hop address in the routing table is the same as the successor in the EIGRP topology table.

Each route also has a time associated with it: the length of time, perhaps days or months, since EIGRP last advertised this network to this router. EIGRP does not refresh routes periodically; it resends the routing information only when neighbor adjacencies change.

The next field in the output is the interface, serial 0/0/1 in this case, from which packets for 172.17.0.0 are sent.

Notice that the routing table includes routes, to null0, for the advertised routes. Cisco IOS software automatically puts these routes in the table; they are called summary routes. Null0 is a directly connected, software-only interface. The use of the null0 interface prevents the router from trying to forward traffic to other routers in search of a more precise, longer match. For example, if the router in the figure receives a packet to an unknown subnet that is part of the summarized range—172.16.3.5, for example—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, or sent to the *bit bucket*), which prevents the router from forwarding the packet to a default route and possibly creating a routing loop.

Verify EIGRP IP Operations

The commands used to verify EIGRP operations are the **show ip protocols**, **show ip eigrp interfaces**, **show ip eigrp neighbors**, **show ip eigrp topology**, and **show ip eigrp traffic** commands. This topic describes these commands, except for the **show ip eigrp neighbors** command, which was detailed in an earlier topic.

Verifying EIGRP: show ip protocols

```
R1#show ip protocols
Routing Protocol is "eigrp 100"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Default networks flagged in outgoing updates
  Default networks accepted from incoming updates
  EIGRP metric weight K1=1, K2=0, K3=1, K4=0, K5=0
  EIGRP maximum hopcount 100
  EIGRP maximum metric variance 1
  Redistributing: eigrp 100
  EIGRP NSF-aware route hold timer is 240s
<output omitted>

Maximum path: 4
Routing for Networks:
  172.16.1.0/24
  192.168.1.0
Routing Information Sources:
  Gateway         Distance      Last Update
  (this router)          90          00:09:38
  Gateway         Distance      Last Update
  192.168.1.102       90          00:09:40
Distance: internal 90 external 170
```

The **show ip protocols** command gives information about any and all dynamic routing protocols running on the router.

As shown in the partial output in the figure, when EIGRP is running, the **show ip protocols** command output displays any routing filtering occurring on EIGRP outbound or inbound updates. It also identifies whether EIGRP is generating a default network or receiving a default network in EIGRP updates.

The full output of this command is as follows:

```
R1#show ip protocols
Routing Protocol is "eigrp 100"

  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Default networks flagged in outgoing updates
  Default networks accepted from incoming updates
  EIGRP metric weight K1=1, K2=0, K3=1, K4=0, K5=0
  EIGRP maximum hopcount 100
  EIGRP maximum metric variance 1
  Redistributing: eigrp 100
```

```

EIGRP NSF-aware route hold timer is 240s
Automatic network summarization is in effect
Automatic address summarization:
  192.168.1.0/24 for FastEthernet0/0
    Summarizing with metric 40512000
  172.16.0.0/16 for Serial0/0/1
    Summarizing with metric 28160
Maximum path: 4
Routing for Networks:
  172.16.1.0/24
  192.168.1.0
Routing Information Sources:
  Gateway          Distance      Last Update
  (this router)          90          00:09:38
  Gateway          Distance      Last Update
  192.168.1.102        90          00:09:40
Distance: internal 90 external 170

```

The command output provides information about additional default settings for EIGRP, such as default K values, hop count, and variance.

Note	Because the routers must have identical K values for EIGRP to establish an adjacency, the show ip protocols command helps to determine the current K value setting before an adjacency is attempted.

This sample output also indicates that automatic summarization is enabled (this is the default) and that the router is allowed to load-balance over a maximum of four paths. (Cisco IOS software allows configuration of up to six paths for equal-cost load balancing, using the **maximum-path** command.)

The networks for which the router is routing are also displayed. As shown in the figure, the format of the output varies, depending on the use of the wildcard mask in the **network** command. If a wildcard mask is used, the network address is displayed with a prefix length. If a wildcard mask is not used, the Class A, B, or C major network is displayed.

The routing information sources portion of this command output identifies all other routers that have an EIGRP neighbor relationship with this router. The **show ip eigrp neighbors** command provides a detailed display of EIGRP neighbors.

The **show ip protocols** command output also provides the two administrative distances. First, an administrative distance of 90 applies to networks from other routers inside the AS; these are considered internal networks. Second, an administrative distance of 170 applies to networks introduced to EIGRP for this AS through redistribution; these are called external networks.

Verifying EIGRP: show ip eigrp interfaces

```
R1#show ip eigrp interfaces
IP-EIGRP interfaces for process 100
```

Interface	Peers	Xmit Queue Un/Reliable	Mean SRTT	Pacing Time Un/Reliable	Multicast Flow Timer	Pending Routes
Fa0/0	0	0/0	0	0/10	0	0
Se0/0/1	1	0/0	10	10/380	424	0

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0--2-12

The **show ip eigrp interfaces** command displays information about interfaces configured for EIGRP. This output includes the following key elements:

- **Interface:** Interface over which EIGRP is configured
- **Peers:** Number of directly connected EIGRP neighbors
- **Xmit Queue Un/Reliable:** Number of packets remaining in the Unreliable and Reliable transmit queues
- **Mean SRTT:** Mean SRTT interval, in milliseconds
- **Pacing Time Un/Reliable:** Pacing time used to determine when EIGRP packets should be sent out the interface (unreliable and reliable packets)
- **Multicast Flow Timer:** Maximum number of seconds in which the router will send multicast EIGRP packets
- **Pending Routes:** Number of routes in the packets in the transmit queue waiting to be sent

Verifying EIGRP: show ip eigrp topology

```
R1#show ip eigrp topology
IP-EIGRP Topology Table for AS(100)/ID(192.168.1.101)
Codes: P - Passive, A - Active, U - Update, Q - Query, R -
Reply,
       r - reply Status, s - sia Status
P 192.168.1.96/27, 1 successors, FD is 40512000
   via Connected, Serial0/0/1
P 192.168.1.0/24, 1 successors, FD is 40512000
   via Summary (40512000/0), Null0
P 172.16.0.0/16, 1 successors, FD is 28160
   via Summary (28160/0), Null0
P 172.16.1.0/24, 1 successors, FD is 28160
   via Connected, FastEthernet0/0
P 172.17.0.0/16, 1 successors, FD is 40514560
   via 192.168.1.102 (40514560/28160), Serial0/0/1
```

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—2-13

Another command used to verify EIGRP operations is the **show ip eigrp topology** command. For example, the figure illustrates that router R1 has an ID of 192.168.1.101 and is in AS 100—the EIGRP ID is the highest IP address on an active interface for this router.

As shown in the figure, this command output lists the networks known by this router through the EIGRP routing process. The codes in the command output are as follows:

- **Passive (P):** This network is available, and installation can occur in the routing table. Passive is the correct state for a stable network.
- **Active (A):** This network is currently unavailable, and installation cannot occur in the routing table. Being active means that there are outstanding queries for this network.
- **Update (U):** This code applies if a network is being updated (placed in an update packet). This code also applies if the router is waiting for an acknowledgment for this update packet.
- **Query (Q):** This code applies if there is an outstanding query packet for this network other than being in the active state. This code also applies if the router is waiting for an acknowledgment for a query packet.
- **Reply (R) status:** This code applies if the router is generating a reply for this network or is waiting for an acknowledgment for the reply packet.
- **Stuck-in-active (SIA) status:** This code signifies an EIGRP convergence problem for the network with which it is associated.

The number of successors available for a route is indicated in the command output. In this example, all networks have one successor. If there were equal-cost paths to the same network, a maximum of six paths would be shown. The number of successors corresponds to the number of best routes with equal cost.

For each network, the FD is displayed, followed by the next-hop address, which is followed by a field similar to (40514560/28160) in the figure. The first number in this field is the FD for that network through this next-hop router, and the second number is the advertised distance (AD) from the next-hop router to the destination network.

Verifying EIGRP: show ip eigrp traffic

```
R1#show ip eigrp traffic
IP-EIGRP Traffic Statistics for AS 100
  Hellos sent/received: 429/192
  Updates sent/received: 4/4
  Queries sent/received: 1/0
  Replies sent/received: 0/1
  Acks sent/received: 4/3
  Input queue high water mark 1, 0 drops
  SIA-Queries sent/received: 0/0
  SIA-Replies sent/received: 0/0
  Hello Process ID: 113
  PDM Process ID: 73
```

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—2-14

To display the number of various EIGRP packets sent and received, use the **show ip eigrp traffic** command, as illustrated in the figure. For example, in this network, router R1 has sent 429 hello messages and received 192 hello messages.

Summary

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

Summary

- **The configuration commands for basic EIGRP include:**
 - **router eigrp** *autonomous-system*
 - **network** *network-number* [*wildcard-mask*]
 - **bandwidth** *kilobits*
- **The optional *wildcard-mask* parameter in the network command is an inverse mask used to determine how to interpret the *network-number* parameter. A wildcard bit of 0 is a match and of 1 is “don’t care.”**
- **Create and advertise a default route in an EIGRP AS with the `ip default-network network-number` command.**

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—2-15

Summary (Cont.)

- **Use the `show ip eigrp neighbors` command to verify that the router recognizes its neighbors. Use the `show ip route eigrp` command to verify that the router recognizes routes from its neighbors.**
- **Use the `show ip protocols`, `show ip eigrp interfaces`, `show ip eigrp neighbors`, `show ip eigrp topology`, and `show ip eigrp traffic` commands to verify EIGRP operations.**

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—2-16

Configuring Advanced EIGRP Options

Overview

For a scalable Enhanced Interior Gateway Routing Protocol (EIGRP) network, configuring manual route summarization at key points on the internetwork is vital to good network performance. Load balancing across multiple links is a viable option for efficient bandwidth utilization. Limiting the amount of bandwidth that EIGRP uses across these WAN links allows user traffic better access to the WAN links.

This lesson provides advanced configuration options for EIGRP, including route summarization, load balancing, and limiting EIGRP bandwidth utilization on WAN links.


Objectives

Upon completing this lesson, you will be able to configure advanced EIGRP features for scalable networks. This ability includes being able to meet these objectives:

- Explain why administrators may need to use manual route summarization over default automatic route summarization
- Configure route summarization
- Describe the features of load balancing across equal paths
- Configure EIGRP load balancing across unequal-cost paths
- Explain why EIGRP defaults may need to be changed to ensure efficient use of bandwidth across WAN links
- Configure EIGRP bandwidth use across WAN links

Route Summarization

This topic explains why administrators may need to use manual route summarization over default automatic route summarization.



Some EIGRP features, such as automatically summarizing routes at major network boundaries, are characteristics of distance vector operation. Traditional distance vector protocols, which are classful routing protocols, cannot assume the mask for networks that are not directly connected because routing updates do not exchange masks.

Summarizing routes at major classful boundaries creates smaller routing tables. Smaller routing tables make the routing update process less bandwidth-intensive.

Automatic summarization is enabled by default for EIGRP.

EIGRP Route Summarization: Manual

Manual summarization has the following characteristics:

- Summarization is configurable on a per-interface basis in any router within a network.
- When summarization is configured on an interface, the router immediately creates a route pointing to null0.
 - Loop-prevention mechanism
- When the last specific route of the summary goes away, the summary is deleted.
- The minimum metric of the specific routes is used as the metric of the summary route.

A drawback to using distance vector protocols is the inability to create summary routes at arbitrary boundaries within a major network, which would be desirable because summarizing routes creates smaller routing tables. EIGRP allows administrators to disable automatic summarization and to create one or more summary routes within the network on any bit boundary, so long as a more specific route exists in the routing table. When the last specific route of the summary goes away, the summary is deleted from the routing table.

The minimum metric of the specific routes is used as the metric of the summary route.

Recall that Cisco IOS software automatically puts summary routes to interface null0 in the routing table for automatically summarized routes to prevent routing loops. For the same reason, Cisco IOS software also creates a summary route to interface null0 when manual summarization is configured. 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 manual summarization to be effective, blocks of contiguous addresses (subnets) must come together at a common router so that the router can advertise a single summary route. The number of subnets that can be represented by a summary route is directly related to the difference in the number of bits between the subnet mask and the summary mask. The formula 2^n , where n equals the difference in the number of bits between the summary and subnet mask, indicates how many subnets can be represented by a single summary route. For example, if the summary mask contains three fewer bits than the subnet mask, eight ($2^3 = 8$) subnets can be aggregated into one advertisement.

For example, if network 10.0.0.0 is divided into /24 subnets and is summarized to the summarization block 10.1.8.0/21, the difference between the /24 networks and the /21 summarizations is 3 bits; therefore, $2^3 = 8$ subnets can be aggregated. The summarized subnets range from 10.1.8.0/24 through 10.1.15.0/24.

When configuring summary routes, the administrator needs to specify the IP address of the summary route and the summary mask. The Cisco IOS software for EIGRP handles many of the details that surround proper implementation, including details about metrics, loop prevention, and removal of the summary route from the routing table if none of the more specific routes are valid.

Configuring Manual Route Summarization

This topic explains how to use the **no auto-summary** command to disable automatic route summarization and the **ip summary-address eigrp** command to create a summary address on an interface.

Configuring Route Summarization

```
(config-router)#  
no auto-summary
```

- Turns off automatic summarization for the EIGRP process

```
(config-if)#  
ip summary-address eigrp as-number address mask  
[admin-distance]
```

- Creates a summary address that this interface will generate

© 2006 Cisco Systems, Inc. All rights reserved. BSCI v3.0—2-4

EIGRP automatically summarizes routes at the classful boundary. In some cases, you may not want automatic summarization to occur. For example, if you have discontinuous networks, you need to disable automatic summarization to minimize router confusion. To disable automatic summarization, use the **no auto-summary** command under the EIGRP router configuration mode, as shown in the figure.


Note that an EIGRP router does not perform automatic summarization of networks in which it does not participate.

Use the **ip summary-address eigrp** interface command, as also shown in the figure, to manually create a summary route at an arbitrary bit boundary as long as a more specific route exists in the routing table. The table describes the **ip summary-address eigrp** parameters.

ip summary-address eigrp Command

Command	Description
<i>as-number</i>	EIGRP autonomous system (AS) number.
<i>address</i>	The IP address advertised as the summary address. This address does not need to be aligned on Class A, B, or C boundaries.
<i>mask</i>	The IP mask used to create the summary address.
<i>admin-distance</i>	(Optional) Administrative distance. A value from 0 to 255.

Manually Summarizing EIGRP Routes



Example: Summarizing EIGRP Routes

The figure shows a discontinuous network 172.16.0.0. In the configuration examples on routers A and B, automatic summarization has been disabled, so the 172.16.1.0 and 172.16.2.0 subnets are advertised into network 10.0.0.0. The routing tables of routers in the 10.0.0.0 network, including router C as shown in the next figure, now include these discontinuous subnets.

An EIGRP router autosummarizes routes for only networks to which it is attached. If a network was not autosummarized at the major network boundary, as is the case in this example on routers A and B because autosummarization is turned off, all the subnet routes are carried into the router C routing table. Router C will not autosummarize the 172.16.1.0 and 172.16.2.0 subnets because it does not own the 172.16.0.0 network. Therefore, router C would send routes to the 172.16.1.0 subnet and the 172.16.2.0 subnet to the WAN. Configuring a summary route on the router C interface serial 0/0/0, as shown in the figure, means that only one route will be sent on the WAN, representing all subnets that belong to network 172.16.0.0.

To configure manual route summarization, use the following procedure:

- Step 1** Select the interface to propagate the summary route.
- Step 2** Specify the EIGRP routing protocol, the AS number, and the summary address and the mask of the routes being summarized.

Note For manual route summarization, the summary route is advertised only if a component (a more specific entry) of the summary route is present in the routing table.

Router C Routing Table

```
RouterC#show ip route
<output omitted>
Gateway of last resort is not set
  172.16.0.0/16 is variably subnetted, 3 subnets, 2 masks
D    172.16.0.0/16 is a summary, 00:00:04, Null0
D    172.16.1.0/24 [90/156160] via 10.1.1.2, 00:00:04, FastEthernet0/0
D    172.16.2.0/24 [90/20640000] via 10.2.2.2, 00:00:04, Serial0/0/1
C    192.168.4.0/24 is directly connected, Serial0/0/0
  10.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
C    10.2.2.0/24 is directly connected, Serial0/0/1
C    10.1.1.0/24 is directly connected, FastEthernet0/0
D    10.0.0.0/8 is a summary, 00:00:05, Null0
RouterC#
```

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—2-6

Example: Router C Routing Table

This figure shows the routing table on router C. Router C has both 172.16.1.0 and 172.16.2.0, the discontinuous subnets, in its routing table.

Load Balancing Across Equal Paths

This topic describes the features of load balancing across equal paths.

EIGRP Load Balancing

- Routes with a metric equal to the minimum metric are installed in the routing table (equal-cost load balancing).
- There can be up to six entries in the routing table for the same destination:
 - The number of entries is configurable.
 - The default is four.
 - Set to 1 to disable load balancing.

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—2-7

Equal-cost load balancing is the ability of a router to distribute traffic over all its network ports that are the same metric from the destination address. Load balancing increases the use of network segments and increases effective network bandwidth.

For IP, Cisco IOS software applies load balancing between a maximum of four equal-cost paths by default. With the **maximum-paths** *maximum-path* router configuration command, up to six equally good routes can be kept in the routing table. (Setting the *maximum-path* option to 1 disables load balancing.) When a packet is process-switched, load balancing over equal-cost paths occurs on a per-packet basis. When packets are fast-switched, load balancing over equal-cost paths occurs on a per-destination basis. (Thus, if you are testing load balancing, do not ping to or from routers with the fast-switching interfaces, because these locally router-generated packets are process-switched rather than fast-switched and might produce confusing results.)

Configuring Load Balancing Across Unequal-Cost Paths

This topic explains how to use the **variance** command to balance traffic across multiple routes that have different metrics.

EIGRP Unequal-Cost Load Balancing

```
Router(config-router)#  
variance multiplier
```

- Allows the router to include routes with a metric smaller than the *multiplier* value times the minimum metric route to that destination

© 2006 Cisco Systems, Inc. All rights reserved. BSCI v3.0—2-8

EIGRP can also balance traffic across multiple routes that have different metrics, which is called unequal-cost load balancing. The degree to which EIGRP performs load balancing is controlled with the **variance** command, as shown in the figure.

The following table lists the parameter for the **variance** command.


variance Command Parameter

Command	Description
<i>multiplier</i>	A value from 1 to 128, used for load balancing. The default is 1, which indicates equal-cost load balancing. The multiplier defines the range of metric values that are accepted for load balancing by the EIGRP process.

Variance Example

Network	Neighbor	FD	AD
Z	B	30	10
	C	20	10
	D	45	25

(config-router)#
variance 2



- Router E chooses router C to get to network Z, because it has lowest FD of 20.
- With a variance of 2, router E chooses router B to get to network Z ($20 + 10 = 30 < [2 * (FD) = 40]$).
- Router D is never considered to get to network Z (because $25 > 20$).

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0-2.9

Example: Variance

In the figure, a variance of 2 is configured, and the range of the metric values, which are the feasible distances (FDs) for router E to get to network Z, is 20 to 45. This range of values determines the feasibility of a potential route.

A route is feasible if the next router in the path is closer to the destination than the current router and if the metric of the alternate path is within the variance. Load balancing can use only feasible paths, and the routing table includes only these paths. The two feasibility conditions are:

- The local best metric (the current FD) must be greater than the best metric (the advertised distance [AD]) learned from the next router. In other words, the next router in the path must be closer to the destination than the current router; this prevents routing loops.
- The variance multiplied by the local best metric (the current FD) must be greater than the metric through the next router (the alternative FD). This condition is true if the metric of the alternate path is within the variance.

If both of these conditions are met, the route is called feasible and can be added to the routing table.

In the figure, there are three paths to network Z with the following metrics:

- Path 1: 30 (via B)
- Path 2: 20 (via C)
- Path 3: 45 (via D)

By default, the router places only path 2, via C, in the routing table, because it is the least-cost path. To load-balance over paths 1 and 2, use a variance of 2, because $20 * 2 = 40$, which is greater than the metric through path 1.

In this example, router E uses router C as the successor because it has the lowest FD (20). With the **variance 2** command applied to router E, the path through router B meets the criteria for load balancing. In this case, the FD through router B is less than twice the FD for the successor (router C).


Router D is not considered for load balancing with this variance, because the FD through router D is greater than twice the FD for the successor (router C). In this example, however, router D would never be a feasible successor no matter what the variance is. This is because the router D AD of 25 is greater than the router E FD of 20; therefore, to avoid a potential routing loop, router D is not considered a feasible successor.

EIGRP Bandwidth Use Across WAN Links

This topic explains why EIGRP defaults may need to be changed to ensure efficient use of bandwidth across WAN links.

Configuring WAN Links

- EIGRP supports different WAN links:
 - Point-to-point links
 - NBMA
 - Multipoint links
 - Point-to-point links
- EIGRP uses up to 50% of bandwidth by default; this bandwidth utilization can be changed.



© 2006 Cisco Systems, Inc. All rights reserved. BSCI v3.0—2-10

EIGRP operates efficiently in WAN environments. It is scalable on both point-to-point links and on point-to-point and multipoint nonbroadcast multiaccess (NBMA) links.

Because of the inherent differences in the operational characteristics of WAN links, the default configuration parameters may not be the best option for all WAN links. A solid understanding of EIGRP operation, coupled with knowledge of available link speeds, can yield an efficient, reliable, and scalable router configuration.

By default, EIGRP may use up to 50 percent of the bandwidth of an interface or subinterface. EIGRP uses the bandwidth of the link set by the **bandwidth** command, or the default bandwidth of the link if none is configured, when calculating how much bandwidth to use. This percentage can be changed on a per-interface basis by using the **ip bandwidth-percent eigrp as-number percent** interface configuration command. In this command, *as-number* is the AS number, and *percent* is the percentage of the configured bandwidth that EIGRP can use. This percentage can be greater than 100, which may be useful if the bandwidth is configured artificially low for routing-policy reasons. Examples of the use of this command are provided later in this lesson.

Bandwidth Utilization over WAN Interfaces

- **Bandwidth utilization over point-to-point subinterfaces using Frame Relay:**
 - Treats bandwidth as T1 by default
 - Should manually configure bandwidth as the CIR of the PVC
- **Bandwidth utilization over multipoint Frame Relay, ATM, and ISDN PRI:**
 - EIGRP uses the bandwidth on the physical interface divided by the number of neighbors on that interface to calculate the bandwidth attributed per neighbor.

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—2-11

Cisco IOS software assumes that point-to-point Frame Relay subinterfaces (like all serial interfaces) operate at full T1 link speed. In many implementations, however, only fractional T1 speeds are available. Therefore, when configuring these subinterfaces, set the bandwidth to match the contracted committed information rate (CIR) of the permanent virtual circuit (PVC).

When configuring multipoint interfaces, especially for Frame Relay (but also for ATM and ISDN PRI), it is important to understand that all neighbors share the bandwidth equally. That is, EIGRP uses the **bandwidth** command on the physical interface divided by the number of Frame Relay neighbors connected on that physical interface to calculate the bandwidth attributed to each neighbor. EIGRP configuration should reflect the correct percentage of the actual available bandwidth on the line.

Bandwidth Utilization over WAN Interfaces (Cont.)

- Each PVC can have a different CIR, creating an EIGRP packet-pacing problem.
- Multipoint interfaces:
 - Convert these to point-to-point configuration or manually configure bandwidth by multiplying the lowest CIR by the number of PVCs.

© 2006 Cisco Systems, Inc. All rights reserved.


BSCI v3.0—2-12

Each installation has a unique topology and requires a unique configuration. Differing CIR values often require a hybrid configuration that blends the characteristics of point-to-point circuits with multipoint circuits.

When configuring multipoint interfaces, configure the bandwidth to represent the minimum CIR multiplied by the number of circuits. This approach may not fully use the higher-speed circuits, but it ensures that the circuits with the lowest CIR are not overdriven. If the topology has a small number of very low-speed circuits, these interfaces are typically defined as point-to-point so that their bandwidth can be set to match the provisioned CIR.

Configuring EIGRP Bandwidth Use Across WAN Links

This topic explains how to use the **ip bandwidth-percent eigrp as-number percent** command to efficiently configure EIGRP bandwidth use across WAN links.



Example: WAN Configuration—Frame Relay Hub-and-Spoke Topology

As described earlier, the **ip bandwidth-percent eigrp as-number percent** command allows the maximum percentage of the bandwidth of an interface that EIGRP will use to be specified.


The figure illustrates a common hub-and-spoke design configuration topology with ten virtual circuits to the ten remote sites (only four of the ten remote sites are shown in the slide). The configurations for routers C and G, using EIGRP AS 200, are also shown in the figure.

The circuits are provisioned as 56-kbps links, but there is insufficient bandwidth at the interface to support this allocation. For example, if the hub tries to communicate to all remote sites at the same time, the bandwidth that is required exceeds the available link speed of 256 kbps for the hub: 10 times the CIR of 56 kbps equals 560 kbps.

In a point-to-point topology, all virtual circuits are treated equally; the interfaces and subinterface are therefore all configured with a bandwidth equal to one-tenth of the available link speed (25 kbps).

On each interface and subinterface, the EIGRP allocation percentage is raised to 110 percent of the specified bandwidth in an attempt to ensure that EIGRP packets are delivered through the Frame Relay network. This adjustment causes EIGRP packets to receive approximately 28 kbps of the provisioned 56 kbps on each circuit. This extra configuration restores the 50-50 ratio that was tampered with when the bandwidth was set to an artificially low value.

EIGRP WAN Configuration: Hybrid Multipoint



- Configure lowest CIR virtual circuit as point-to-point, specify bandwidth = CIR.
- Configure higher CIR virtual circuits as multipoint, combine CIRs.

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.6-2-14

Example: WAN Configuration—Hybrid Multipoint

This figure presents an example of a hybrid solution. There is only one lower-speed circuit, and the other circuits are all provisioned to the same CIR.

The preferred configuration on router C shows the low-speed circuit configured as point-to-point with the bandwidth set to the CIR value. The remaining circuits are designated as a multipoint subinterface, and their CIRs are added together to set the bandwidth for the subinterface.

In multipoint interfaces, the bandwidth is shared equally among all circuits. In this case, the bandwidth is set to 768 kbps, which is the sum of the three CIRs ($3 * 256 = 768$). Each link will be allocated one-third of this bandwidth, resulting in 256 kbps each.

Summary

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

Summary

- **EIGRP performs automatic network-boundary summarization, but administrators can disable automatic summarization and perform manual route summarization on an interface-by-interface basis. Summarizing routes creates smaller routing tables.**
- **Use the `no auto-summary` command to disable automatic summarization. Use the `ip summary-address eigrp` command to create a summary address.**
- **EIGRP performs equal-cost load balancing by default for up to four paths (up to six paths can be supported).**
- **Use the `variance` command to configure unequal-cost load balancing.**

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—2-15

Summary (Cont.)

- **EIGRP uses up to 50 percent of the bandwidth of an interface by default. Because of the inherent differences in the operational characteristics of WAN links, this default may not be the best option for all WAN links.**
- **Use the `ip bandwidth-percent eigrp` command to configure EIGRP bandwidth use across WAN links.**

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—2-16

Configuring EIGRP Authentication

Overview

You can prevent your router from receiving fraudulent route updates by configuring neighbor router authentication. Enhanced Interior Gateway Routing Protocol (EIGRP) neighbor authentication (also called neighbor router authentication or route authentication) can be configured such that routers can participate in routing based on predefined passwords.

This lesson describes EIGRP Message Digest 5 (MD5) authentication and how to configure and troubleshoot it.

Objectives

Upon completing this lesson, you will be able to implement authentication in an EIGRP network. This ability includes being able to meet these objectives:

- Describe router authentication
- Describe the MD5 authentication used in EIGRP
- Configure MD5 authentication
- Troubleshoot MD5 authentication

Router Authentication

This topic describes router authentication used by routing protocols.

Router Authentication

- **Many routing protocols support authentication such that a router authenticates the source of each routing update packet that it receives.**
- **Simple password authentication is supported by:**
 - IS-IS
 - OSPF
 - RIPv2
- **MD5 authentication is supported by:**
 - OSPF
 - RIPv2
 - BGP
 - EIGRP

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—2.2

Neighbor router authentication (also called route authentication) can be configured such that routers only participate in routing based on predefined passwords.

By default, no authentication is used for routing protocol packets. When neighbor router authentication has been configured on a router, the router authenticates the source of each routing update packet that it receives, which is accomplished by the exchange of an authentication key (also called a password) that is known to both the sending and the receiving router.

There are two types of authentication: simple password authentication (also called plain-text authentication) and MD5 authentication.

Simple password authentication is supported by Integrated System-Integrated System (IS-IS), Open Shortest Path First (OSPF), and Routing Information Protocol version 2 (RIPv2). MD5 authentication is supported by OSPF, RIPv2, Border Gateway Protocol (BGP), and EIGRP.

Note Authentication for EIGRP, OSPF, and BGP is covered in this course.

Simple Password vs. MD5 Authentication

- **Simple password authentication:**
 - Router sends packet and key.
 - Neighbor checks whether key matches its key.
 - Process not secure.
- **MD5 authentication:**
 - Configure a key (password) and key ID; router generates a message digest, or hash, of the key, key ID and message.
 - Message digest is sent with packet; key is not sent.
 - Process OS secure.

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—2-3

Both forms of authentication work in the same way, with the exception that MD5 sends a message digest instead of the authenticating key itself. The message digest is created using the key and a message, but the key itself is not sent, preventing it from being read while it is being transmitted. Simple password authentication sends the authenticating key itself over the wire.

Note	Note that simple password authentication is not recommended for use as part of your security strategy, because it is vulnerable to passive attacks. Anybody with a link analyzer could easily view the password on the wire. The primary use of simple password authentication is to avoid accidental changes to the routing infrastructure. Using MD5 authentication, however, is a recommended security practice.

Note	As with all keys, passwords, and other security secrets, it is imperative that you closely guard authenticating keys used in neighbor authentication. The security benefits of this feature rely on your keeping all authenticating keys confidential. Also, when performing router management tasks via Simple Network Management Protocol (SNMP), do not ignore the risk associated with sending keys using unencrypted SNMP.

With simple password authentication, a password (key) is configured on a router; each participating neighbor router must be configured with the same key.

MD5 authentication is a cryptographic authentication. A key (password) and key ID are configured on each router. The router uses an algorithm based on the routing protocol packet, the key, and the key ID to generate a message digest (also called a hash) that is appended to the packet. Unlike simple authentication, the key is not exchanged over the wire—the message digest is sent instead of the key, which ensures that nobody can eavesdrop on the line and learn keys during transmission.

MD5 Authentication

This topic describes MD5 authentication used in EIGRP.

EIGRP MD5 Authentication

- **EIGRP supports MD5 authentication.**
- **Router generates and checks every EIGRP packet. Router authenticates the source of each routing update packet that it receives.**
- **Configure a key (password) and key ID; each participating neighbor must have same key configured.**

© 2006 Cisco Systems, Inc. All rights reserved. BSCI v3.0—2-4

EIGRP neighbor authentication (also called neighbor router authentication or route authentication) can be configured such that routers can participate in routing based on predefined passwords.

By default, no authentication is used for EIGRP packets. EIGRP can be configured to use MD5 authentication.

When neighbor authentication has been configured on a router, the router authenticates the source of each routing update packet that it receives. For EIGRP MD5 authentication, an authenticating key (sometimes referred to as a password) and a key ID must be configured on both the sending and the receiving router.

MD5 Authentication

EIGRP MD5 authentication:

- Router generates a message digest, or hash, of the key, key ID, and message.
- EIGRP allows keys to be managed using key chains.
- Specify key ID (number), key, and lifetime of key.
- First valid activated key, in order of key numbers, is used.

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—2-6

The MD5 keyed digest in each EIGRP packet prevents the introduction of unauthorized or false routing messages from unapproved sources.

Each key has its own key ID, which is stored locally. The combination of the key ID and the interface associated with the message uniquely identifies the authentication algorithm and MD5 authentication key in use.

EIGRP allows keys to be managed using key chains. Each key definition within the key chain can specify a time interval for which that key will be activated (its lifetime). Then, during the lifetime of a given key, routing update packets are sent with this activated key. Only one authentication packet is sent, regardless of how many valid keys exist. The software examines the key numbers in order from lowest to highest, and it uses the first valid key that it encounters.

Keys cannot be used during time periods for which they are not activated. Therefore, it is recommended that for a given key chain, key activation times overlap to avoid any period of time for which no key is activated. If there is a time period during which no key is activated, neighbor authentication cannot occur, and therefore routing updates will fail.

Note that the router needs to know the time to be able to rotate through keys in synchronization with the other participating routers, so that all routers are using the same key at the same moment. Refer to the Network Time Protocol (NTP) and calendar commands in the “Performing Basic System Management” chapter of the *Cisco IOS Configuration Fundamentals Configuration Guide* for information about configuring time at your router.

Configuring MD5 Authentication

This topic describes how to configure MD5 authentication for EIGRP.

Configuring EIGRP MD5 Authentication

```
Router(config-if)#  
ip authentication mode eigrp autonomous-system md5
```

- **Specifies MD5 authentication for EIGRP packets**

```
Router(config-if)#  
ip authentication key-chain eigrp autonomous-system  
name-of-chain
```

- **Enables authentication of EIGRP packets using key in the keychain**

© 2006 Cisco Systems, Inc. All rights reserved. BSCI v3.0—2-6

To configure MD5 authentication for EIGRP, complete the following steps:

- Step 1** Enter configuration mode for the interface on which you want to enable authentication.
- Step 2** Specify MD5 authentication for EIGRP packets using the **ip authentication mode eigrp md5** command, as shown in the figure. The table describes the parameter in this command.

ip authentication mode eigrp md5 Parameter

Parameter	Description
<i>autonomous-system</i>	EIGRP autonomous system (AS) number in which authentication is to be used

- Step 3** Enable the authentication of EIGRP packets with a key specified in a key chain by using the **ip authentication key-chain eigrp** command, as shown in the figure. The parameters for this command are described in the table.

ip authentication key-chain eigrp Parameters

Parameter	Description
<i>autonomous-system</i>	EIGRP AS number in which authentication is to be used
<i>name-of-chain</i>	Name of the authentication key chain from which a key is to be obtained

Configuring EIGRP MD5 Authentication (Cont.)

```
Router(config)#  
key chain name-of-chain
```

- Enters configuration mode for the keychain

```
Router(config-keychain)#  
key key-id
```

- Identifies key and enters configuration mode for the keyid

© 2006 Cisco Systems, Inc. All rights reserved. BSCI v2.0--2.7

Step 4 Enter the configuration mode for the key chain using the **key chain** command, as shown in the figure. The table describes the parameter in this command.

key chain Parameter

Parameter	Description
<i>name-of-chain</i>	Name of the authentication key chain from which a key is to be obtained

Step 5 Identify a key ID to use, and enter configuration mode for that key using the **key** command, as shown in the figure. The table describes the parameter in this command.

key Parameter

Parameter	Description
key-id	ID number of an authentication key on a key chain. The range of keys is from 0 to 2147483647. The key ID numbers need not be consecutive.

Configuring EIGRP MD5 Authentication (Cont.)

```
Router(config-keychain-key) #  
key-string text
```

- Identifies key string (password)

```
Router(config-keychain-key) #  
accept-lifetime start-time {infinite | end-time | duration  
seconds}
```

- Optional: Specifies when key will be accepted for received packets

```
Router(config-keychain-key) #  
send-lifetime start-time {infinite | end-time | duration  
seconds}
```

- Optional: Specifies when key can be used for sending packets

© 2006 Cisco Systems, Inc. All rights reserved. BSCI v3.0-2.8

Step 6 Identify the key string (password) for this key using the **key-string** command, as shown in the figure. The table describes the parameter in this command.

key-string Parameter

Parameter	Description
text	Authentication string that is to be used to authenticate sent and received EIGRP packets. The string can contain from 1 to 80 uppercase and lowercase alphanumeric characters, except that the first character cannot be a number.

Step 7 Optionally, specify the time period during which this key will be accepted for use on received packets using the **accept-lifetime** command, as shown in the figure. The table describes the parameters in this command.

accept-lifetime Parameters

Parameter	Description
<i>start-time</i>	Beginning time that the key specified by the key command is valid for use on received packets. The syntax can be either of the following: <i>hh:mm:ss month date year</i> <i>hh:mm:ss date month year</i> <i>hh</i> —hours <i>mm</i> —minutes <i>ss</i> —seconds <i>month</i> —first three letters of the name of the month <i>date</i> —date (1–31) <i>year</i> —year (four digits) The default start time and the earliest acceptable date is January 1, 1993.
infinite	Key is valid for use on received packets from the <i>start-time</i> value on.
<i>end-time</i>	Key is valid for use on received packets from the <i>start-time</i> value until the <i>end-time</i> value. The syntax is the same as that for the <i>start-time</i> value. The <i>end-time</i> value must be after the <i>start-time</i> value. The default end time is an infinite time period.
<i>seconds</i>	Length of time (in seconds) that the key is valid for use on received packets. The range is from 1 to 2147483646.


Step 8 Optionally, specify the time period during which this key can be used for sending packets using the **send-lifetime** command, as shown in the figure. The table describes the parameters in this command.

send-lifetime Parameters

Parameter	Description
<i>start-time</i>	Beginning time that the key specified by the key command is valid to be used for sending packets. The syntax can be either of the following: <i>hh:mm:ss month date year</i> <i>hh:mm:ss date month year</i> <i>hh</i> —hours <i>mm</i> —minutes <i>ss</i> —seconds <i>month</i> —first three letters of the name of the month <i>date</i> —date (1 to 31) <i>year</i> —year (four digits) The default start time and the earliest acceptable date is January 1, 1993.
infinite	Key is valid to be used for sending packets from the <i>start-time</i> value on.
<i>end-time</i>	Key is valid to be used for sending packets from the <i>start-time</i> value until the <i>end-time</i> value. The syntax is the same as that for the <i>start-time</i> value. The <i>end-time</i> value must be after the <i>start-time</i> value. The default end time is an infinite time period.
<i>seconds</i>	Length of time (in seconds) that the key is valid to be used for sending packets. The range is from 1 to 2147483646.

Note	If the service password-encryption command is not used when implementing EIGRP authentication, the key string will be stored as plaintext in the router configuration. If you configure the service password-encryption command, the key string will be stored and displayed in an encrypted form; when it is displayed, there will be an encryption type of 7 specified before the encrypted key string.

Example MD5 Authentication Configuration



© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—2-9

Example: MD5 Authentication Configuration

The figure shows the network used to illustrate the configuration, verification, and troubleshooting of MD5 authentication.

R1 Configuration for MD5 Authentication

```
<output omitted>
key chain R1chain
  key 1
    key-string firstkey
    accept-lifetime 04:00:00 Jan 1 2006 infinite
    send-lifetime 04:00:00 Jan 1 2006 04:01:00 Jan 1 2006
  key 2
    key-string secondkey
    accept-lifetime 04:00:00 Jan 1 2006 infinite
    send-lifetime 04:00:00 Jan 1 2006 infinite
<output omitted>
interface FastEthernet0/0
  ip address 172.16.1.1 255.255.255.0
!
interface Serial0/0/1
  bandwidth 64
  ip address 192.168.1.101 255.255.255.224
  ip authentication mode eigrp 100 md5
  ip authentication key-chain eigrp 100 R1chain
!
router eigrp 100
  network 172.16.1.0 0.0.0.255
  network 192.168.1.0
  auto-summary
```

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—2-10

Example: R1 Configuration for MD5 Authentication

The configuration of the R1 router is shown in this figure.

MD5 authentication is configured on the serial 0/0/1 interface with the **ip authentication mode eigrp 100 md5** command. The **ip authentication key-chain eigrp 100 R1chain** command specifies that the key chain R1chain is to be used.

The **key chain R1chain** command enters configuration mode for the R1chain key chain. Two keys are defined. Key 1 is set to “first key” with the **key-string firstkey** command. This key is acceptable for use on packets received by R1 from January 1, 2006 onward, as specified in the **accept-lifetime 04:00:00 Jan 1 2006 infinite** command. However, the **send-lifetime 04:00:00 Jan 1 2006 04:01:00 Jan 1 2006** command specifies that this key is valid for use only when sending packets for one minute on January 1, 2006; it is no longer valid for use in sending packets.

Key 2 is set to “second key” with the **key-string secondkey** command. This key is acceptable for use on packets received by R1 from January 1, 2006 onward, as specified in the **accept-lifetime 04:00:00 Jan 1 2006 infinite** command. This key can also be used when sending packets from January 1, 2006 onward, as specified in the **send-lifetime 04:00:00 Jan 1 2006 infinite** command.

R1 will therefore accept and attempt to verify the MD5 digest of any EIGRP packets with a key ID equal to 1. R1 will also accept a packet with a key ID equal to 2. All other MD5 packets will be dropped. R1 will send all EIGRP packets using key 2, because key 1 is no longer valid for use when sending packets.

R2 Configuration for MD5 Authentication

```
<output omitted>
key chain R2chain
  key 1
    key-string firstkey
    accept-lifetime 04:00:00 Jan 1 2006 infinite
    send-lifetime 04:00:00 Jan 1 2006 infinite
  key 2
    key-string secondkey
    accept-lifetime 04:00:00 Jan 1 2006 infinite
    send-lifetime 04:00:00 Jan 1 2006 infinite
<output omitted>
interface FastEthernet0/0
  ip address 172.17.2.2 255.255.255.0
!
interface Serial0/0/1
  bandwidth 64
  ip address 192.168.1.102 255.255.255.224
  ip authentication mode eigrp 100 md5
  ip authentication key-chain eigrp 100 R2chain
!
router eigrp 100
  network 172.17.2.0 0.0.0.255
  network 192.168.1.0
  auto-summary
```

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—2-11

Example: R2 Configuration for MD5 Authentication

The configuration of the R2 router is shown in this figure.

MD5 authentication is configured on the serial 0/0/1 interface with the **ip authentication mode eigrp 100 md5** command. The **ip authentication key-chain eigrp 100 R2chain** command specifies that the key chain R2chain is to be used.

The **key chain R2chain** command enters configuration mode for the R2chain key chain. Two keys are defined. Key 1 is set to “first key” with the **key-string firstkey** command. This key is acceptable for use on packets received by R2 from January 1, 2006 onward, as specified in the **accept-lifetime 04:00:00 Jan 1 2006 infinite** command. This key can also be used when sending packets from January 1, 2006 onward, as specified in the **send-lifetime 04:00:00 Jan 1 2006 infinite** command.

Key 2 is set to “second key” with the **key-string secondkey** command. This key is acceptable for use on packets received by R2 from January 1, 2006 onward, as specified in the **accept-lifetime 04:00:00 Jan 1 2006 infinite** command. This key can also be used when sending packets from January 1, 2006 onward, as specified in the **send-lifetime 04:00:00 Jan 1 2006 infinite** command.

R2 will therefore accept and attempt to verify the MD5 digest of any EIGRP packets with a key ID equal to 1 or 2. R2 will send all EIGRP packets using key 1, because it is the first valid key in the key chain.

Verifying MD5 Authentication

```
R1#
*Jan 21 16:23:30.517: %DUAL-5-NBRCHANGE: IP-EIGRP(0) 100: Neighbor
192.168.1.102 (Serial0/0/1) is up: new adjacency

R1#show ip eigrp neighbors
IP-EIGRP neighbors for process 100
H   Address                Interface      Hold Uptime    SRTT   RTO  Q  Seq
                               (sec)          (ms)          Cnt  Num
0   192.168.1.102           Se0/0/1       12 00:03:10    17   2280 0   14

R1#show ip route
<output omitted>
Gateway of last resort is not set
D   172.17.0.0/16 [90/40514560] via 192.168.1.102, 00:02:22, Serial0/0/1
    172.16.0.0/16 is variably subnetted, 2 subnets, 2 masks
D   172.16.0.0/16 is a summary, 00:31:31, Null0
C   172.16.1.0/24 is directly connected, FastEthernet0/0
    192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks
C   192.168.1.96/27 is directly connected, Serial0/0/1
D   192.168.1.0/24 is a summary, 00:31:31, Null0
R1#ping 172.17.2.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.17.2.2, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 12/15/16 ms
```

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—2-12

Verifying MD5 Authentication


The figure shows the output of the **show ip eigrp neighbors** and **show ip route** commands on the R1 router.

The neighbor table indicates that the two routers have successfully formed an EIGRP adjacency. The routing table verifies that the 172.17.0.0 network has been learned via EIGRP over the serial connection.

The results of a ping to the R2 Fast Ethernet interface address are also displayed to illustrate that the link is working.

Troubleshooting MD5 Authentication

This topic describes how to troubleshoot MD5 authentication for EIGRP.



Example: Successful MD5 Authentication

The output of the **debug eigrp packets** command on R1 shown in the figure illustrates that R1 is receiving EIGRP packets with MD5 authentication, with a key ID equal to 1, from R2.

Similarly, the output of the **debug eigrp packets** command on R2 shown in the figure illustrates that R2 is receiving EIGRP packets with MD5 authentication, with a key ID equal to 2, from R1.

Troubleshooting MD5 Authentication Problem

MD5 authentication on both R1 and R2, but R1 key 2 (that it uses when sending) changed

```
R1(config-if)#key chain Rlchain
R1(config-keychain)#key 2
R1(config-keychain-key)#key-string wrongkey

R2#debug eigrp packets
EIGRP Packets debugging is on
  (UPDATE, REQUEST, QUERY, REPLY, HELLO, IPXSAP, PROBE, ACK, STUB, SIAQUERY,
  SIAREPLY)
R2#
*Jan 21 16:50:18.749: EIGRP: pkt key id = 2, authentication mismatch
*Jan 21 16:50:18.749: EIGRP: Serial0/0/1: ignored packet from 192.168.1.101, op
ode = 5 (invalid authentication)
*Jan 21 16:50:18.749: EIGRP: Dropping peer, invalid authentication
*Jan 21 16:50:18.749: EIGRP: Sending HELLO on Serial0/0/1
*Jan 21 16:50:18.749:   AS 100, Flags 0x0, Seq 0/0 idbQ 0/0 iidbQ un/rely 0/0
*Jan 21 16:50:18.753: %DUAL-5-NBRCHANGE: IP-EIGRP(0) 100: Neighbor 192.168.1.101
  (Serial0/0/1) is down: Auth failure

R2#show ip eigrp neighbors
IP-EIGRP neighbors for process 100
R2#
```

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—2-14

Example: Troubleshooting MD5 Authentication Problems

In this example, the key string for key 2 of router R1, the one that it uses when sending EIGRP packets, is changed to be different from the key string that router R2 is expecting.

The output of the **debug eigrp packets** command on R2 shown in the figure illustrates that R2 is receiving EIGRP packets with MD5 authentication, with a key ID equal to 2, from R1, but that there is an authentication mismatch. The EIGRP packets from R1 are ignored, and the neighbor relationship is declared to be down. The output of the **show ip eigrp neighbors** command confirms that R2 does not have any EIGRP neighbors.

The two routers keep trying to re-establish their neighbor relationship. Because of the different keys used by each router in this scenario, R1 will authenticate hello messages sent by R2 using key 1. However, when R1 sends a hello message back to R2 using key 2, there will be an authentication mismatch. From the perspective of R1, the relationship appears to be up for awhile, but then it times out, as illustrated by the following messages received on R1. The output of the **show ip eigrp neighbors** command on R1 also illustrates that R1 does have R2 in its neighbor table for a short time.

```
R1#
*Jan 21 16:54:09.821: %DUAL-5-NBRCHANGE: IP-EIGRP(0) 100:
Neighbor 192.168.1.102 (Serial0/0/1) is down: retry limit
exceeded
*Jan 21 16:54:11.745: %DUAL-5-NBRCHANGE: IP-EIGRP(0) 100:
Neighbor 192.168.1.102 (Serial0/0/1) is up: new adjacency
R1#show ip eigrp neighbors
```

H	Address	Interface	Hold Uptime	SRTT	RTO	Q	Seq
			(sec)	(ms)		Cnt	Num
0	192.168.1.102	Se0/0/1	13 00:00:38	1	5000	1	0

Summary

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

Summary

- **There are two types of router authentication: simple password and MD5.**
- **When EIGRP authentication is configured, the router generates and checks every EIGRP packet and authenticates the source of each routing update packet that it receives. EIGRP supports MD5 authentication.**
- **To configure MD5 authentication, use the ip authentication mode eigrp and ip authentication key-chain interface commands. The key chain must also be configured, starting with the key chain command.**
- **Use debug eigrp packets to verify and troubleshoot MD5 authentication.**

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—2-15

Lesson 5

Using EIGRP in an Enterprise Network

Overview

Network administrators benefit from understanding how to configure Enhanced Interior Gateway Routing Protocol (EIGRP) to prevent common routing problems that hinder network scalability. For example, you can implement EIGRP stub routers to limit the EIGRP query range, making EIGRP more scalable with fewer complications.

EIGRP is a scalable routing protocol that ensures that as a network grows larger, it operates efficiently and adjusts rapidly to changes. This lesson describes practical EIGRP-specific design and configuration techniques to implement an effective scalable network.

Objectives

Upon completing this lesson, you will be able to describe, recognize, and correct common EIGRP issues and problems. This ability includes being able to meet these objectives:

- Explain factors affecting scalability in large internetworks
- Explain how EIGRP uses queries to update its routing tables in the event that a route is lost and there is no feasible successor
- Explain how to mark the spokes of large network as stubs to reduce EIGRP queries and thus improve network scaling
- Explain why SIA connections occur
- Explain how to minimize active routes
- Describe how graceful shutdown prevents loss of packets when routers go down

Scalability in Large Networks

This topic explains factors affecting scalability in large internetworks.

Factors That Influence EIGRP Scalability

- **Quantity of routing information exchanged between peers;** without proper route summarization, this can be excessive.
- **Number of routers that must be involved when a topology change occurs.**
- **Depth of topology:** the number of hops that information must travel to reach all routers.
- **Number of alternate paths through the network.**

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—2.2

Some of the factors that affect network scalability are as follows:

- **Amount of information exchanged between neighbors:** If more information than necessary for routing to function correctly is exchanged between EIGRP neighbors, the routers have to work harder at neighbor startup and to react to changes in the network.
- **Number of routers:** When a topology change occurs in the network, EIGRP resource consumption directly relates to the number of routers that must be involved in the change.
- **Depth of the topology:** The topology depth can affect the convergence time. Depth refers to the number of hops that information must travel to reach all routers. A multinational network without route summarization is an example of a network with large depth and therefore increasing convergence time. A three-tiered network design (as described in Module 1) is highly recommended for all IP routing environments. There should never be more than seven hops between any two routing devices on an internetwork. The propagation delay and query process across multiple hops when changes occur may slow down convergence of the network.
- **Number of alternate paths through the network:** A network should provide alternate paths to avoid single points of failure. However, too much complexity (too many alternate paths) can also create EIGRP convergence problems because the EIGRP routing process, using queries, needs to explore all possible paths for lost routes. This complexity creates an ideal condition for a router to become stuck in active (SIA) as it awaits a response to queries that are being propagated through these many alternate paths.

EIGRP Queries

This topic explains how EIGRP uses queries to converge rapidly when a route is lost.

EIGRP Query Process

- **Queries are sent when a route is lost and no feasible successor is available.**
- **The lost route is now in active state.**
- **Queries are sent to all neighboring routers on all interfaces except the interface to the successor.**
- **If the neighbors do not have the lost-route information, queries are sent to their neighbors.**
- **If a router has an alternate route, it answers the query; this stops the query from spreading in that branch of the network.**

© 2006 Cisco Systems, Inc. All rights reserved.


BSCI v3.9—2-3

As an advanced distance vector protocol, EIGRP relies on neighboring routers to provide routing information. Recall that when a router loses a route and does not have a feasible successor in its topology table, it looks for an alternate path to the destination. This is known as *going active* on a route; a route is considered passive when a router is not performing recompilation on that route.

The router sends query packets to all neighbors on interfaces other than the one used to reach the previous successor (split horizon), inquiring whether they have a route to the given destination. If a router has an alternate route, it answers the query and does not propagate it further. If a neighbor does not have an alternate route, it queries each of its own neighbors for an alternate path. The queries then propagate through the network, creating an expanding tree of queries. When a router answers a query, it stops the spread of the query through that branch of the network; however, the query can still spread through other branches of the network as other routers attempt to find alternate paths, which might not exist.

EIGRP Stubs

The stability of large-scale EIGRP networks is often dependent on the range of queries through the network. This topic explains how to mark the spokes of a large network as stubs to reduce EIGRP queries and thus improve network scaling.



Hub-and-spoke network topologies commonly use stub routing. In a hub-and-spoke topology, the remote router forwards all traffic that is not local to a hub router; the remote router does not need to retain a complete routing table. Generally, the hub router needs to send only a default route to the remote routers.

In a hub-and-spoke topology, having a full routing table on the remote router serves no functional purpose because the path to the corporate network and the Internet is always through the hub router. Additionally, having a full routing table at the spoke router increases the amount of memory required. Route summarization and route filtering can be used to conserve bandwidth and memory requirements on the spoke routers.

Traffic from a hub router should not use a remote router as a transit path. A typical connection from a hub router to a remote router has significantly less bandwidth than a connection at the network core; attempting to use the connection to a remote router as a transit path typically results in excessive congestion, as illustrated in the figure. The EIGRP stub routing feature can prevent this problem by restricting the remote router from advertising the hub router routes back to other hub routers. For example, routes recognized by the remote router E from hub router A are not advertised to hub router B. Because the remote router does not advertise the hub routes back to the hub routers, the hub routers do not use the remote routers as a transit path. Using the EIGRP stub routing feature improves network stability, reduces resource utilization, and simplifies stub router configuration.

EIGRP Stub

- The EIGRP stub routing feature improves network stability, reduces resource utilization, and simplifies remote router (spoke) configuration.
- Stub routing is commonly used in a hub-and-spoke topology.
- A stub router sends a special peer information packet to all neighboring routers to report its status as a stub router.
- A neighbor that receives a packet informing it of the stub status does not query the stub router for any routes.

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—2-8

The EIGRP stub feature was first introduced in Cisco IOS Software Release 12.0(7)T.

Only the remote routers are configured as stubs. A stub router sends a special peer information packet to all neighboring routers to report its status as a stub router. Any neighbor that receives a packet informing it of the stub status does not query the stub router for any routes. Therefore, a router that has a stub peer does not query that peer; instead, hub routers connected to the stub router answer the query on behalf of the stub router. The stub routing feature does not prevent routes from being advertised to the remote router.

The EIGRP stub routing feature also simplifies the configuration and maintenance of hub-and-spoke networks. When stub routing is enabled in dual-homed remote configurations, you do not have to configure filtering on remote routers to prevent them from appearing as transit paths to the hub routers.

Caution EIGRP stub routing should be used on stub routers only. A stub router is defined as a router connected to the network core or hub layer through which core transit traffic should not flow. A stub router should only have hub routers for EIGRP neighbors; ignoring this restriction may cause undesirable behavior.

Configuring EIGRP Stub

Router(config-router)#

eigrp stub [receive-only | connected | static | summary]

- **receive-only:** Prevents the stub from sending any type of route.
- **connected:** Permits stub to send connected routes (may still need to redistribute).
- **static:** Permits stub to send static routes (must still redistribute).
- **summary:** Permits stub to send summary routes.
- **Default is** connected and summary.

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0–2-8

To configure a router as an EIGRP stub, use the **eigrp stub** command, as illustrated in the figure. A router configured as a stub with this command shares information about connected and summary routes with all neighboring routers by default.

The table describes the parameters of the **eigrp stub** command.

eigrp stub Command Parameters

Command	Description
receive-only	The receive-only keyword restricts the router from sharing any of its routes with any other router within an EIGRP autonomous system (AS). This keyword does not permit any other option to be specified, because it prevents any type of route from being sent. The three other optional keywords (connected , static , and summary) cannot be used with the receive-only keyword. Use this option if there is a single interface on the router.
connected	The connected keyword permits the EIGRP stub routing feature to send connected routes. If a network command does not include the connected routes, it might be necessary to redistribute connected routes with the redistribute connected command under the EIGRP process. This option is enabled by default and is the most widely practical stub option.
static	The static keyword permits the EIGRP stub routing feature to send static routes. Redistributing static routes with the redistribute static command is still necessary.
summary	The summary keyword permits the EIGRP stub routing feature to send summary routes. Summary routes can be created manually with the ip summary-address command or automatically at a major network border router with the auto-summary command enabled. This option is enabled by default.

The parameters of this command can be used in any combination, with the exception of the **receive-only** keyword. If one of these keywords, except **receive-only**, is used individually, then the connected and summary routes are not sent automatically.

The EIGRP stub routing feature does not automatically enable route summarization on the hub router. In most cases, the network administrator should configure route summarization on the hub routers.


If a true stub network is required, the hub router can be configured to send a default route to the spoke routers. This approach is the most simple and conserves the most bandwidth and memory on the spoke routers.

Note	Although EIGRP is a classless routing protocol, it has classful behavior by default, such as having automatic summarization on by default. When you configure the hub router to send a default route to the remote router, ensure that the ip classless command is issued on the remote router. By default, the ip classless command is enabled in all Cisco IOS images that support the EIGRP stub routing feature.

The EIGRP stub routing feature allows a network administrator to prevent sending queries to the spoke router under any condition. It is highly recommended that you use both EIGRP route summarization and EIGRP stub features to provide the best scalability.

Without the stub feature, a hub router will send a query to the spoke routers if a route is lost somewhere in the network. If there is a communication problem over the WAN link between the hub router and the spoke router, replies may not be received for all queries (this is known as being SIA), and the network may become unstable.

Limiting Updates and Queries: Using EIGRP Stub



Example: Limiting Updates and Queries: Using EIGRP Stub

The figure illustrates how using the EIGRP stub feature affects the network shown earlier. Each of the remote routers is configured as a stub. Queries for network 10.1.8.0/24 are not sent to routers C, D, or E, thus reducing the bandwidth used and the chance of the routes being SIA.

Using the EIGRP stub feature at the remote sites allows the hub (regional offices) sites to immediately answer queries without propagating the queries to the remote sites, saving CPU cycles and bandwidth and lessening convergence time even when the remote sites are dual-homed to two or more hub (regional) sites.


Example: eigrp stub Parameters

If stub connected is configured:

- B will advertise 10.1.2.0/24 to A.
- B will not advertise 10.1.2.0/23, 10.1.3.0/23, or 10.1.4.0/24.

If stub summary is configured:

- B will advertise 10.1.2.0/23 to A.
- B will not advertise 10.1.2.0/24, 10.1.3.0/24, or 10.1.4.0/24.



Example: eigrp stub Parameters

This figure illustrates an example network and configuration for router B. The networks that will be advertised when the various options of the **eigrp stub** command are used are also shown.


In the first scenario, with the **eigrp stub connected** command, router B will advertise only 10.1.2.0/24. Notice that although 10.1.3.0/24 is also a connected network, it is not advertised to router A because it is not advertised in a **network** command, and connected routes are not redistributed.

In the second scenario, with the **eigrp stub summary** command, router B will advertise only 10.1.2.0/23, the summary route that is configured on the router.

Example: eigrp stub Parameters (Cont.)

If stub static is configured:

- B will advertise 10.1.4.0/24 to A.
- B will not advertise 10.1.2.0/24, 10.1.2.0/23, or 10.1.3.0/24.



If stub receive-only is configured:

- B will not advertise anything to A, so A needs to have a static route to the networks behind B to reach them.

```
Router B
ip route 10.1.4.0 255.255.255.0 10.1.1.10
!
interface serial 0
 ip summary-address eigrp 10.1.2.0 255.255.254.0
!
router eigrp 100
 redistribute static 1000 1 255 1 1500
 network 10.2.2.2 0.0.0.1
 network 10.1.2.0 0.0.0.255
!
eigrp stub static
eigrp stub receive-only
```

Fig. 19

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0-2.9

In the third scenario, with the **eigrp stub static** command, router B will advertise only 10.1.4.0/24, the static route that is configured on the router.

In the final scenario, with the **eigrp stub receive-only** command, router B will not advertise anything.

SIA Connections

SIA routes can be some of the most challenging problems to resolve in an EIGRP network. This topic explains why SIA connections occur.

EIGRP Query Process Stuck in Active

- **The router has to get all the replies from the neighbors with an outstanding query before the router calculates the successor information.**
- **If any neighbor fails to reply to the query within 3 minutes by default, the route is SIA, and the router resets the neighbor relationship with the neighbor that fails to reply.**

© 2006 Cisco Systems, Inc. All rights reserved. BSCI v3.0--2-10

EIGRP uses a reliable multicast approach to search for an alternate to a lost route; therefore, it is imperative that EIGRP receive a reply for each query it generates in the network.

Once a route goes active and the query sequence is initiated, the route can only come out of the active state and move to passive state when it receives a reply for every generated query. If the router does not receive a reply to all the outstanding queries within 3 minutes (the default time), the route goes to the SIA state.


When the route goes to SIA state, the querying router resets the neighbor relationship to the neighbor that fails to reply. This setting causes the router to go active on all routes known through the lost neighbor and to readvertise all the routes that it knows about to the lost neighbor.

The most common reasons for SIA routes are as follows:

- The router is too busy to answer the query because of high CPU usage or memory problems and cannot allocate the memory to process the query or build the reply packet.
- The link between the two routers is not good; therefore, some packets are lost between the routers. While the router receives enough packets to maintain the neighbor relationship, the router does not receive all queries or replies.
- A failure causes traffic on a link to flow in only one direction—this is called a unidirectional link.

Preventing SIA Connections

A route becomes active when it goes down or its metric worsens and there are no feasible successors. This topic explains how to minimize active routes.



SIA-Query and SIA-Reply are two new additions to the Type, Length, Value (TLV) triplets in the EIGRP packet header. These packets are generated automatically with no configuration required, from Cisco IOS Software Release 12.1(5) and later with the active process enhancement feature. This feature enables an EIGRP router to monitor the progression of the search for a successor route and ensure that the neighbor is still reachable. Improved network reliability results from reducing the unintended termination of the neighbor adjacency.


The figure on the left illustrates what would happen before this feature was introduced. Router A sends a query for network 10.1.1.0/24 to router B. Router B has no entry for this network, so it queries router C. If problems exist between router B and C, the reply packet from router C to router B may be delayed or lost. Router A has no visibility of downstream progress and assumes that the lack of response indicates problems with router B. After the router A 3-minute active timer expires, the neighbor relationship with router B is reset, along with all known routes from router B.

By contrast, with the active process enhancement feature, router A queries downstream router B (with an SIA-Query) at the midway point of the active timer (1.5 minutes by default) about the status of the route. Router B responds (with an SIA-Reply) that it is searching for a replacement route. Upon receiving this SIA-Reply response packet, router A validates the status of router B and does not terminate the neighbor relationship.

Meanwhile router B sends up to three SIA-Queries to router C. If they go unanswered, router B terminates the neighbor relationship with router C. Router B then updates router A with an SIA-Reply indicating that the network 10.1.1.0/24 is unreachable. Routers A and B remove the active route from their topology tables. The neighbor relationship between routers A and B remains intact.

Graceful Shutdown

This topic describes how graceful shutdown prevents loss of packets when routers go down.



Graceful shutdown, implemented with the goodbye message feature, is designed to improve EIGRP network convergence.

In the figure, router A is using router B as the successor for a number of routes; router C is the feasible successor for the same routes. Router B normally would not tell router A if the EIGRP process on router B was going down; for example, if router B was being reconfigured. Router A would have to wait for its hold timer to expire before it would discover the change and react to it. Packets sent during this time would be lost.

With graceful shutdown, the goodbye message is broadcast when an EIGRP routing process is shut down to inform adjacent peers about the impending topology change. This feature allows supporting EIGRP peers to synchronize and recalculate neighbor relationships more efficiently than would occur if the peers discovered the topology change after the hold timer expired.

The goodbye message is supported in Cisco IOS Software Release 12.3(2), 12.3(3)B, and 12.3(2)T and later. Goodbye messages are sent in hello packets. EIGRP sends an interface goodbye message with all K values set to 255 when taking down all peers on an interface. The following message is displayed by routers that support goodbye messages when one is received:

```
*Apr 26 13:48:42.523: %DUAL-5-NBRCHANGE: IP-EIGRP(0) 1:
Neighbor 10.1.1.1 (Ethernet0/0) is down: Interface Goodbye
received
```


A Cisco router that runs a software release that does not support the goodbye message will misinterpret the message as a K-value mismatch and therefore display the following message:

```
*Apr 26 13:48:41.811: %DUAL-5-NBRCHANGE: IP-EIGRP(0) 1:
Neighbor 10.1.1.1 (Ethernet0/0) is down: K-value mismatch
```

Note	The receipt of a goodbye message by a peer that does not support this feature does not disrupt normal network operation. The peer will terminate the session when the hold timer expires. The sending and receiving routers will reconverge normally after the sender reloads.

Summary

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

Summary

- **Factors that affect network scalability include these:**
 - Amount of information exchanged between neighbors
 - Number of routers
 - Depth of the topology
 - Number of alternate paths through the network
- **When a route is lost and no feasible successor is available, queries are sent to all neighboring routers on all interfaces.**
- **The `eigrp stub` command is used to enable the stub routing feature, which improves network stability, reduces resource utilization, and simplifies stub router configuration.**

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—2-13

Summary (Cont.)

- **After a route goes active and the query sequence is initiated, it can only come out of the active state and move to passive state when it receives a reply for every generated query. If the router does not receive a reply to all the outstanding queries within 3 minutes (the default time), the route goes to the SIA state.**
- **The active process enhancement feature enables an EIGRP router to monitor the progression of the search for a successor route so that neighbor relationships are not reset unnecessarily.**
- **With graceful shutdown, a goodbye message is broadcast when an EIGRP routing process is shut down, to inform adjacent peers about the impending topology change.**

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—2-14

Module Summary

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

Module Summary

- **EIGRP starts by building a table of adjacent neighbors. Route exchanges with these neighbors result in an EIGRP topology table. The best EIGRP routes are moved into the IP routing table.**
- **The configuration commands for basic EIGRP include these:**
 - **router eigrp** *autonomous-system*
 - **network** *network-number* [*wildcard-mask*]
- **Use the no auto-summary command to disable automatic summarization. Use the ip summary-address eigrp command to create a summary address. Use the variance command to configure unequal-cost load balancing.**

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.9—2-1

Module Summary (Cont.)

- **To configure EIGRP MD5 authentication, use the ip authentication mode eigrp and ip authentication key-chain interface commands. The key chain must also be configured, starting with the key chain command.**
- **Features such as stub routing, active process enhancement, and graceful shutdown help improve network stability and performance.**

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.9—2-2

Configuring Enhanced Interior Gateway Routing Protocol (EIGRP) for your routing environment enables you to achieve such benefits as rapid convergence, lower bandwidth utilization, and multiple routed protocol support. Using EIGRP ensures that as a network grows larger, it will still operate efficiently and adjust to changes rapidly.

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 three features are benefits of EIGRP? (Choose three.) (Source: Introducing EIGRP)
- A) fast convergence
 - B) support for VLSM and discontinuous subnets
 - C) same metric algorithm as OSPF
 - D) manual route summarization at any point in the network
- Q2) What is listed in the EIGRP topology table? (Source: Introducing EIGRP)
- A) directly connected routers that have formed an EIGRP adjacency
 - B) best routes to a destination network
 - C) all routes learned from each EIGRP neighbor
- Q3) What is listed in the EIGRP neighbor table? (Source: Introducing EIGRP)
- A) directly connected routers that have formed an EIGRP adjacency
 - B) best routes to a destination network
 - C) all routes learned from each EIGRP neighbor
- Q4) What is listed in the IP routing table? (Source: Introducing EIGRP)
- A) directly connected routers that have formed an EIGRP adjacency
 - B) best routes to a destination network
 - C) all routes learned from each EIGRP neighbor
- Q5) Which two statements are true of the EIGRP metric calculation? (Choose two.) (Source: Introducing EIGRP)
- A) The following are the default K values: K1 = 1, K2 = 1, K3 = 0, K4 = 0, K5 = 0.
 - B) To convert an IGRP metric to an EIGRP metric, multiply the IGRP metric by 256.
 - C) To convert an EIGRP metric to an IGRP metric, multiply the EIGRP metric by 256.
 - D) The following are the default K values: K1 = 1, K2 = 0, K3 = 1, K4 = 0, K5 = 0.
- Q6) Which three characteristics are key features of EIGRP? (Choose three.) (Source: Introducing EIGRP)
- A) fast convergence
 - B) partial updates
 - C) support for multiple Layer 3 protocols
 - D) backward compatibility with RIP

- Q7) Which of the following are key EIGRP technologies? (Choose three.) (Source: Introducing EIGRP)
- A) RTP
 - B) protocol-dependent modules
 - C) protocol-independent modules
 - D) DUAL
 - E) RMTP
- Q8) Which two characteristics are features of EIGRP? (Choose two.) (Source: Introducing EIGRP)
- A) support for load balancing across unequal-cost paths
 - B) manual summarization at any point on the internetwork
 - C) provision of highly structured area design requirements
- Q9) Which type of database is a list of all EIGRP adjacencies? (Source: Introducing EIGRP)
- A) EIGRP topology table
 - B) EIGRP neighbor table
 - C) IP routing table
- Q10) Which type of database contains a list of all the best EIGRP routes to reach a destination? (Source: Introducing EIGRP)
- A) EIGRP topology table
 - B) EIGRP neighbor table
 - C) IP routing table
- Q11) Which type of database contains a list of all possible EIGRP routes to reach a destination? (Source: Introducing EIGRP)
- A) EIGRP topology table
 - B) EIGRP neighbor table
 - C) IP routing table
- Q12) Which five criteria can be considered by EIGRP to calculate the metric? (Choose five.) (Source: Introducing EIGRP)
- A) MTU
 - B) bandwidth
 - C) cost
 - D) delay
 - E) loading
 - F) hop count
 - G) reliability
- Q13) Which two criteria are used by EIGRP to calculate the metric by default? (Choose two.) (Source: Introducing EIGRP)
- A) MTU
 - B) bandwidth
 - C) cost
 - D) delay
 - E) loading
 - F) hop count
 - G) reliability

- Q14) Which packet type establishes neighbor relationships? (Source: Introducing EIGRP)
- A) ACK
 - B) hello
 - C) query
 - D) reply
 - E) update
- Q15) Which packet type is responsible for sending routing advertisements? (Source: Introducing EIGRP)
- A) ACK
 - B) hello
 - C) query
 - D) reply
 - E) update
- Q16) Which packet type is used to acknowledge a reliable packet? (Source: Introducing EIGRP)
- A) ACK
 - B) hello
 - C) query
 - D) reply
 - E) update
- Q17) Which packet type is used to respond to a query? (Source: Introducing EIGRP)
- A) ACK
 - B) hello
 - C) query
 - D) reply
 - E) update
- Q18) Which packet type is used to ask neighbors about routing information? (Source: Introducing EIGRP)
- A) ACK
 - B) hello
 - C) query
 - D) reply
 - E) update
- Q19) Which three packets must be explicitly acknowledged? (Choose three.) (Source: Introducing EIGRP)
- A) ACK
 - B) hello
 - C) reply
 - D) query
 - E) update
- Q20) How does DUAL select the successor for a specific destination network? (Source: Introducing EIGRP)
- A) by selecting the next-hop router with the highest FD
 - B) by selecting the next-hop router with the lowest FD
 - C) by selecting the next-hop router with the highest AD
 - D) by selecting the next-hop router with the lowest AD

- Q21) What is the formula for selecting a feasible successor? (Source: Introducing EIGRP)
- A) The AD of the current successor route is less than the FD of the feasible successor route.
 - B) The FD of the current successor route is less than the AD of the feasible successor route.
 - C) The FD of the feasible successor route is less than the AD of the current successor route.
 - D) The AD of the feasible successor route is less than the FD of the current successor route.
- Q22) What does EIGRP do when a successor fails and there are no feasible successors, but there are alternate paths available? (Source: Introducing EIGRP)
- A) It immediately uses the alternate pathway with the lowest FD and sends queries and updates to ensure that this pathway is loop-free.
 - B) It automatically uses the alternate pathway with the lowest FD.
 - C) It sends queries to see if the alternate paths are still viable. When a loop-free path is found, the path is installed in the routing table.
 - D) It removes the network from the routing table and waits for the periodic update from EIGRP neighbors to see if an alternate route exists.
- Q23) Which two conditions signify the active state for EIGRP? (Choose two.) (Source: Introducing EIGRP)
- A) The route can be used and is stable.
 - B) The route cannot be used.
 - C) EIGRP queries are outstanding and the router is waiting for EIGRP replies.
 - D) The state signifies that this is the best route with the lowest FD.

- Q24) Test your understanding of EIGRP by matching terms with statements. Write the letter of the statement in front of the term that the statement describes. (Source: Introducing EIGRP)

Term

- _____ 1. successor
- _____ 2. feasible successor
- _____ 3. hello
- _____ 4. topology table
- _____ 5. IP
- _____ 6. update
- _____ 7. routing table
- _____ 8. DUAL

Statement

- A) a network protocol that EIGRP supports
- B) a database that contains successor and feasible successor information
- C) a database that includes administrative distance
- D) a neighbor router that has the best path to a destination
- E) a neighbor router that has the best alternate loop-free path to a destination
- F) an algorithm used by EIGRP to ensure fast convergence
- G) a multicast packet used to discover neighbors
- H) a packet sent by EIGRP routers when a new neighbor is discovered and a change occurs

- Q25) What is the purpose of the **network** command for EIGRP? (Source: Implementing and Verifying EIGRP)

- I) to determine which router interfaces participate in EIGRP and which networks the router advertises
- J) to specify the AS number to which the router belongs
- K) to define the EIGRP neighbors
- L) to tell EIGRP which networks to advertise, those directly connected and those learned through EIGRP

- Q26) Which command creates a default route for EIGRP? (Source: Implementing and Verifying EIGRP)

- A) **ip default-network network-number**
- B) **ip route 0.0.0.0 0.0.0.0 outbound-interface**
- C) **ip route 0.0.0.0 255.0.0.0 outbound-interface**
- D) **ip route 0.0.0.0 255.255.255.255 outbound-interface**

- Q27) Which command displays an indication if a network is SIA? (Source: Implementing and Verifying EIGRP)

- A) **show ip route**
- B) **show ip protocol**
- C) **show ip eigrp topology**

- Q28) Which command displays the K value settings for EIGRP? (Source: Implementing and Verifying EIGRP)
- A) **show ip route**
 - B) **show ip protocols**
 - C) **show ip eigrp topology**
- Q29) Which statement about the EIGRP AS system number is true? (Source: Implementing and Verifying EIGRP)
- A) is an optional parameter that can be left blank
 - B) must be different on each router that wants to exchange EIGRP updates
 - C) must be the same on each router that wants to exchange EIGRP updates
 - D) is an IP address
- Q30) Which is the correct **network** command to allow updates to propagate only out of interfaces that are part of subnet 10.1.0.0/16? (Source: Implementing and Verifying EIGRP)
- A) **network 10.1.0.0 mask 255.255.0.0**
 - B) **network 10.1.0.0 mask 0.0.255.255**
 - C) **network 10.1.0.0 255.255.0.0**
 - D) **network 10.1.0.0 0.0.255.255**
- Q31) Which three are true of configuring the **ip default-network** command for EIGRP? (Choose three.) (Source: Implementing and Verifying EIGRP)
- A) The network must be reachable by the router using this command.
 - B) The command will set the gateway of last resort to 0.0.0.0 on the router issuing this command.
 - C) The network must be advertised to other neighbors as an EIGRP route.
 - D) The network will be flagged by other EIGRP routers as a candidate default route.
- Q32) What does the passive state in the EIGRP topology table signify? (Source: Implementing and Verifying EIGRP)
- A) There are outstanding queries for this network.
 - B) The network is unreachable.
 - C) The network is up and operational, and this state signifies normal conditions.
 - D) A feasible successor has been selected.
- Q33) Which command indicates the number of EIGRP peer routers on an interface? (Source: Implementing and Verifying EIGRP)
- A) **show ip eigrp interfaces**
 - B) **show ip eigrp neighbors**
 - C) **show ip route**
 - D) **show ip eigrp topology**
- Q34) By default, how many equal-cost paths to the same destination network can EIGRP place in the routing table? (Source: Configuring Advanced EIGRP Options)
- A) one
 - B) two
 - C) four
 - D) six

- Q35) Between headquarters and remote site A, there are two dedicated serial PPP connections, one at 64 kbps and the other at 128 kbps. What is the appropriate variance to allow for unequal-cost load balancing across these links? (Source: Configuring Advanced EIGRP Options)
- A) 1
 - B) 2
 - C) 3
 - D) 4
- Q36) What is the default bandwidth percentage that EIGRP will use on WAN links? (Source: Configuring Advanced EIGRP Options)
- A) 25 percent
 - B) 50 percent
 - C) 75 percent
 - D) 100 percent
- Q37) Which command is used to disable automatic EIGRP network-boundary summarization, and where is it applied? (Source: Configuring Advanced EIGRP Options)
- A) **no boundary-summarization** at the interface level
 - B) **no auto-summary** under the routing process
 - C) **no auto-summary** at the interface level
 - D) **no boundary-summarization** under the routing process
- Q38) Which command is used to configure manual summarization of all the subnets in network 10.1.32.0/21 for EIGRP in AS 101? (Source: Configuring Advanced EIGRP Options)
- A) **ip summary-address eigrp 101 10.1.32.0 255.255.248.0**
 - B) **ip eigrp 101 summary-address 10.1.32.0 255.255.240.0**
 - C) **ip summary-address eigrp 101 10.1.32.0 255.255.240.0**
 - D) **ip eigrp 101 summary-address 10.1.32.0 255.255.248.0**
- Q39) A router is running EIGRP in AS 200 on a subinterface with a CIR of 64 kbps. The subinterface is configured with the **bandwidth 50** command. Which command will result in the router using a maximum of 40 kbps for EIGRP on that subinterface? (Source: Configuring Advanced EIGRP Options)
- A) **ip bandwidth-percent eigrp 200 62**
 - B) **ip bandwidth-percent eigrp 62 200**
 - C) **ip bandwidth-percent eigrp 80 200**
 - D) **ip bandwidth-percent eigrp 200 80**
- Q40) Which authentication does EIGRP support? (Source: Configuring EIGRP Authentication)
- A) MD5
 - B) MD5 and simple password
 - C) simple password
 - D) none

- Q41) When EIGRP authentication is configured between two routers, each router has its own unique password. (Source: Configuring EIGRP Authentication)
- A) true
 - B) false
- Q42) Which three of the following are used to generate the message digest when EIGRP MD5 authentication is configured? (Choose three.) (Source: Configuring EIGRP Authentication)
- A) packet
 - B) sequence number
 - C) key ID
 - D) key
 - E) router ID
- Q43) What does the **accept-lifetime 04:00:00 Jan 1 2006 infinite** command do? (Source: Configuring EIGRP Authentication)
- A) specifies that a key is acceptable for use on received packets from the first of January 2006 onward
 - B) specifies that a key is acceptable for use on sent packets from the first of January 2006 onward
 - C) specifies that a key is acceptable for use on received packets until the first of January 2006
 - D) specifies that a key is acceptable for use on sent packets until the first of January 2006 onward
- Q44) Which command is used to specify that EIGRP MD5 authentication in AS 100 is to be utilized? (Source: Configuring EIGRP Authentication)
- A) **ip authentication mode eigrp 100 md5**
 - B) **ip eigrp 100 authentication mode md5**
 - C) **ip authentication-key eigrp 100**
 - D) **ip message-digest-key eigrp 100**
 - E) **ip eigrp 100 authentication message-digest**
- Q45) Which command is used to troubleshoot EIGRP authentication? (Source: Configuring EIGRP Authentication)
- A) **debug ip eigrp adj**
 - B) **debug ip eigrp packets**
 - C) **debug eigrp packets**
 - D) **debug ip eigrp adjacency events**
 - E) **debug eigrp adj**
- Q46) When a router gets a query from a neighboring router that is not a successor for the network listed in the query, and that network is in a passive state on this router, what does the router do? (Source: Using EIGRP in an Enterprise Network)
- A) The router replies that the destination is unreachable.
 - B) The router attempts to find a new successor; if successful, it replies with new information. If the router is not successful, it marks the destination unreachable and queries all neighboring routers except the previous successor.
 - C) The router replies with the current successor information.
 - D) The router marks the destination unreachable and queries all neighboring routers except the previous successor.

- Q47) Which are three factors that impact network scalability? (Choose three.) (Source: Using EIGRP in an Enterprise Network)
- A) number of alternate paths through the network
 - B) amount of information exchanged between neighbors
 - C) AS number
 - D) depth of the topology
- Q48) Which three statements are true for implementing EIGRP stub routers? (Choose three.) (Source: Using EIGRP in an Enterprise Network)
- A) Stub routing is commonly used on networks with a hub-and-spoke topology.
 - B) The EIGRP stub feature should be configured only on remote spoke routers.
 - C) EIGRP stub routers can and should be used at a transit point to other parts of the network and other autonomous systems.
 - D) Queries are not propagated to EIGRP stub routers. EIGRP updates are sent to stub routers, or a default route is passed.
- Q49) How long does a querying router wait to reset a neighbor that fails to reply to a query? (Source: Using EIGRP in an Enterprise Network)
- A) 15 seconds
 - B) 40 seconds
 - C) 1 minute
 - D) 3 minutes
- Q50) Which command configures an EIGRP stub router to not send any routing updates? (Source: Using EIGRP in an Enterprise Network)
- A) **eigrp stub**
 - B) **eigrp stub receive-only**
 - C) **eigrp stub no-send**
 - D) **eigrp stub none**
- Q51) With graceful shutdown, how is a goodbye message sent? (Source: Using EIGRP in an Enterprise Network)
- A) in an update packet
 - B) in a query packet
 - C) in a goodbye packet
 - D) in a hello packet

Module Self-Check Answer Key

- Q1) A, B, D
- Q2) C
- Q3) A
- Q4) B
- Q5) B, D
- Q6) A, B, C
- Q7) A, B, D
- Q8) A, B
- Q9) B
- Q10) C
- Q11) A
- Q12) A, B, D, E, G
- Q13) B, D
- Q14) B
- Q15) E
- Q16) A
- Q17) D
- Q18) C
- Q19) C, D, E
- Q20) B
- Q21) D
- Q22) C
- Q23) B, C
- Q24) 1 = D, 2 = E, 3 = G, 4 = B, 5 = A, 6 = H, 7 = C, 8 = F
- Q25) A
- Q26) A
- Q27) C
- Q28) B
- Q29) C
- Q30) D
- Q31) A, C, D
- Q32) C
- Q33) A
- Q34) C

- Q35) B
- Q36) B
- Q37) B
- Q38) A
- Q39) D
- Q40) A
- Q41) B
- Q42) A, C, D
- Q43) A
- Q44) A
- Q45) C
- Q46) C
- Q47) A, B, D
- Q48) A, B, D
- Q49) D
- Q50) B
- Q51) D

Configuring OSPF

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 based primarily on RFC 2328. OSPF is a fairly complex protocol made up of several protocol handshakes, database advertisements, and packet types.

Configuration and verification of OSPF in a Cisco Systems 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. 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:

- Describe how OSPF operates
- Explain how information flows between routers to maintain OSPF links
- Explain how to configure OSPF single-area and multiarea routing
- Describe the features of various OSPF network architectures
- Describe how LSAs and OSPF databases maintain links through the network
- Describe the procedure for configuring OSPF route summarization for interarea and external routes
- Implement and verify OSPF area parameters including stub, NSSA, totally stubby, and backbone
- Implement authentication in an OSPF network

Introducing the OSPF Protocol

Overview

Open Shortest Path First Protocol (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, shortest path first (SPF) calculations, and how OSPF verifies that its links are still active.


Objectives

Upon completing this lesson, you will be able to describe how OSPF operates. This ability includes being able to meet these objectives:

- Describe link-state routing protocols
- Describe the two-tier hierarchy structure of OSPF
- Describe how routers running a link-state routing protocol establish neighbor adjacencies with their neighboring routers
- Describe how OSPF calculates the best path to each destination network
- Describe how routers use LSUs to verify that links are still active

Link-State Routing Protocols

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



The need to overcome limitations of distance vector routing protocols led to the development of link-state routing protocols. Link-state routing protocols have the following characteristics:

- Respond quickly to network changes
- Send triggered updates when a network change occurs
- Send periodic updates, known as link-state refresh, at long intervals, such as every 30 minutes

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 a link-state advertisement (LSA) concerning that link.

The LSA propagates to all neighboring devices using a special multicast address. Each routing device takes 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.

Link-State Data Structures

- **Neighbor table:**
 - Also known as the adjacency database
 - Contains list of recognized neighbors
- **Topology table:**
 - Typically referred to as LSDB
 - Contains all routers and their attached links in the area or network
 - Identical LSDB for all routers within an area
- **Routing table:**
 - Commonly named a forwarding database
 - Contains list of best paths to destinations

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—3-3

OSPF and Intermediate System-to-Intermediate System (IS-IS) are classified as link-state routing protocols because of the manner in which they distribute routing information and calculate routes.

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 paths 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 best paths to each destination:** Each router independently calculates 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 held in the routing table.

Link-State Routing Protocols

- **Link-state routers recognize more information about the network than their distance vector counterparts.**
- **Each router has a full picture of the topology.**
- **Consequently, link-state routers tend to make more accurate decisions.**

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0–3.4

The memory resources that are needed to maintain these tables represent one drawback to link-state protocols. Because the topology table is identical for all OSPF routers in an area and contains full information about all the routers and links in an area, each router is able to independently select a loop-free and efficient pathway, based on cost, to reach every network in the area. This benefit overcomes the “routing by rumors” limitations of distance vector routing.

With distance vector routing protocols, the routers rely on routing decisions from the neighbors. Routers do not have a full picture of the network topology.

With link-state routing protocols, each router has a full picture of the network topology, and each router can independently make a decision based on an accurate picture of the network topology.

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 used.

Link-State Data Structure: Network Hierarchy

- **Link-state routing requires a hierarchical network structure that is enforced by OSPF.**
- **This two-level hierarchy consists of the following:**
 - **Transit area (backbone or area 0)**
 - **Regular areas (nonbackbone areas)**

© 2006 Cisco Systems, Inc. All rights reserved.


BSCI v3.0—3-5

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

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

- **Transit area:** An OSPF area whose primary function is the fast and efficient movement of IP packets. Transit areas interconnect with other OSPF area types. Generally, end users are not found within a transit area. OSPF area 0, also known as the backbone area, is by definition a transit area.
- **Regular area:** An OSPF area whose primary function is to connect users and resources. Regular areas are usually set up along functional or geographical groupings. By default, a regular 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. An area that does not allow traffic to pass through it is known as a regular area, or nonbackbone area, and can have a number of subtypes, including standard area, stub area, totally stubby area, and not-so-stubby area (NSSA).

OSPF forces a rigid two-layer area hierarchy. The underlying physical connectivity of the network must map to the two-layer area structure, with all nonbackbone areas attaching directly to area 0.



In link-state routing protocols, all routers must keep a copy of the LSDB; the more OSPF routers, the larger the LSDB. It can be advantageous to have all information in all routers, but this approach does not scale to large network sizes.

The area concept is a compromise. Routers inside an area maintain detailed information about the links and only general or summary information about routers and links in other areas.


When a router or link fails, that information is flooded along adjacencies only to the routers in the local area. Routers outside the area do not receive this information. By maintaining a hierarchical structure and limiting the number of routers in an area, an OSPF autonomous system (AS) can scale to very large sizes.

OSPF areas require a hierarchical structure, meaning that all areas must connect directly to area 0, the backbone. In the figure, notice that links between area 1 routers and area 2 or 3 routers are not allowed.

All interarea traffic must pass through the backbone area, area 0. 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

- Routers A and B are backbone routers.
- Backbone routers make up area 0.
- Routers C, D, and E are known as area border routers (ABRs).
- ABRs attach all other areas to area 0.



Routers that make up area 0 are known as backbone routers. 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. 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.

OSPF Adjacency Databases

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



A router 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. In general, routers establish adjacencies as follows:

1. The router sends and receives hello packets to and from its neighboring routers. The format of the destination address is typically multicast.
2. The routers exchange hello packets subject to protocol-specific parameters, such as checking whether the neighbor is in the same AS and area. Routers declare the neighbor up when the exchange is complete.
3. After two routers establish neighbor adjacency using the hello packets, they synchronize their LSDBs by exchanging LSAs and confirming the receipt of LSAs from the adjacent router. The two neighboring routers now recognize that they have synchronized their LSDBs with each other. For OSPF, the routers are now in full adjacency state with each other.
4. If necessary, the routers forward any new LSAs to other neighboring routers, ensuring complete synchronization of link-state information inside the area.

Forming OSPF Adjacencies

- **Point-to-point WAN links:**
 - Both neighbors become fully adjacent.
- **LAN links:**
 - Neighbors form a full adjacency with the DR and BDR.
 - Routers maintain two-way state with the other routers (DROTHERs).
- Routing updates and topology information are passed only between adjacent routers.
- Once an adjacency is formed, LSDBs are synchronized by exchanging LSAs.
- LSAs are flooded reliably throughout the area (or network).

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—3-9

The two OSPF routers on a point-to-point serial link, usually encapsulated in High-Level Data Link Control (HDLC) or PPP, form a full adjacency with each other.

LAN links elect one router as the designated router (DR) and another as the backup designated router (BDR). All other routers on the LAN form full adjacencies with these two routers and pass LSAs only to them. The DR forwards the updates from one neighbor on the LAN to all other neighbors on that LAN.

One of the main functions of a DR is to ensure that all of the routers on the same LAN have identical databases, and the DR passes its database to any new routers that come up. It is inefficient to have all the routers on that LAN pass the same information to the new router, so one router represents the other routers to a new router on the LAN or to other routers in the area.

Routers on the LAN also maintain a partial-neighbor relationship, a two-way adjacency state, with the other routers on the LAN that are not the DR or BDR (DROTHERs).

The exchange of link-state information occurs through LSAs, which are also known as link-state protocol data units (PDUs). LSAs report the state of routers and the links between routers, hence the term *link state*. Link-state information must be synchronized between routers, which means the following:

- LSAs are reliable; there is a method for acknowledging the delivery of LSAs.
- LSAs are flooded throughout the area (or throughout the domain if there is only one area).
- LSAs have a sequence number and a set lifetime so that each router recognizes that it has the most up-to-date version of the LSA.
- LSAs are periodically refreshed to confirm topology information before it ages out of the link-state database.

Only by reliably flooding link-state information can every router in the area or domain ensure that it 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.

ABRs are, by definition, in two or more areas. ABRs maintain an LSDB for each area they are in. Most LSAs are flooded only throughout an area; the ABR may also generate LSAs into another area. This process is described later in this module.

Calculating the OSPF Metric

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

OSPF Calculation

Routers find the best paths to destinations by applying Dijkstra's SPF algorithm to the link-state database as follows:

- **Every router in an area has the identical link-state database.**
- **Each router in the area places itself into the root of the tree that is built.**
- **The best path is calculated with respect to the lowest total cost of links to a specific destination.**
- **Best routes are put into the forwarding database (routing table).**

© 2006 Cisco Systems, Inc. All rights reserved.


BSCI v3.0—3-10

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.

By assigning a cost to each link in the network, and by placing the specific node at the root of a tree and summing the costs toward each given destination, the branches of the tree can be calculated to determine the best path to each destination. The best paths are put in the forwarding database (routing table).

For OSPF, the default behavior is that the interface cost is calculated based on its configured bandwidth. An OSPF cost can also be manually defined for each interface, which overrides the default cost value.

SPF Calculation




Example: SPF Calculation

The figure illustrates an example of a Dijkstra calculation. The Dijkstra calculation occurs as follows:

- Router H advertises its presence to router E; router E then passes on the advertisements of H and its own advertisements to its neighbors (C and G); router G passes these and its own advertisements on to D; and so on.
- These LSAs follow the split horizon rule, which dictates that a router should never advertise an LSA to the router from which it came. In the example, router E does not advertise the LSAs of router H back to H.
- Router x has four neighboring routers: A, B, C, and D. From these routers, it 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 depicted in the figure.
- 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 resulting best paths (the SPF tree). From these best paths, shown with solid lines, routes to destination networks attached to each router are offered to the routing table; for each route, the next-hop address is the appropriate neighboring router (A, B, C, or D).

Link-State Data Structures

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



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 to distance vector routers, which send their 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 a link-state acknowledgment (LSAck) back, 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 an LSack back, 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.

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 sizes.**
- **Adjacencies are built by OSPF routers using the Hello protocol. Over these logical adjacencies, LSUs are sent to exchange database information between adjacent OSPF routers.**

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—3-13

Summary (Cont.)

- **Dijkstra's SPF algorithm is used to calculate best paths for all destinations. SPF is run against the LSDB, and the outcome is a table of best paths, known as the routing table.**
- **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.**

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—3-14

OSPF Packet Types

Overview

The Open Shortest Path First Protocol (OSPF) protocol has five packet types: 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. This ability includes being able to meet these objectives:

- Describe the five OSPF packet types
- Describe how OSPF neighbor adjacencies are established
- Describe the process of exchanging and synchronizing the LSDBs (topology tables) between routers
- Describe how OSPF maintains synchronization of the LSDBs of all routers in the network
- Describe the process of maintaining a database of only the most recent link-state sequence numbers
- Describe how to verify that OSPF packets are flowing properly between two routers

OSPF Packet Types

This topic describes the five OSPF packet types.




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 router to router
4	LSU	Sends specifically requested link-state records
5	LSAck	Acknowledges the other packet types

OSPF Packet Header Format




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 TCP is not used, 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:** For OSPF version 2
- **Type:** Differentiates the five OSPF packet types
- **Packet length:** Length of 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 either no authentication, clear-text passwords or encrypted Message Digest 5 (MD5) formats for router authentication
- **Authentication:** Used in authentication scheme
- **Data (for hello packet):** Includes a list of known neighbors
- **Data (for DBD packet):** Contains a summary of the link-state database (LSDB), which includes all known router IDs and their last sequence number, among a number of other fields
- **Data (for LSR packet):** Contains the type of LSU needed and the router ID that has the needed LSU
- **Data (for LSU packet):** Contains the full link-state advertisement (LSA) entries. Multiple LSA entries can fit in one OSPF update packet
- **Data (for LSack packet):** Is empty

Establishing OSPF Neighbor Adjacencies

This topic describes how OSPF neighbor adjacencies are established.



Neighbor OSPF routers must recognize each other on the network before they can share information because OSPF routing depends on the status of the link between two routers. This process is done using the Hello protocol. The Hello protocol establishes and maintains neighbor relationships by ensuring bidirectional (two-way) communication between neighbors. Bidirectional communication occurs when a router recognizes itself listed in the hello packet received from a neighbor.

Each interface 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 at 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 routing updates by replacing them with a default route.

Note	After a DR and BDR are selected, any router added to the network will establish adjacencies with the DR and BDR only.

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.




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

1. Router A 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 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 running OSPF receive the hello packet from router A and add router A to their list of neighbors. This state is the initial state (init).
3. All routers that received the hello packet send a unicast reply hello packet to router A with their corresponding information. The neighbor field in the hello packet includes all neighboring routers and router A.
4. When router A receives these hello packets, it adds all the routers that had its router ID in their hello packets to its own neighbor relationship database. This state is referred to as the two-way state. At this point, all routers that have each other in their lists of neighbors have established bidirectional communication.
5. If the link type is a broadcast network, generally a LAN link like Ethernet, then 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.

6. Periodically (every 10 seconds by default on broadcast networks) 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 whose hello packets have been received by the router. Remember that “received” means that the receiving router recognizes its name as one of the entries in the received hello packet.

Discovering the Network Routes



© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0-3.6

After the DR and BDR have been selected, the routers are considered to be in the exstart state, and they are ready to discover the link-state information about the internetwork and create their LSDBs. The process used to discover the network routes is the exchange protocol, and it gets the routers 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.

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.

Step 2 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 about 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 3 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 it received with the information it has. If the DBD has a more up-to-date link-state entry, then the router sends an LSR to the other router. The process of sending LSRs is called 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 4 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 and 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.

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.

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

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 (that includes the updated LSA entry) to all OSPF DRs and BDRs (at 224.0.0.6). An LSU packet may contain several distinct LSAs.
- Step 2** The DR acknowledges the 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 on the network.

Step 4 The router updates its LSDB using the LSU that includes the changed LSA. It then recomputes the shortest path first (SPF) algorithm against the updated database after a short delay (the SPF delay) and updates the routing table as necessary.

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 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 the SPF algorithm is calculating. However, if the SPF is calculating a new route, the new route is used only after the SPF calculation is complete.

Maintaining Link-State Sequence Numbers

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

LSA Sequence Numbering

- Each LSA in the LSDB maintains a sequence number.
- The sequence numbering scheme is a 4-byte number that begins with 0x80000001 and ends with 0x7FFFFFFF.
- OSPF floods each LSA every 30 minutes to maintain proper database synchronization. Each time the LSA is flooded, the sequence number is incremented by one.
- Ultimately, an LSA sequence number will wrap around to 0x80000001. When this occurs, the existing LSA is prematurely aged to the maximum age (one hour) and flushed.
- When a router encounters two instances of an LSA, it must determine which is more recent. The LSA having the newer (higher) LS a sequence number is more recent.

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0–3-9

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

The link-state sequence number field in an LSA header is 32 bits in length. Beginning with the leftmost bit set, the first legal sequence number is 0x80000001. It is used to detect old or redundant LSAs; the larger the number, the more recent the LSA.

To ensure an accurate database, OSPF floods (refreshes) each LSA every 30 minutes. Each time 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 longer in the database 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, being refreshed every 30 minutes. At some point, the sequence number will need to wrap back to the starting sequence number. When this process occurs, the existing LSA will be prematurely aged out (the maxage timer is immediately set to one hour) and flushed. The LSA will then begin its sequencing at 0x80000001 again.

LSA Sequence Numbers and Maximum Age

```
RTC# 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>
```

- **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.**

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0--3-10

Example: LSA Sequence Numbers and Maximum Age

The output of the **show ip ospf database** command 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 announces a router LSA for those interfaces that it owns in that area. 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.

Verifying Packet Flow

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

debug ip ospf packet

Debug of 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
```

- Shows fields in OSPF header

© 2006 Cisco Systems, Inc. All rights reserved. BSCI v3.0—3-11

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

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.

OSPF Packet Header Fields

The table lists the OSPF packet header fields 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

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, and exchange) before finally reaching the goal of full state. Full state means that databases are synchronized with adjacent routers.**

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—3-12

Summary (Cont.)

- **LSAs are sent on change 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. The LSA sequence number is incremented every time it is advertised.**
- **Each LSA in the LSDB has a sequence number, which is incremented by one each time the LSA is flooded. When a router encounters two instances of an LSA, it must determine which is more recent. The LSA having the newer (higher) LSA sequence number is more recent.**
- **Use the debug ip ospf packet command to verify that OSPF packets are flowing properly between two routers.**

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—3-13

Configuring OSPF Routing

Overview

This lesson discusses the primary configuration commands for a single-area and multiarea Open Shortest Path First Protocol (OSPF) 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.

Objectives

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

- Describe the procedure to configure basic single-area and multiarea OSPF
- Configure a router ID
- Verify the OSPF router ID
- Verify an OSPF configuration

Configuring Basic Single-Area and Multiarea OSPF

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

Configuring Basic OSPF

```
Router(config)#  
router ospf process-id [vrf vpn-name]
```

- Enables one or more OSPF routing processes

```
Router(config-router)#  
network ip-address wildcard-mask area area-id
```

- Defines the interfaces that OSPF will run on

```
Router(config-if)#  
ip ospf process-id area area-id [secondaries none]
```

- Optional method to enable OSPF explicitly on an interface

© 2006 Cisco Systems, Inc. All rights reserved. BSCI v3.0—3-2

To configure the OSPF process, complete the following steps:

- Step 1** Enable the OSPF process on the router using the **router ospf** command as shown in the figure.

router ospf Parameters

The table describes the parameters of the **router ospf** command.

Parameter	Description
<i>process-id</i>	An internally used number to identify the OSPF routing process. The process ID does not need to match process IDs on other routers. Running multiple OSPF processes on the same router is not recommended because it creates multiple database instances that add extra overhead.
vrf <i>vpn-name</i>	(Optional) Specifies the name of the VPN routing and forwarding (VRF) instance to associate with OSPF VRF processes.

- Step 2** Identify which interfaces on the router are part of the OSPF process, using the **network** command, as shown in the figure. This command also identifies the OSPF area to which the network belongs.

network Parameters

The table describes the parameters of the **network** command.

Parameter	Description
<i>ip-address</i>	Is either the network address, the subnet, or the address of the interface. This address instructs the router to recognize which links to advertise to, which links to check for advertisements, and which networks to advertise.
<i>wildcard-mask</i>	Determines how to interpret the IP address. The mask has wildcard bits, in which 0 is a match and 1 is “don’t care.” For example, 0.0.255.255 indicates a match in the first two octets. If specifying the interface address, use the mask 0.0.0.0 to match all four octets of the address. An address and wildcard mask combination of 0.0.0.0 255.255.255.255 matches all interfaces on the router.
<i>area-id</i>	Specifies the OSPF area to be associated with the address. This parameter can be a decimal number or can be in dotted-decimal notation similar to an IP address, such as A.B.C.D.


Starting with Cisco IOS Software 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.

ip ospf area Parameters

The table describes the parameters of the **ip ospf area** command.

Parameter	Description
<i>process-id</i>	The process ID is entered as a decimal in the range from 1 to 65535.
<i>area-id</i>	The area ID is entered either as a decimal value in the range from 0 to 4294967295 or as an IP address.
secondaries none	(Optional) Prevents secondary IP addresses on the interface from being advertised.

Configuring OSPF on Internal Routers of a Single Area



Example: Configuring OSPF on Internal Routers of a Single Area

The figure shows the OSPF configuration for Fast Ethernet broadcast networks and serial point-to-point links. All three routers in the figure are assigned to area 0 and configured for network 10.0.0.0.


Router A uses a general **network 10.0.0.0 0.255.255.255** statement. This technique assigns all interfaces defined in the 10.0.0.0 network to OSPF process 1.

Router B uses a specific host address technique. The wildcard mask of 0.0.0.0 requires a match on all four octets of the address. This technique allows the operator to define which specific interfaces will run OSPF.

Although the two examples shown are a commonly used combination of a network statement and a wildcard mask, others could also work. For instance, a range of subnets could be specified.

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.

Configuring OSPF for Multiple Areas



Example: Configuring OSPF for Multiple Areas

The figure shows an example of multiarea OSPF configuration. Router A is in area 0, router C is in area 1, and router B is the area border router (ABR) between the two areas.

The configuration for router A is the same as in the previous example.

Router B has a network statement for area 0. The configuration for area 1 in this example uses the **ip ospf 50 area 1** command; alternatively a separate **network** router configuration command could have been used.

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

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—3-5

The OSPF database uses the OSPF router ID to uniquely describe each router in the network. 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 up IP interface on the router for OSPF to use as router ID. If no up interface with an IP address is available when the OSPF process starts, the following error message occurs:

```
p5r2(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 instead of a physical interface address, because a loopback interface never goes down. If there is more than one loopback interface, then 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.

Loopback Interfaces

Unadvertised loopback address


Ex. 192.168.255.254

- Not in OSPF table
- Saves address space
- Cannot use ping command

Advertised loopback address

Ex. 172.16.17.5

- In OSPF table
- Uses address space
- Can use ping command



```
Router(config)#interface loopback 0
Router(config-if)#ip address 172.16.17.5 255.255.255.255
```

- 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.

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

```
Router(config)#interface loopback number
```

Configuring an IP address on a loopback interface overrides the highest IP address 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, then 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.

OSPF router-id Command

```
Router(config-router)#
```

```
router-id ip-address
```

- This command is configured under the `router ospf [process-id]` command.
- Any unique arbitrary 32-bit value in an IP address format (dotted decimal) can be used.
- If this command is used on an OSPF process that is already active, then the new router ID is used after the next reload or manual OSPF process restart using:

```
Router#
```

```
clear ip ospf process
```

```
Router(config)#router ospf 1
```

```
Router(config-router)#router-id 172.16.1.1
```

```
Router#clear ip ospf process
```

© 2006 Cisco Systems, Inc. All rights reserved.BSCI v3.0–3.7

Use the OSPF **router-id** command to ensure that OSPF selects 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 have to be unique throughout the autonomous system (AS), no matter how they are configured.

Verifying the OSPF Router ID

This topic explains how to verify the OSPF router ID with the **show ip ospf** command.



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 the shortest path first (SPF) algorithm has been executed. This command also has optional parameters so that you can further specify the information to be displayed.

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

```
RouterB#sh ip ospf

Routing Process "ospf 50" with ID 10.64.0.2
Supports only single TOS(TOS0) routes
Supports opaque LSA
Supports Link-local Signaling (LLS)
Supports area transit capability
It is an area border router
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
  Area BACKBONE(0)
    Area BACKBONE(0)
      Area has no authentication
      SPF algorithm last executed 00:01:25.028 ago
      SPF algorithm executed 7 times
      Area ranges are
      Number of LSA 6. Checksum Sum 0x01FE3E
      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 00:00:54.636 ago
      SPF algorithm executed 3 times
      Area ranges are
      Number of LSA 4. Checksum Sum 0x01228A
      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

RouterB#
```

Verifying OSPF Operation

This topic explains how an OSPF configuration is verified using **show ip protocols**, **show ip route ospf**, **show ip ospf interface**, **show ip ospf**, **show ip ospf neighbor**, and **show ip route ospf** commands.

Verifying OSPF Operation

Router#
`show ip protocols`

- Verifies the configured IP routing protocol processes, parameters, and statistics

Router#
`show ip route ospf [process-id]`

- Displays all OSPF routes learned by the router

Router#
`show ip ospf interface [type number]`

- Displays the OSPF router ID, area ID, and adjacency information

© 2005 Cisco Systems, Inc. All rights reserved. BSCI v3.0—3-9

To verify that OSPF has been properly configured, use the following **show** commands:

- The **show ip protocols** command displays IP routing protocol parameters about timers, filters, metrics, networks, and other information for the entire router.
- The **show ip route ospf** command displays the OSPF routes 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 to be displayed, including the OSPF process ID.
- The **show ip ospf interface** command verifies 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.

Verifying OSPF Operation (Cont.)

Router#

```
show ip ospf
```

- Displays the OSPF router ID, timers, and statistics

Router#

```
show ip ospf neighbor [type number] [neighbor-id]  
[detail]
```

- Displays information about the OSPF neighbors, including DR and BDR information on broadcast networks

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—3-10

- As shown in the last topic, the **show ip ospf** command displays the OSPF router ID, OSPF timers, the number of times the SPF algorithm has been executed, and link-state advertisement (LSA) information.
- The **show ip ospf neighbor** command displays a list of neighbors, including their OSPF router ID, their OSPF priority, their neighbor adjacency state (for example, init, exstart, or full), and the dead timer.

Example: The show ip route ospf Command

```
RouterA#show ip route ospf
      10.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
O IA   10.2.1.0/24 [110/782] via 10.64.0.2, 00:03:05, FastEthernet0/0
RouterA#
```

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0--3-11

Example: The show ip route ospf Command

Use the **show ip route ospf** command to verify the OSPF routes in the IP routing table. The O code represents OSPF routes, and IA is “interarea.” In the figure, the 10.2.1.0 subnet is recognized on FastEthernet 0/0 via neighbor 10.64.0.2.

The entry [110/782] in the routing table represents the administrative distance assigned to OSPF (110), and the total cost of the route to subnet 10.2.1.0 (cost of 782).

Example: The show ip ospf interface Command

```
RouterA#show ip ospf interface fastEthernet 0/0
FastEthernet0/0 is up, line protocol is up
  Internet Address 10.64.0.1/24, Area 0
  Process ID 1, Router ID 10.64.0.1, Network Type BROADCAST, Cost: 1
  Transmit Delay is 1 sec, State DROTHER, Priority 0
  Designated Router (ID) 10.64.0.2, Interface address 10.64.0.2
  No backup designated router on this network
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    oob-resync timeout 40
    Hello due in 00:00:04
  Supports Link-local Signaling (LLS)
  Index 1/1, flood queue length 0
  Next 0x0(0)/0x0(0)
  Last flood scan length is 1, maximum is 4
  Last flood scan time is 0 msec, maximum is 4 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 10.64.0.2 (Designated Router)
  Suppress hello for 0 neighbor(s)
```

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—3-12

Example: The show ip ospf interface Command

The **show ip ospf interface** [*type number*] [*brief*] command displays OSPF-related interface information.

show ip ospf interface Parameters

The table contains information about the parameters of the **show ip ospf interface** command.

Parameter	Description
<i>type</i>	(Optional) Interface type
<i>number</i>	(Optional) Interface number
brief	Optional) Displays brief overview information for OSPF interfaces, states, addresses and masks, and areas on the router

The **show ip ospf interface** command output in the figure is from router A from the last configuration example and details the OSPF status of FastEthernet 0/0 interface. This command verifies that OSPF is running on this particular interface and gives the OSPF area that it is in.

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

Example: The show ip ospf neighbor Command

```
RouterB# show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
10.64.0.1	0	FULL/DROTHER	00:00:30	10.64.0.1	FastEthernet0/0
10.2.1.1	0	FULL/-	00:00:34	10.2.1.1	Serial0/0/1

```
RouterB# show ip ospf neighbor detail
```

```
Neighbor 10.64.0.1, interface address 10.64.0.1
  In the area 0 via interface FastEthernet0/0
  Neighbor priority is 0, State is FULL, 16 state changes
  DR is 10.64.0.2 BDR is 0.0.0.0
<output omitted>
```

```
Neighbor 10.2.1.1, interface address 10.2.1.1
  In the area 1 via interface Serial0/0/1
  Neighbor priority is 0, State is FULL, 6 state changes
  DR is 0.0.0.0 BDR is 0.0.0.0
<output omitted>
```

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0--3-13

Example: The show ip ospf neighbor Command

One of the most important OSPF 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** [*type number*] [*neighbor-id*] [**detail**] command displays OSPF neighbor information for each interface.

show ip ospf neighbor Parameters

The table contains information about the parameters of the **show ip ospf neighbor** command.

Parameter	Description
<i>type</i>	(Optional) Interface type
<i>number</i>	(Optional) Interface number
<i>neighbor-id</i>	(Optional) Neighbor ID
detail	(Optional) Displays details of all neighboring routers

In the first output in the figure, router B (from the last configuration example) has two neighbors. The first entry in the table represents the adjacency formed on the Fast Ethernet interface. A FULL state means that the link-state database (LSDB) has been exchanged successfully. The DROTHER entry means that a router other than this neighboring router is the designated router. (Note that the OSPF priority on the router A FastEthernet 0/0 interface has been set to 0, indicating that it cannot be the DR or BDR on that interface.)

The second line in the table represents the neighbor of router B 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 a partial output of the **show ip ospf neighbor detail** command output, providing details of the neighbors of router B. The full output is as follows:

```
RouterB#show ip ospf neighbor detail
```

```
Neighbor 10.64.0.1, interface address 10.64.0.1
```

```
  In the area 0 via interface FastEthernet0/0
```

```
  Neighbor priority is 0, State is FULL, 16 state changes
```

```
  DR is 10.64.0.2 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 00:07:14
```

```
  Index 2/2, 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
```

```
Neighbor 10.2.1.1, interface address 10.2.1.1
```

```
  In the area 1 via interface Serial0/0/1
```

```
  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 00:00:39
```

```
  Neighbor is up for 00:01:50
```

```
  Index 1/1, 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
```


Summary

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

Summary

- **Configuration of OSPF is a two-step process:**
 - **Enter OSPF configuration with the `router ospf` command.**
 - **Use the `network` command to describe which interfaces will run OSPF in which area.**
- **OSPF selects a router ID at startup time:**
 - **The router ID's specified in the `router-id` command under the OSPF process.**
 - **Otherwise, the highest IP address of a loopback interface, if there are any, is used.**
 - **By default, the highest IP address of all active interfaces**
- **Use the `show ip ospf` command to verify the router ID.**
- **Use the `show ip protocols`, `show ip route ospf`, `show ip ospf interface`, `show ip ospf`, and `show ip ospf neighbor` commands to verify OSPF operation.**

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—3-14

OSPF Network Types

Overview

Understanding that an Open Shortest Path First Protocol (OSPF) area is made up of different types of network links is important because the adjacency behavior is different for each network type, and OSPF must be properly configured to function correctly 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 under some network topologies.

This lesson describes each OSPF network type, how the adjacencies are formed over 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:

- Describe the three types of networks defined by OSPF
- Describe the adjacency behavior for point-to-point serial links
- Describe the adjacency behavior for a broadcast network link
- Describe how OSPF routers apply conditions to the OSPF priority values of the other routers during the hello packet exchange process to elect a DR and BDR
- Describe the adjacency behavior for an NBMA network
- Describe the configuration options for OSPF over Frame Relay
- Describe how to configure an NBMA topology in an OSPF over Frame Relay network
- Describe how to configure point-to-multipoint and point-to-multipoint nonbroadcast topologies in an OSPF over Frame Relay network
- Describe how to configure a physical interface into multiple subinterfaces
- Describe how to track OSPF adjacencies

Introducing OSPF Network Types

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

OSPF Network Types

The three types of networks defined by OSPF are:

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

© 2006 Cisco Systems, Inc. All rights reserved. BSCI v3.0–3.2

OSPF defines distinct types of networks, based on their physical link type. OSPF operation on each type is different, including how adjacencies are established and the configuration 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 that has no broadcast capability. Frame Relay, ATM, and X.25 are examples of NBMA networks. There are five modes of OSPF operation available for NBMA networks, as described later in this lesson.

OSPF operation and configuration on each of these network types is the focus of this lesson.

Adjacency Behavior for a Point-to-Point Link

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.
- No DR or BDR election required.
- OSPF autodetects this interface type.
- OSPF packets are sent using multicast 224.0.0.5.

© 2006 Cisco Systems, Inc. All rights reserved. BSCI v3.0—3-3

A point-to-point network joins a single pair of routers. A T1 serial line configured with a 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 because 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 for a Broadcast Network Link

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



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, for example, two routers connected on the same Ethernet segment. 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 a 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. Instead of each router exchanging link-state information with every other router on the segment, each router 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 added to the network establishes adjacencies with the DR and BDR only.

Selecting the DR and BDR

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



To elect a DR and BDR, the routers view the OSPF priority value of the 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 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 gets added to the network, it does not preempt 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, then 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 connected to multiple multiaccess broadcast networks can be a DR on one segment and a regular router on another segment.

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

Setting Priority for DR Election

```
Router(config-if)#
```

```
ip ospf priority number
```

- 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.

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—3-6

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 set to zero priority cannot be involved in the DR or BDR election process.

Here is a configuration example:


```
Router(config)#interface FastEthernet 0/0
```

```
Router(config-if)#ip ospf priority 10
```

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.

Adjacency Behavior for an NBMA Network

This topic describes the adjacency behavior for an NBMA network.



When a single interface interconnects multiple sites over an NBMA network, the nonbroadcast nature of the network can create 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 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 on 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 need a list of neighbors.**
- **OSPF neighbors are not automatically discovered by the router.**

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—3-8


OSPF considers that the NBMA environment functions in a way similar to other broadcast media 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 full 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.

OSPF over Frame Relay Configuration Options

This topic describes the configuration options for OSPF over Frame Relay.



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 a variety of 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

- **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**

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0--3-10

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 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

Selecting the OSPF Network Type for NBMA Networks

```
Router(config-if)#
```

```
ip ospf network [{broadcast | non-broadcast | point-to-  
multipoint [non-broadcast] | point-to-point}]
```

- Defines OSPF network type

Example: Broadcast Mode

```
Router(config)#interface serial 0/0/0  
Router(config-if)#encapsulation frame-relay  
Router(config-if)#ip ospf network broadcast
```

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—3-11

The **ip ospf network** interface command has several options.

ip ospf network Command Options

The table describes the **ip ospf network** interface command options.

Command Options	Description
broadcast	Cisco extension. ■ Makes the WAN interface appear to be a LAN. ■ One IP subnet. ■ Uses multicast OSPF hello packet to automatically discover the neighbors. ■ DR and BDR elected. ■ Requires a full-mesh or a partial-mesh topology.
non-broadcast	RFC-compliant mode. ■ One IP subnet. ■ Neighbors must be manually configured. ■ DR and BDR elected. ■ DR and BDR need to have full connectivity with all other routers. ■ Typically used in a full-mesh or a partial-mesh topology.
point-to-multipoint	RFC-compliant mode. ■ One IP subnet. ■ Uses multicast OSPF hello packet to automatically discover the neighbors. ■ DR and BDR not required—router sends additional LSAs with more information about neighboring routers. ■ Typically used in partial-mesh or star topology.
point-to-multipoint non-broadcast	Cisco extension. ■ If 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 hello multicast packets; this Cisco mode should be used instead. ■ Neighbors must be manually configured. ■ DR and BDR election is not required.
point-to-point	Cisco extension. ■ Different IP subnet on each subinterface. ■ No DR or BDR election. ■ Used when only two routers need to form an adjacency on a pair of interfaces. ■ Interfaces can be either LAN or WAN.

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.

Example: Sample Configuration of a Router Using OSPF Broadcast Mode


The previous figure shows a sample configuration of a Frame Relay router in a full-mesh topology, with the broadcast mode of operation defined.

OSPF over Frame Relay NBMA Configuration

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 (acts like a LAN).
- All serial ports are part of the same IP subnet.
- Frame Relay, X.25, and ATM networks default to nonbroadcast mode.
- Neighbors must be statically configured.
- Duplicates LSA updates.
- Complies with RFC 2328.



© 2006 Cisco Systems, Inc. All rights reserved. BSCI v3.0—3-12

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, then you should select the DR and BDR manually to ensure that the selected DR and BDR have full 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 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 the point-to-multipoint mode.

Frame Relay, ATM, and X.25 networks default to OSPF nonbroadcast mode.

Using the neighbor Command

```
Router(config-router)#
```

```
neighbor ip-address [priority number] [poll-interval  
number] [cost number] [database-filter all]
```

- Used to statically define neighbor relationships in an NBMA network

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0--3-13


The OSPF **neighbor** command is used to statically define adjacent relationships in NBMA networks using the nonbroadcast mode.

neighbor Command Options

The **neighbor** command has the options shown in the table.

Command	Description
<i>ip-address</i>	IP address of the neighboring router.
priority <i>number</i>	(Optional) Specifies priority of neighbor. The default is 0, which means that the neighboring router does not become the DR.
poll-interval <i>number</i>	(Optional) Amount of time an NBMA interface waits before sending hellos to the neighbor even if the neighbor is inactive. The poll interval is defined in seconds.
cost <i>number</i>	(Optional) Assigns a cost to the neighbor, in the form of an integer from 1 to 65535. Neighbors with no specific cost configured will assume the cost of the interface, based on the ip ospf cost command. For point-to-multipoint interfaces, the cost keyword and the <i>number</i> argument are the only options that are applicable. This keyword does not apply to NBMA networks.
database-filter all	(Optional) Filters outgoing LSAs to an OSPF neighbor.

neighbor Command Example



```
RouterA(config)# router ospf 100
RouterA(config-router)# network 192.168.0.0 0.0.255.255 area 0
RouterA(config-router)# neighbor 192.168.1.2 priority 0
RouterA(config-router)# neighbor 192.168.1.3 priority 0
RouterA(config-router)# network 172.16.0.0 0.0.255.255 area 0
```

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—3-14

Example: neighbor Command


This figure illustrates an example of statically defining adjacencies. 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 B and C because a full-mesh topology does not exist. This configuration ensures that router A becomes the DR because only router A has full connectivity to the other two routers. No BDR will be elected in this case.

Note The default priority on the neighbor command is 0. However, during testing it was noted that configuring the priority in this way resulted in a priority of 1, not 0. Setting the OSPF priority to 0 at the interface level on routers B and C did result in a priority of 0 and the routers not being elected as 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



```
RouterA# 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

© 2006 Cisco Systems, Inc. All rights reserved.

BSOI v3.0--3-15

The **show ip ospf neighbor** [*type number*] [*neighbor-id*] [**detail**] command displays OSPF neighbor information, as illustrated in the figure.

show ip ospf neighbor Parameters

The table describes the parameters of the **show ip ospf neighbor** command.

Parameter	Description
<i>type</i>	(Optional) Interface type
<i>number</i>	(Optional) Interface number
<i>neighbor-id</i>	(Optional) ID for neighboring router
detail	(Optional) Displays details of all neighboring routers

Example: show ip ospf neighbor Command


Router A in the figure 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 Router A is the DR and there is no BDR in this network).

The neighbor learned on FastEthernet 0/0 show 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.

OSPF over Frame Relay Point-to-Multipoint Configuration

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. In 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 the point-to-multipoint mode, DRs are not used, and a type 2 network 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 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 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

Router A

```
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
```

Router C

```
interface Serial0/0/0
ip address 192.168.1.3 255.255.255.0
encapsulation frame-relay
ip ospf network point-to-multipoint
ip ospf priority 0
```

© 2006 Cisco Systems, Inc. All rights reserved.BSCI v3.0—3-17

Example: Point-to-Multipoint Configuration

This figure shows partial configurations of routers A and C in point-to-multipoint mode. This configuration does not require subinterfaces and uses only a single subnet.

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

Point-to-Multipoint Example

```
RouterA#sh ip ospf int 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_MULTIPOINT,
  Cost: 781
  Transmit Delay is 1 sec, State POINT_TO_MULTIPOINT
  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)
RouterA#
```

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—3-18

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 in order for the neighbors to form successful adjacencies.

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

Point-to-Multipoint Nonbroadcast

- **Cisco extension to RFC-compliant point-to-multipoint mode**
- **Must statically define neighbors, like nonbroadcast mode**
- **Like point-to-multipoint mode, DR and BDR not elected**
- **Used in special cases where neighbors cannot be automatically discovered**

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0--3-19

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. If 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.

Using Subinterfaces in OSPF over Frame Relay Configuration

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

Using Subinterfaces

```
Router(config)#  
interface serial number.subinterface-number {multipoint |  
point-to-point}
```

- The physical serial port becomes multiple logical ports.
- Each subinterface requires an IP subnet.

© 2006 Cisco Systems, Inc. All rights reserved. BSCI v3.0—3-20

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 caused by split horizon over an NBMA for distance vector-based routing protocols. Each subinterface requires an IP subnet.

Use the **interface serial** command to define subinterfaces.

interface serial Parameters

The table lists the parameters of the **interface serial** command.

Parameter	Description
<i>number.subinterface-number</i>	■ Interface number and subinterface number. ■ Subinterface number is in the range of 1 to 4294967293. ■ Interface number that precedes the period (.) must match the interface number to which this subinterface belongs.
multipoint	■ On multipoint subinterfaces routing IP, all routers are in the same subnet.
point-to-point	■ On point-to-point subinterfaces routing IP, each pair of point-to-point routers is in its own subnet.

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.

Point-to-Point Subinterfaces

Router(config)#

```
interface serial number.subinterface-number point-to-point
```

- **Each PVC and SVC gets its own subinterface.**
- **OSPF point-to-point mode is the default on point-to-point Frame Relay subinterfaces.**
 - **No DR/BDR**
 - **Do not need to configure neighbors**


© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—3-21

When point-to-point subinterfaces are configured, each virtual circuit (PVC and SVC) gets its own subinterface.

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

Point-to-Point Subinterface Example



- PVCs are treated like point-to-point links.
- Each subinterface requires a subnet.

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—3-22

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

Multipoint Subinterfaces

Router(config)#

```
interface serial number.subinterface-number multipoint
```

- **Multiple PVCs and SVCs are on a single subinterface.**
- **OSPF nonbroadcast mode is the default.**
 - **DR and BDR are required.**
 - **Neighbors need to be statically configured.**


© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0--3-23

When multipoint subinterfaces are configured, there are multiple virtual circuits (PVCs or SVCs) on a single subinterface.

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

Multipoint Subinterface Example



- Single interface serial 0/0/0 has been logically separated into two subinterfaces: one point-to-point (S0/0/0.1) and one point-to-multipoint (S0/0/0.2).
- Each subinterface requires a subnet.
- OSPF defaults to point-to-point mode on point-to-point subinterfaces.
- OSPF defaults to nonbroadcast mode on point-to-multipoint subinterfaces.

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—3-24

Example: Multipoint Subinterface

In this figure, router A has one point-to-point subinterface and a second point-to-multipoint subinterface. The multipoint subinterface supports two other routers in a single subnet.

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

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

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/BDR elected	Cisco
Nonbroadcast (NBMA)	Full or partial mesh	Same	30 sec	Manual configuration, DR/BDR elected	RFC
Point-to-multipoint	Partial-mesh or star	Same	30 Sec	Automatic, no DR/BDR	RFC
Point-to-multipoint nonbroadcast	partial-mesh or star	Same	30 sec	Manual configuration, no DR/BDR	Cisco
Point-to-point	Partial-mesh or star, using subinterface	Different for Each Subinterface	10 sec	Automatic, no DR/BDR	Cisco

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—3-25

Example: OSPF over NBMA Topology Summary

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

Tracking OSPF Adjacencies

Understanding OSPF adjacency protocol handshakes is important in troubleshooting. This topic describes how to interpret the output from the **debug** command.

Creation of Adjacencies for Point-to-Point Mode

```
RouterA# debug ip ospf adj
OSPF: Interface Serial0/0/0.1 going Up
OSPF: Build router LSA for area 0, router ID 192.168.1.1, seq 0x80000023
OSPF: Rcv DBD from 192.168.1.2 on Serial0/0/0.1 seq 0xCF0 opt 0x52 flag 0x7 len 32
mtu 1500 state INIT
OSPF: 2 Way Communication to 192.168.1.2 on Serial0/0/0.1, state 2WAY
OSPF: Send DBD to 192.168.1.2 on Serial0/0/0.1 seq 0xF4D opt 0x52 flag 0x7 len 32
OSPF: NBR Negotiation Done. We are the SLAVE
OSPF: Send DBD to 192.168.1.2 on Serial0/0/0.1 seq 0xCF0 opt 0x52 flag 0x2 len 132
OSPF: Rcv DBD from 192.168.1.2 on Serial0/0/0.1 seq 0xCF1 opt 0x52 flag 0x3 len 132
mtu 1500 state EXCHANGE
OSPF: Send DBD to 192.168.1.2 on Serial0/0/0.1 seq 0xCF1 opt 0x52 flag 0x0 len 32
OSPF: Database request to 192.168.1.2
OSPF: sent LS REQ packet to 192.168.1.2, length 12
OSPF: Rcv DBD from 192.168.1.2 on Serial0/0/0.1 seq 0xCF2 opt 0x52 flag 0x1 len 32
mtu 1500 state EXCHANGE
OSPF: Exchange Done with 192.168.1.2 on Serial0/0/0.1
OSPF: Send DBD to 192.168.1.2 on Serial0/0/0.1 seq 0xCF2 opt 0x52 flag 0x0 len 32
OSPF: Synchronized with 192.168.1.2 on Serial0/0/0.1, state FULL
%OSPF-5-ADJCHG: Process 100, Nbr 192.168.1.2 on Serial0/0/0.1 from LOADING to FULL,
Loading Done
OSPF: Build router LSA for area 0, router ID 192.168.1.1, seq 0x80000024
```

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0-3.25

Use the **debug ip ospf adj** command to track OSPF adjacencies as they come up or go down. Debugging allows you to see exactly which OSPF packets are being sent between routers. The ability to see packets as they are sent over a link is an invaluable troubleshooting tool.

Example: debug Output for Point-to-Point Mode

The figure shows output from the **debug ip ospf adj** command that describes a serial interface in point-to-point mode. No DR election occurs; however, the adjacency forms, allowing database description (DBD) packets to be sent during the exchange process.

Notice that the neighbor relationship passes through the two-way phase and into the exchange phase. After DBD packets are sent between routers, the neighbors will move into the final state: full adjacency.

Creation of Adjacencies for Broadcast Mode

```
RouterA# debug ip ospf adj
OSPF: Interface FastEthernet0/0 going Up
OSPF: Build router LSA for area 0, router ID 192.168.1.1,seq 0x80000008
OSPF: 2 Way Communication to 172.16.1.1 on FastEthernet0/0, state 2WAY
OSPF: end of Wait on interface FastEthernet0/0
<output omitted>

OSPF: Neighbor change Event on interface FastEthernet0/0
OSPF: DR/BDR election on FastEthernet0/0
OSPF: Elect BDR 172.16.1.1
OSPF: Elect DR 192.168.1.1
DR: 192.168.1.1 (Id)   BDR: 172.16.1.1 (Id)
OSPF: Rcv DBD from 172.16.1.1 on FastEthernet0/0 seq 0x14B 7 opt 0x52 flag 0x7
len 32  mtu 1500 state EXSTART
OSPF: First DBD and we are not SLAVE-if)#
OSPF: Send DBD to 172.16.1.1 on FastEthernet0/0 seq 0xDCE opt 0x52 flag 0x7
len 32
OSPF: Retransmitting DBD to 172.16.1.1 on FastEthernet0/0[1]
OSPF: Rcv DBD from 172.16.1.1 on FastEthernet0/0 seq 0xDCE
opt 0x52 flag 0x2 len 152  mtu 1500 state EXSTART
<output omitted>
```

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0--3-27

Example: debug ip ospf adj Output for Broadcast Mode

The figure shows partial **debug ip ospf adj** output illustrating the DR and BDR election process on a Fast Ethernet interface. The OSPF default behavior on a Fast Ethernet link is broadcast mode. First, the DR and BDR are selected, and then the exchange process occurs.

The full output from this command is as follows:

```
RouterA# debug ip ospf adj

OSPF: Interface FastEthernet0/0 going Up

OSPF: Build router LSA for area 0, router ID 192.168.1.1,seq
0x80000008

%LINEPROTO-5-UPDOWN: Line protocol on Interface
FastEthernet0/0, changed state to up

OSPF: 2 Way Communication to 172.16.1.1 on FastEthernet0/0,
state 2WAY

OSPF: end of Wait on interface FastEthernet0/0

OSPF: DR/BDR election on FastEthernet0/0

OSPF: Elect BDR 192.168.1.1

OSPF: Elect DR 192.168.1.1

OSPF: Elect BDR 172.16.1.1

OSPF: Elect DR 192.168.1.1

DR: 192.168.1.1 (Id)   BDR: 172.16.1.1 (Id)

OSPF: Send DBD to 172.16.1.1 on FastEthernet0/0 seq 0xDCE opt
0x52 flag 0x7 len 32

OSPF: No full nbrs to build Net Lsa for interface
FastEthernet0/0
```

```

OSPF: Neighbor change Event on interface FastEthernet0/0
OSPF: DR/BDR election on FastEthernet0/0
OSPF: Elect BDR 172.16.1.1
OSPF: Elect DR 192.168.1.1
DR: 192.168.1.1 (Id)    BDR: 172.16.1.1 (Id)
OSPF: Neighbor change Event on interface FastEthernet0/0
OSPF: DR/BDR election on FastEthernet0/0
OSPF: Elect BDR 172.16.1.1
OSPF: Elect DR 192.168.1.1
DR: 192.168.1.1 (Id)    BDR: 172.16.1.1 (Id)
OSPF: Rcv DBD from 172.16.1.1 on FastEthernet0/0 seq 0x14B 7
opt 0x52 flag 0x7 len 32  mtu 1500 state EXSTART
OSPF: First DBD and we are not SLAVE-if)#
OSPF: Send DBD to 172.16.1.1 on FastEthernet0/0 seq 0xDCE opt
0x52 flag 0x7 len 32
OSPF: Retransmitting DBD to 172.16.1.1 on FastEthernet0/0[1]
OSPF: Rcv DBD from 172.16.1.1 on FastEthernet0/0 seq 0xDCE
    opt 0x52 flag 0x2 len 152  mtu 1500 state EXSTART
OSPF: NBR Negotiation Done. We are the MASTER
OSPF: Send DBD to 172.16.1.1 on FastEthernet0/0 seq 0xDCf
opt 0x52 flag 0x3 len 132
OSPF: Database request to 172.16.1.1
OSPF: sent LS REQ packet to 172.16.1.1, length 24
OSPF: Rcv DBD from 172.16.1.1 on FastEthernet0/0 seq 0xDCf
    opt 0x52 flag 0x0 len 32  mtu 1500 state EXCHANGE
OSPF: Send DBD to 172.16.1.1 on FastEthernet0/0 seq 0xDD0
opt 0x52 flag 0x1 len 32
OSPF: No full nbrs to build Net Lsa for interface
FastEthernet0/0
OSPF: Build network LSA for FastEthernet0/0, router ID
192.168.1.1
OSPF: Build network LSA for FastEthernet0/0, router ID 192
.168.1.1
OSPF: Rcv DBD from 172.16.1.1 on FastEthernet0/0 seq 0xDD0
    opt 0x52 flag 0x0 len 32  mtu 1500 state EXCHANGE
OSPF: Exchange Done with 172.16.1.1 on FastEthernet0/0
OSPF: Synchronized with 172.16.1.1 on FastEthernet0/0, state
FULL
%OSPF-5-ADJCHG: Process 100, Nbr 172.16.1.1 on FastEthernet0/0
from LOADING to FULL, Loading Done
OSPF: Build router LSA for area 0, router ID 192.168.1.1,
seq 0x80000009

```



```
OSPF: Build network LSA for FastEthernet0/0, router ID 192
.168.1.1
OSPF: Build network LSA for FastEthernet0/0, router ID 192
.168.1.1
```

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, adjacency is dynamic, uses multicast addresses, and has no DR or BDR.**
- **On broadcast links, 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.**
- **The router with the highest OSPF priority is selected as the DR. The router with the second-highest priority value is the BDR.**
- **By default on NBMA links, adjacency requires manual definition of neighbors for the DR and BDR because OSPF will consider the network similar to broadcast media.**

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—3-28

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.**
- **In nonbroadcast mode, a DR and BDR are elected, and neighbors must be statically configured.**
- **In point-to-multipoint mode, no DR and BDR are needed and neighbors are automatically discovered. In point-to-multipoint nonbroadcast mode, no DR and BDR are needed, but neighbors must be statically configured.**
- **A physical interface can be split into multiple logical interfaces called subinterfaces. Each subinterface requires an IP subnet.**
- **Using the debug ip ospf adj command enables you to see OSPF packet exchanges and the status of neighbor adjacencies.**

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—3-29

Link-State Advertisements

Overview

There is nothing more important in Open Shortest Path First Protocol (OSPF) than understanding how the topology database is built. Troubleshooting OSPF often requires analyzing the OSPF database and routing table; therefore, a solid understanding of the OSPF database and link-state advertisements (LSAs) is essential.

The two core concepts of OSPF are the link-state database (LSDB) and LSAs. This lesson describes each of the common LSA types and how they form the layout of the OSPF LSDB.

This lesson also describes OSPF virtual links and OSPF router types, including internal routers, backbone routers, area border routers (ABRs), and autonomous system boundary routers (ASBRs). OSPF LSDB overload protection is explored, as is how the cost metric is changed.


Objectives

Upon completing this lesson, you will be able to describe how LSAs and OSPF databases maintain links through the network. This ability includes being able to meet these objectives:

- Describe the different OSPF router types
- Describe OSPF virtual links
- Describe the LSAs defined by OSPF
- Describe how to interpret the OSPF LSDB and routing table
- Configure OSPF LSDB overload protection
- Change the cost metric from default values

OSPF Router Types

OSPF LSDBs are often very large. For this reason, an area hierarchical structure has been imposed that defines several router types. This topic describes the different OSPF router types.




OSPF can usually operate within a single area; however, certain issues arise if this single area expands into hundreds of networks. If an expansion occurs, the following issues need to be addressed:

- **Frequent shortest path first (SPF) algorithm calculations:** In a large network, changes are inevitable; therefore, the routers spend many CPU cycles recalculating the SPF algorithm and updating the routing table.
- **Large routing table:** OSPF does not perform route summarization by default. If the routes are not summarized, the routing table can become very large, depending on the size of the network.
- **Large LSDB:** Because the LSDB covers the topology of the entire network, each router must maintain an entry for every network in the area, even if not every route is selected for the routing table.

A solution to these issues is to divide the network into multiple OSPF areas. OSPF allows the separation of a large area into smaller, more manageable areas that are still able to exchange routing information.

The Solution: OSPF Hierarchical Routing



- **Consists of areas and autonomous systems**
- **Minimizes routing update traffic**

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—3-3

Hierarchical area routing is the ability of OSPF to separate a large internetwork into multiple areas. When you use this technique, interarea routing still occurs, but many of the internal routing operations, such as SPF calculations, remain within individual areas.


Example: OSPF Hierarchical Routing

For example, in the figure, if area 1 is having problems with a link going up and down, routers in other areas do not need to continually run their SPF calculation, because they are isolated from the problem in area 1.

Given a proper IP addressing hierarchy, using multiple OSPF areas has several important advantages:

- **Reduced frequency of SPF calculations:** Because detailed route information exists within each area, it is not necessary to flood all link-state changes to all other areas. Therefore, only the routers that are affected by the change need to recalculate SPF.
- **Smaller routing tables:** With multiple areas, detailed route entries for specific networks within an area can remain in the area. Instead of advertising these explicit routes outside the area, routers can be configured to summarize the routes into one or more summary addresses. Advertising these summaries reduces the number of LSAs propagated between areas but keeps all networks reachable.
- **Reduced link-state update (LSU) overhead:** LSUs contain a variety of LSA types, including link-state and summary information. Rather than send an LSU about each network within an area, a router can advertise a single summarized route or small number of routes between areas, reducing the overhead associated with LSUs when they cross areas.

Types of OSPF Routers



Certain types of OSPF routers control the traffic types that go in and out of various areas. The following are the four router types:

- **Internal routers:** Routers that have all their interfaces in the same area and have identical LSDBs.
- **Backbone routers:** Routers that sit in the perimeter of the backbone area and have at least one interface connected to area 0. Backbone routers maintain OSPF routing information using the same procedures and algorithms as internal routers.
- **ABRs:** Routers that have interfaces attached to multiple areas, maintain separate LSDBs for each area to which they connect, and route traffic destined for or arriving from other areas. ABRs are exit points for the area, which means that routing information destined for another area can get there only via the ABR of the local area.

ABRs can be configured to summarize the routing information from the LSDBs of their attached areas. ABRs distribute the routing information into the backbone. The backbone routers then forward the information to the other ABRs. In a multiarea network, an area can have one or more ABRs.

- **ASBRs:** Routers that have at least one interface attached to an external internetwork (another autonomous system [AS]), such as a non-OSPF network. ASBRs can import non-OSPF network information to the OSPF network and vice versa; this process is called route redistribution.


A router can exist as more than one router type. For example, if a router interconnects to area 0 and area 1, as well as to a non-OSPF network, it is both an ABR and an ASBR.

A router has a separate LSDB for each area to which it connects; therefore, an ABR could have one LSDB for area 0 and another LSDB for another area in which it participates. Two routers belonging to the same area maintain identical LSDBs for that area.

An LSDB is synchronized between pairs of adjacent routers. On broadcast networks like Ethernet, an LSDB is synchronized between the router that is not a designated router (DR) or a backup designated router (BDR) (that is, a DROTHER) and its DR and BDR.

OSPF Virtual Links

This topic describes OSPF virtual links and how to configure them.



The two-tiered area hierarchy of OSPF requires that all areas directly connect to the backbone area, area 0, and that area 0 be contiguous.

A virtual link is a link that allows discontinuous area 0s 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 Software Release 10.0. In the figure, area 0 is discontinuous because of a network failure. A logical link (virtual link) is built between the two ABRs, routers A and B. 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 learned through a virtual link have the DoNotAge (DNA) option set, so that the LSA does not age out. This DNA technique is required to prevent excessive flooding over the virtual link.

Configuring Virtual Links

```
Router(config-router)#
```

```
area area-id virtual-link router-id [authentication  
[message-digest | null]] [hello-interval seconds]  
[retransmit-interval seconds] [transmit-delay  
seconds] [dead-interval seconds] [[authentication-  
key key] | [message-digest-key key-id md5 key]]
```

Creates a virtual link

```
remoterouter#sh ip ospf  
Routing Process "ospf 1000" with ID 10.2.2.2  
Supports only single TOS(TOS0) routes  
Supports opaque LSA  
Supports Link-local Signaling (LLS)  
Supports area transit capability  
It is an area border router  
<output omitted>
```

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—3-6

Use the **area area-id virtual-link router-id** router configuration command, along with any necessary optional parameters, to define an OSPF virtual link. To remove a virtual link, use the **no** form of this command.


The **area virtual-link** command includes the router ID of the far-end router. To find the router ID in the far-end router, use the **show ip ospf** command, **show ip ospf interface** command, or **show ip protocol** command on that remote router, as illustrated in the figure.

area virtual-link Command Parameters

The table describes the options available with the **area virtual-link** command.

Parameter	Description
<i>area-id</i>	Assigns an area ID to the transit area for the virtual link. This ID can be either a decimal value or in dotted-decimal format like a valid IP address. There is no default. The transit area cannot be a stub area.
<i>router-id</i>	Router ID of the virtual link neighbor. The router ID appears in the show ip ospf command display. There is no default.
authentication	(Optional) Specifies an authentication type.
message-digest	(Optional) Specifies the use of message-digest authentication.
null	(Optional) Overrides password or message-digest authentication if configured for the area. Authentication is not used.
hello-interval <i>seconds</i>	(Optional) Specifies the time (in seconds) between the hello packets that the Cisco IOS software sends on an interface. This unsigned integer value is advertised in the hello packets. The value must be the same for all routers and access servers attached to a common network. The default is 10 seconds.
retransmit-interval <i>seconds</i>	(Optional) Specifies the time (in seconds) between LSA retransmissions for adjacencies belonging to the interface. The value must be greater than the expected round-trip delay between any two routers on the attached network. The default is 5 seconds.
transmit-delay <i>seconds</i>	(Optional) Specifies the estimated time (in seconds) to send an LSU packet on the interface. This integer value must be greater than zero. LSAs in the update packet have their age incremented by this amount before transmission. The default value is 1 second.
dead-interval <i>seconds</i>	(Optional) Specifies the time (in seconds) that must pass without hello packets being seen before a neighboring router declares the router down. There is an unsigned integer value. The default is four times the default hello interval, or 40 seconds. As with the hello interval, this value must be the same for all routers and access servers attached to a common network.
authentication-key <i>key</i>	(Optional) Specifies the password used by neighboring routers.
message-digest-key <i>key-id md5 key</i>	(Optional) Identifies the key ID and password used between this router and neighboring routers for MD5 authentication. The <i>key-id</i> argument is a number in the range from 1 to 255. The <i>key</i> argument is an alphanumeric string of up to 16 characters. All neighboring routers on the same network must have the same key ID and key to route OSPF traffic. There is no default value.

OSPF Virtual Link Configuration Example



Example: OSPF Virtual Link Configuration

In the figure, area 0 is discontinuous because of a network failure. A virtual link is used as a backup strategy to temporarily reconnect area 0; area 1 is used as the transit area.

Router A builds a virtual link to router B, and router B builds a virtual link to the router A. Each router points at the router ID of the other router.

The show ip ospf virtual-links Command

```
RouterA#sh 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
RouterA#
```

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0–3-8

Example: show ip ospf virtual-links Command

Use the **show ip ospf virtual-links** command to verify that the configured virtual link works properly.

show ip ospf virtual-links Command Fields

The table describes the fields displayed in the command output in the figure.

Field	Description
Virtual Link OSPF_VL0 to router 10.2.2.2 is up	Specifies the OSPF neighbor and whether the link to that neighbor is up or down
Transit area 1	Specifies the transit area through which the virtual link is formed
via interface Serial0/0/1	Specifies the interface through which the virtual link is formed
Cost of using 781	Specifies the cost of reaching the OSPF neighbor through the virtual link
Transmit Delay is 1 sec	Specifies the transmit delay on the virtual link
State POINT_TO_POINT	Specifies the state of the OSPF neighbor
Timer intervals configured	Specifies the various timer intervals that are configured for the link
Hello due in 0:00:07	Specifies when the next hello is expected from the neighbor
Adjacency State FULL	Specifies the adjacency state between the neighbors

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

```
RouterA#sh ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
10.200.200.13	1	FULL/DR	00:00:33	10.1.1.3	FastEthernet0/0
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

```
RouterA#
```

```
RouterA#sh 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

RouterA#

OSPF LSA Types

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 defined by OSPF.

LSA Types	
LSA Type	Description
1	Router LSAs
2	Network LSAs
3 or 4	Summary LSAs
5	Autonomous system external LSAs
6	Multicast OSPF LSA
7	Defined for not-so-stubby areas
8	External attributes LSA for Border Gateway Protocol (BGP)
9, 10, 11	Opaque LSAs

© 2006 Cisco Systems, Inc. All rights reserved. BSCI v3.0—3-9

The following are descriptions of each type of LSA. LSA types 1 through 5 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 links of the router to the area and are flooded only within a 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 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. These 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 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 is an LSA type that is used in NSSAs.


Type 8

Type 8 is a specialized LSA that is 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 Systems 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.

LSA Type 1: Router LSA



- **One router LSA (type 1) for every router in an area**
 - Includes list of directly attached links
 - Each link identified by IP prefix assigned to link and link type
- **Identified by the router ID of the originating router**
- **Floods within its area only; does not cross 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 types: type 1, 2, 3, or 4. The LSA includes a link ID field that identifies, by the network number and mask, the object that this link connects to.

LSA Type 1 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/subnet number
4	Virtual link	Neighboring router ID

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

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

LSA Type 2: Network LSA



- **One network (type 2) LSA for each transit broadcast or NBMA network in an area**
 - Includes list of attached routers on the transit link
 - Includes subnet mask of link
- **Advertised by the DR of the broadcast network**
- **Floods within its area only; does not cross ABR**


© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—3-11

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, as well as the subnet mask 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.

LSA Type 3: Summary LSA



- **Type 3 LSAs are used to flood network information to areas outside the originating area (interarea)**
 - Describes network number and mask of link.
- **Advertised by the ABR of originating area.**
- **Regenerated by subsequent ABRs to flood throughout the autonomous system.**
- **By default, routes are not summarized, and type 3 LSA is advertised for every subnet.**

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0--3-12

The ABR sends type 3 summary LSAs. Type 3 LSAs advertise any networks 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, or even 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 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	Summary LSAs do not, by default, contain summarized routes.

LSA Type 4: Summary LSA



- Summary (type 4) LSAs are used to advertise an ASBR to all other areas in the autonomous system.
- They are generated by the ABR of the originating area.
- They are regenerated by all subsequent ABRs to flood throughout the autonomous system.
- Type 4 LSAs contain the router ID of the ASBR.

© 2006 Cisco Systems, Inc. All rights reserved.


BSCI v3.0—3-13

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 destined to an external AS requires routing table knowledge of the ASBR that originated the external routes.

Example: LSA Type 4—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.

LSA Type 5: External LSA



- External (type 5) LSAs are used to advertise networks from other autonomous systems.
- Type 5 LSAs are advertised and owned by the originating ASBR.
- Type 5 LSAs flood throughout the entire autonomous system.
- The advertising router ID (ASBR) is unchanged throughout the autonomous system.
- Type 4 LSA is needed to find the ASBR.
- By default, routes are not summarized.

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0--3-14

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.

Interpreting the OSPF LSDB and Routing Table

OSPF LSDB and the IP routing table are two of the most important concepts to understand for OSPF operation. This topic describes how to interpret the OSPF LSDB and routing table.

Interpreting the OSPF Database

```
RouterA#show ip ospf database
      OSPF Router with ID (10.0.0.11) (Process ID 1)
      Router Link States (Area 0)
Link ID      ADV Router      Age      Seq#          Checksum Link count
10.0.0.11    10.0.0.11          548      0x80000002    0x00401A  1
10.0.0.12    10.0.0.12          549      0x80000004    0x003A1B  1
100.100.100.100 100.100.100.100 548      0x800002D7    0x00EEA9  2
      Net Link States (Area 0)
Link ID      ADV Router      Age      Seq#          Checksum
172.31.1.3   100.100.100.100 549      0x80000001    0x004EC9
      Summary Net Link States (Area 0)
Link ID      ADV Router      Age      Seq#          Checksum
10.1.0.0     10.0.0.11      654      0x80000001    0x00FB11
10.1.0.0     10.0.0.12      601      0x80000001    0x00F516
<output omitted>
```

© 2006 Cisco Systems, Inc. All rights reserved. BSCI v3.0—3-15

Example: Interpreting the OSPF Database

The figure illustrates use of the **show ip ospf database** command to get information about an OSPF LSDB. The router link states are type 1 LSAs, the net link states are type 2 LSAs, and the summary net link states are type 3 LSAs.

The database columns are as follows:

- **Link ID:** Identifies each LSA.
- **ADV router:** Advertising router; the source router of the LSA.
- **Age:** The maximum age counter in seconds; the maximum age is 1 hour, or 3,600 seconds.
- **Seq#:** Sequence number of the LSA; this number begins at 0x80000001 and increases with each update of the LSA.
- **Checksum:** Checksum of the individual LSA to ensure reliable receipt of that LSA.
- **Link count:** Total number of directly attached links, 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. The full command output from this router is as follows:

RouterA#show ip ospf database

OSPF Router with ID (10.0.0.11) (Process ID 1)

Router Link States (Area 0)

Link ID count	ADV Router	Age	Seq#	Checksum	Link
10.0.0.11	10.0.0.11	548	0x80000002	0x00401A	1
10.0.0.12	10.0.0.12	549	0x80000004	0x003A1B	1
100.100.100.100	100.100.100.100	548	0x800002D7	0x00EEA9	2

Net Link States (Area 0)

Link ID	ADV Router	Age	Seq#	Checksum
172.31.1.3	100.100.100.100	549	0x80000001	0x004EC9

Summary Net Link States (Area 0)

Link ID	ADV Router	Age	Seq#	Checksum
10.1.0.0	10.0.0.11	654	0x80000001	0x00FB11
10.1.0.0	10.0.0.12	601	0x80000001	0x00F516
10.1.1.0	10.0.0.11	7	0x80000009	0x004DC5
10.1.1.0	10.0.0.12	9	0x80000007	0x00E81B
10.1.1.0	172.31.1.1	1111	0x80000003	0x00DD82
10.1.2.0	10.0.0.11	599	0x80000003	0x00EB1C
10.1.2.0	10.0.0.12	603	0x80000001	0x004CCC
10.1.3.0	10.0.0.11	14	0x80000002	0x00E225
10.1.3.0	10.0.0.12	69	0x80000001	0x00DE29
10.200.200.13	172.31.1.1	1108	0x80000001	0x00764E

Router Link States (Area 1)

Link ID count	ADV Router	Age	Seq#	Checksum	Link
10.0.0.11	10.0.0.11	19	0x80000009	0x00B6C3	3
10.0.0.12	10.0.0.12	601	0x80000005	0x0085F0	3
10.200.200.13	10.200.200.13	20	0x80000003	0x000AB2	3

10.200.200.14	10.200.200.14	62	0x8000004D	0x003C2E	3

Net Link States (Area 1)

Link ID	ADV Router	Age	Seq#	Checksum
10.1.1.1	10.0.0.11	19	0x80000001	0x00D485
10.1.2.4	10.200.200.14	622	0x80000001	0x009F20

Summary Net Link States (Area 1)

Link ID	ADV Router	Age	Seq#	Checksum
172.31.1.0	10.0.0.11	540	0x80000003	0x004108
172.31.1.0	10.0.0.12	542	0x80000003	0x003B0D
172.31.1.0	172.31.1.1	1399	0x80000003	0x00C5CA
172.31.2.0	10.0.0.11	536	0x80000001	0x00D762
172.31.2.0	10.0.0.12	537	0x80000001	0x00D167
172.31.2.0	172.31.1.1	1394	0x80000001	0x005C25

Summary ASB Link States (Area 1)

Link ID	ADV Router	Age	Seq#	Checksum
100.100.100.100	10.0.0.11	536	0x80000001	0x007213
100.100.100.100	10.0.0.12	537	0x80000001	0x006C18
100.100.100.100	172.31.1.1	1394	0x80000001	0x00F6D5

Type-5 AS External Link States

Link ID	ADV Router	Age	Seq#	Checksum	Tag
10.254.0.0	100.100.100.100	1351	0x8000010A	0x00C518	0

Interpreting the Routing Table: Types of Routes

Router Designator		Description
O	OSPF intra-area (router LSA) and network LSA	• Networks from within the area of the router • Advertised by way of router LSAs and network LSA
O IA	OSPF interarea (summary LSA)	• Networks from outside the area of the router, but within the OSPF autonomous system • Advertised by way of summary LSAs
O E1	Type 1 external routes	• Networks outside of the autonomous system of the router • Advertised by way of external LSAs
O E2	Type 2 external routes	

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0--3-16

The figure defines each of the routing table descriptors for OSPF. 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 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 area (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 (OSPF).
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 (interarea).
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 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 configured on the ASBR. The following external packet types can be configured, as shown in the figure:

- **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 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 show ip route Command

```
RouterB>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
```

The **show ip route** command example in this 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 a default of 110 for all OSPF routes, and the total cost of the route has been calculated as 50.

Configuring OSPF LSDB Overload Protection

This topic explains how to configure OSPF LSDB overload protection using the **max-lsa** command.

OSPF LSDB Overload Protection

```
Router (config-router) #  
max-lsa maximum-number [threshold-percentage] [warning-  
only] [ignore-time minutes] [ignore-count count-number]  
[reset-time minutes]
```

- **Excessive LSAs generated by other routers can drain local router resources.**
- **This feature can limit the processing of non-self-generated LSAs for a defined OSPF process.**

© 2006 Cisco Systems, Inc. All rights reserved. BSCI v3.0—3-19

If other routers are misconfigured, causing, for example, a redistribution of a large number of prefixes, 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 Software Release 12.3(7)T and later (and some specific earlier releases) by using the **max-lsa** command.

max-lsa Parameters

The table lists the parameters of the **max-lsa** command.

Parameter	Description
<i>maximum-number</i>	Maximum number of non-self-generated LSAs that the OSPF process can keep in the OSPF LSDB.
<i>threshold-percentage</i>	(Optional) The percentage of the maximum LSA number, as specified by the <i>maximum-number</i> argument, at which a warning message is logged. The default is 75 percent.
warning-only	(Optional) Specifies that only a warning message is sent when the maximum limit for LSAs is exceeded; the OSPF process never enters ignore state. Disabled by default.
ignore-time <i>minutes</i>	(Optional) Specifies the time, in minutes, to ignore all neighbors after the maximum limit of LSAs has been exceeded. The default is 5 minutes.
ignore-count <i>count-number</i>	(Optional) Specifies the number of times that the OSPF process can consecutively be placed into the ignore state. The default is five times.
reset-time <i>minutes</i>	(Optional) Specifies the time, in minutes, after which the ignore count is reset to 0. The default is 10 minutes.

When this feature is enabled, the router keeps count of the number of received (non-self-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.

Changing the Cost Metric

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

Changing the Cost Metric

- Dijkstra's algorithm determines the best path by adding all link costs along a path.
- The cost, or metric, is an indication of the overhead to send packets over an interface. Default = (100 Mbps) / (bandwidth in Mbps).
`RouterA(config-if)#`
`ip ospf cost interface-cost`
- Overrides the default cost calculation. Values from 1 to 65535 can be defined.
`RouterA(config-router)#`
`auto-cost reference-bandwidth ref-bw`
- Sets the reference bandwidth to values other than 100 Mbps (legal values range from 1 to 4,294,967 in megabits per second).

© 2006 Cisco Systems, Inc. All rights reserved. BSCI v3.0—3-20

By default, OSPF calculates the OSPF metric for an interface according to the inverse bandwidth of the interface. In general, the cost in Cisco routers is calculated using the formula (100 Mbps) / (bandwidth in Mbps). For example, a 64-kbps link gets a metric of 1562, while a T1 link gets a metric of 64. However, the cost is calculated based on a maximum bandwidth of 100 Mbps, 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** *value* interface command to accurately define the bandwidth of the interface (in kbps).

If interfaces that are faster than 100 Mbps are being used, you should use the **auto-cost reference-bandwidth** *ref-bw* command on all routers in the network to ensure accurate route calculations. The *ref-bw* is a reference bandwidth in megabits per second, and ranges from 1 to 4,294,967.

To override the default cost, manually define the cost using the **ip ospf cost** *interface-cost* 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 strongly preferred the link.

Summary

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

Summary

- **There are four OSPF router types: internal routers, backbone routers, ABRs, and ASBRs.**
- **A virtual link allows discontinuous area 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.**
- **There are 11 OSPF LSA types. The first five are the most commonly used:**
 - **Type 1 router**
 - **Type 2 network**
 - **Type 3 and 4 summary**
 - **Type 5 external**

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—3-21

Summary (Cont.)

- **In the IP routing table, OSPF routes are classified as either intra-area, interarea, or external; external routes are subdivided into E1 and E2.**
- **OSPF LSDB overload protection limits the processing of non-self-generated LSAs.**
- **The OSPF cost defaults to (100 Mbps) / (bandwidth in megabits per second). The cost can be changed on a per-interface basis, and the reference bandwidth (100 Mbps) can also be changed.**

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—3-22

OSPF Route Summarization

Overview

Scalability, 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 Open Shortest Path First Protocol (OSPF) is the ability to summarize routes at area and autonomous system (AS) boundaries.

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

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


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 the functions of interarea route summarization and external route summarization
- Describe how to configure route summarization in OSPF
- Describe the benefits of a default route in OSPF
- Describe how to configure a default route injection into OSPF

OSPF Route Summarization

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



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.

Route summarization is the consolidation of multiple routes into a single advertisement. Route summarization directly affects the amount of bandwidth, CPU, and memory resources that are consumed by the OSPF routing process.

Without route summarization, every specific-link LSA is propagated into the OSPF backbone and beyond, causing unnecessary network traffic and router overhead. Whenever an LSA is sent, all affected OSPF routers have to recompute their LSDB and the shortest path first (SPF) tree using the shortest path first (SPF) algorithm.


With route summarization, only summarized routes propagate into the backbone (area 0). This summarization is important because it prevents every router from having to rerun the 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.

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

The two types of summarization are as follows:

- **Interarea route summarization:** Interarea route summarization occurs on area border routers (ABRs) and applies to routes from within each area. It does not apply to external routes injected into OSPF via redistribution. To perform effective interarea route summarization, network numbers within areas should be assigned contiguously so that these addresses can be summarized into a minimal number of summary addresses. The figure illustrates interarea summarization at the area border routers (ABRs).
- **External route summarization:** External route summarization is specific to external routes that are injected into OSPF via route redistribution. Again, it is important to ensure the contiguity of the external address ranges that are being summarized. Summarizing overlapping ranges from two different routers could cause packets to be sent to the wrong destination. Only autonomous system boundary routers (ASBRs) generally summarize external routes.

Using Route Summarization



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

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0-3.3

Example: Using Route Summarization

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, known as variable-length subnet masking (VLSM).

Discontiguous subnets are also supported by OSPF, because subnet masks are part of the LSDB. Some protocols (such as Routing Information Protocol Version 1 [RIPv1] and Interior Gateway Routing Protocol [IGRP]) do not support VLSM or discontiguous subnets.

Therefore, for example, if a major network crosses the boundaries of an OSPF and a RIPv1 or IGRP domain, then VLSM information redistributed into RIPv1 or IGRP is lost, and static routes may have to be configured in the RIPv1 or IGRP domains.

Network numbers in areas should be assigned contiguously to ensure that these addresses can be summarized into a minimal number of summary addresses.

For example, in the figure, the list of 12 networks in the routing table of router B 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 OSPF Route Summarization

This topic describes how to configure route summarization in OSPF.

Configuring Route Summarization

```
Router(config-router)#  
area area-id range address mask [advertise | not-  
advertise] [cost cost]
```

- Consolidates interarea routes on an ABR

```
Router(config-router)#  
summary-address ip-address mask [not-advertise] [tag tag]
```

- Consolidates external routes, usually on an ASBR

© 2006 Cisco Systems, Inc. All rights reserved. BSCI v3.0—3-4

OSPF does not perform automatic summarization on major network boundaries. To configure manual interarea route summarization on an ABR, use the following procedure:

1. Configure OSPF.
2. Use the **area range** command to instruct the ABR to summarize routes for a specific area before injecting them into a different area via the backbone as type 3 summary LSAs.

Cisco IOS software creates a summary route to interface null0 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.

area range Parameters

The table describes the parameters for the **area range** command.

Parameter	Description
<i>area-id</i>	Identifies the area subject to route summarization.
<i>address</i>	Summary address that is designated for a range of addresses.
<i>mask</i>	IP subnet mask that is used for the summary route.
advertise	(Optional) Sets the address range status to advertise and generates a type 3 summary LSA.
not-advertise	(Optional) Sets the address range status to DoNotAdvertise. The type 3 summary LSA is suppressed, and the component networks remain hidden from other networks.
<i>cost</i>	(Optional) Metric or cost for this summary route, which is used during OSPF SPF calculation to determine the shortest paths to the destination. The value can be 0 to 16777215.

To configure manual route summarization on an ASBR to summarize external routes, complete the following steps:


1. Configure OSPF.
2. Use the **summary-address** command to instruct the ASBR (or the ABR) to summarize external routes before injecting them into the OSPF domain as type 5 external LSA.

summary-address Parameters

The table describes the parameters of the **summary-address** command.

Parameter	Description
<i>ip-address</i>	Summary address that is designated for a range of addresses
<i>mask</i>	IP subnet mask that is used for the summary route
not-advertise	(Optional) Used to suppress routes that match the prefix and mask pair
<i>tag</i>	(Optional) Tag value that can be used as a match value for controlling redistribution via route maps

Route Summarization Configuration Example at ABR



Example: Route Summarization Configuration at 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 router 1 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. ABR router R1 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 172.16.32.0 255.255.224.0:** Identifies area 1 as the area containing the range of networks to be summarized into area 0. ABR router R1 summarizes the range of subnets from 172.16.32.0 to 172.16.63.0 into one range: 172.16.32.0 255.255.224.0.

The configuration for router R2 works similarly.

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 an area and the backbone area, sending a summary (type 3) LSA 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.

Route Summarization Configuration Example at ASBR



© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0-3.6

Example: Route Summarization Configuration at ASBR


The figure depicts route summarization on the ASBR. On the left, an external AS running RIPv2 has its routes redistributed into OSPF.

Because of the contiguous subnet block in the external RIPv2 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 RIPv2 routes must also be redistributed into OSPF in this example; redistribution is covered in the “Manipulating Routing Updates” module.

Benefits of a Default Route in OSPF

Occasionally, you will be required to configure OSPF to advertise a default route into its AS. This topic describes the benefits of a default route in OSPF.



Example: Default Routes in OSPF

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 default routes injected into a standard area, but OSPF routers do not, by default, generate a default route into the OSPF domain. In order for OSPF to generate a default route, you must use the **default-information originate** command.

There are two ways to advertise a default route into a standard area. The first is to advertise 0.0.0.0 into the OSPF domain, provided that the advertising router already has a default route. The second is to advertise 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.

A default route shows up in the OSPF database as an external LSA type 5. 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

Configuring a Default Route in OSPF

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

Configuring OSPF Default Routes

```
Router(config-router)#  
default-information originate [always] [metric metric-  
value] [metric-type type-value] [route-map map-name]
```

- Normally, this command advertises a 0.0.0.0 default into the OSPF network only if the default route already exists in the routing table.
- The **always** keyword allows the 0.0.0.0 default to be advertised even when the default route does not exist in the routing table.

© 2006 Cisco Systems, Inc. All rights reserved. BSCI v3.0–3.8


To generate a default external route into an OSPF routing domain, use the **default-information originate** router configuration command, as shown in the figure. To disable this feature, use the **no** form of this command.

default-information originate Parameters

The table describes the parameters of the **default-information originate** command.

Parameter	Description
always	(Optional) Always advertises the default route regardless of whether the router has a default route in the routing table.
metric <i>metric-value</i>	(Optional) Metric that is used for generating the default route. If you omit a value and do not specify a value by using the default-metric router configuration command, the default metric value is 1. Note: IOS documentation indicates that the default metric value is 10; testing shows that it is 1.
metric-type <i>type-value</i>	(Optional) External link type that is associated with the default route that is advertised into the OSPF routing domain. It can be one of the following values: 1: Type 1 external route 2: Type 2 external route The default is type 2 external route (O *E2).
route-map <i>map-name</i>	(Optional) Routing process generates the default route if the route map is satisfied.

Default Route Configuration Example



```
R1#  
router ospf 100  
network 10.1.1.1 0.0.0.0 area 0  
default-information originate metric 10  
ip route 0.0.0.0 0.0.0.0 198.1.1.2
```

```
R2#  
router ospf 100  
network 10.2.1.1 0.0.0.0 area 0  
default-information originate metric 100  
ip route 0.0.0.0 0.0.0.0 198.2.1.2
```

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0-3.9

Example: Default Route Configuration

In the figure, an OSPF network is multihomed to dual Internet service providers (ISPs). Provider A is preferred, and provider B is used as a backup.

The optional **metric** parameter has been used to establish a preference for the default route to ISP A.

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 all destination networks. The benefits include a much smaller routing table and LSDB, with complete reachability.**
- **OSPF uses the default-information originate command to inject a default route.**

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—3-10

Configuring OSPF Special Area Types

Overview

Open Shortest Path First Protocol (OSPF) defines several special-case area types, including stub areas, totally stubby areas, and not-so-stubby areas (NSSAs).

The purpose behind 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 in. 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 large 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 OSPF area parameters, including stub, NSSA, totally stubby, and backbone. This ability includes being able to meet these objectives:

- Describe the OSPF area types
- Configure OSPF stub areas
- Configure OSPF totally stubby areas
- Interpret information shown on routing tables for stub areas and totally stubby areas
- Configure OSPF NSSAs
- Verify all types of OSPF stub areas

Configuring OSPF Area Types

This topic describes the OSPF area types:

- Standard area
- Backbone area (transit area)
- Stub area
- Totally stubby area
- NSSA



The characteristics assigned to an area control the type of route information that it receives. The possible area types are as follows:

- **Standard 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 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, noted as 0.0.0.0. Stub areas cannot contain autonomous system boundary routers (ASBRs) (except that the area border routers [ABRs] may also be ASBRs).
- **Totally stubby area:** This area does not accept external AS routes or summary routes from other areas internal to the AS. If the router needs to send a packet to a network external to the area, it sends the packet using a default route. Totally stubby areas cannot contain ASBRs (except that the ABRs may also be ASBRs).

- **NSSA:** NSSA is an addendum to the OSPF RFC. This area defines a special LSA type 7. An NSSA offers benefits that are similar to those of a stub or totally stubby area. However, NSSAs allow ASBRs, which is against the rule in a stub area.

Stub and Totally Stub Area Rules

An area can be stub or totally stub if:

- **There is a single ABR, or if there is more than one ABR, suboptimal routing paths to other areas or external autonomous systems are acceptable.**
- **All routers in the area are configured as stub routers.**
- **There is no ASBR in the area.**
- **The area is not area 0.**
- **No virtual links go through the area.**

© 2006 Cisco Systems, Inc. All rights reserved.


BSCI v3.0—3-9

Stub and totally stubby areas do not carry any external routes, known as type 5 LSAs. An area can be qualified as a stub or totally stubby if it has the following characteristics:

- There is 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 stub 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 OSPF primary design.

Configuring Stub Areas

This topic describes how to configure OSPF stub areas.



Configuring a stub area reduces the size of the LSDB inside an area, resulting in reduced memory requirements for routers in that area. External network LSAs (type 5), such as those 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 that 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 the spoke being the 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.

Stub Area Configuration

RouterA(config-router) #

```
area area-id stub [no-summary]
```

- This command turns on stub area networking.
- All routers in a stub area must use the stub command.

RouterA(config-router) #

```
area area-id default-cost cost
```

- This command defines the cost of a default route sent into the stub area.
- The default cost is 1.

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—3-5

To configure an area as a stub, complete the following steps:

1. Configure OSPF.
2. Define the area as stub by issuing the **area area-id stub** command to all routers within the area.

area stub Parameters

The table lists the parameters of the **area stub** command.

Parameter	Description
<i>area-id</i>	Identifier for the stub or totally stubby area. The identifier can be either a decimal value or a value in dotted-decimal format like an IP address.
no-summary	(Optional) Prevents an ABR from sending summary LSAs into the stub area. This parameter defines an area to be a totally stubby area, as discussed in the next topic.


By default, the ABR will advertise a default route with a cost of 1. Optionally, the cost of the default route can be changed by using the **area default-cost** command.

area default-cost Parameters

The following table lists the parameters of the **area default-cost** command.

Parameter	Description
<i>area-id</i>	Identifier for the stub area, totally stubby area, or NSSA. The identifier can be either a decimal value or a value in dotted-decimal format like an IP address.
<i>cost</i>	Cost for the default summary route. The acceptable values are 0 through 16777215.

OSPF Stub Area Configuration Example



Example: OSPF Stub Area Configuration

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 3 (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 routes that appear in the routing table of router R4 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)

Note The hello packet exchanged between OSPF routers contains a stub area flag that must match on neighboring routers. The **area area-id 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.

Configuring Totally Stubby Areas

This topic describes how to configure OSPF totally stubby areas.



The totally stubby area technique is 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.

By blocking these routes, the 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 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.

Totally Stubby Configuration

```
RouterA(config-router)#
```

```
area area-id stub no-summary
```

- The addition of **no-summary** on the **ABR** creates a totally stubby area and prevents all summary LSAs from entering the stub area.


© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—3-8

To configure an area as totally stubby, complete the following steps:

1. Configure OSPF.
2. Define the area as a type of stub area by issuing the **area area-id stub** command to all routers within the area.
3. At the ABR only, add the **no-summary** keyword to the **area area-id stub** command. This command is described in the previous topic.

Totally Stubby Configuration Example



Example: Totally Stubby Configuration

The figure shows an example of a totally stubby area configuration. All routes 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 2 and to 10 on router 4.

Both default routes are advertised into area 1. However, the default route from router R2 is advertised with a lower cost to make it preferable if the internal cost from router R3 to router R4 is the same as the internal cost from router R3 to router R2.

Notice that router R3 requires the **area 1 stub** command, yet the **no-summary** extension is not required. Only ABRs use **no-summary** to keep summary LSAs from being propagated into another area.

Note Remember that all routers in a stub or totally stubby area must be configured as stub. An OSPF adjacency will not form between stub and nonstub routers.

Interpreting Routing Tables

This topic interprets information shown on routing tables for stub areas and totally stubby areas.

Routing Table in a Standard Area

```
P1R3#sh ip route
<output omitted>

Gateway of last resort is not set
 172.31.0.0/32 is subnetted, 4 subnets
O IA   172.31.22.4 [110/782] via 10.1.1.1, 00:02:44, FastEthernet0/0
O IA   172.31.11.1 [110/1] via 10.1.1.1, 00:02:44, FastEthernet0/0
O IA   172.31.11.2 [110/782] via 10.1.3.4, 00:02:52, Serial0/0/0
       [110/782] via 10.1.1.1, 00:02:52, FastEthernet0/0
O IA   172.31.11.4 [110/782] via 10.1.1.1, 00:02:44, FastEthernet0/0
 10.0.0.0/8 is variably subnetted, 7 subnets, 2 masks
O      10.11.0.0/24 [110/782] via 10.1.1.1, 00:03:22, 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, 00:03:23, 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, 00:03:23, FastEthernet0/0
O E2   10.254.0.0/24 [110/50] via 10.1.1.1, 00:02:39, FastEthernet0/0
P1R3#
```

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—3-10

Example: Routing Table in a Standard Area

This figure shows how the routing table of an OSPF router in a standard area without stub or totally stubby configuration might look. Intra-area, interarea, and external routes are all maintained in a standard area.

Routing Table in a Stub Area

```
P1R3#sh 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
P1R3#
```

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—3-11

Example: Routing Table in a Stub Area

This 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. External routes are not visible in the routing table but are accessible via the intra-area default route.

Routing Table in a Stub Area with Summarization

```
P1R3#sh 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
P1R3#
```

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0--3-12

Example: Routing Table in a Stub Area with Summarization

This figure shows how the same routing table looks if summarization is performed on the ABR; the area is still configured as a stub area. Intra-area and summarized interarea routes are all maintained. External routes are not visible in the routing table but are accessible via the intra-area default route.

Routing Table in a Totally Stubby Area

```
P1R3#sh 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
P1R3#
```

© 2006 Cisco Systems, Inc. All rights reserved.


BSCI v3.0—3-13

Example: Routing Table in a Totally Stubby Area

This 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. Intra-area routes are maintained. Interarea and external routes are not visible in the routing table but are accessible via the intra-area default route.

Configuring NSSAs

This topic describes how to configure OSPF not-so-stubby areas (NSSAs).



The OSPF NSSA feature is described by RFC 3101 and was first introduced in Cisco IOS Software 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 type 7, which can exist only in an NSSA. An NSSA ASBR generates this LSA, and an NSSA ABR translates it into a type 5 LSA, which gets propagated into the OSPF domain. The NSSA retains the other stub area features; the ABR sends a default route into the NSSA instead of external routes from other ASBRs.

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; N2 means that the metric is calculated like external type 2. The default is O N2.

NSSA Configuration

RouterA(config-router)#

```
area area-id nssa [no-redistribution] [default-  
information-originate [metric metric-value] [metric-  
type type-value]] [no-summary]
```

- Use this command instead of the **area stub** command to define the area as NSSA.
- The **no-summary** keyword creates an NSSA totally stubby area; this is a Cisco proprietary feature.

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—3-16

To configure an area as an NSSA, complete the following steps:


1. Configure OSPF.
2. Define the area as NSSA by issuing the **area area-id nssa** command to all routers within the area. The **area area-id nssa** command is used in place of the **area area-id stub** command. Remember that all routers in the NSSA must have this command configured; routers will not form an adjacency unless both are configured as NSSA.

area nssa Parameters

The table lists the parameters of the **area nssa** command.

Parameter	Description
<i>area-id</i>	Identifier for the NSSA. The identifier can be either a decimal value or a value in dotted-decimal format like an IP address.
no-redistribution	(Optional) Used when the router is an NSSA ABR and you want the redistribute command to import routes only into the standard areas, but not into the NSSA area.
default-information-originate	(Optional) Used to generate a type 7 default into the NSSA area. This keyword takes effect only on an NSSA ABR or an NSSA ASBR.
metric <i>metric-value</i>	(Optional) Metric that is used for generating the default route. Acceptable values are 0 through 16777214.
metric-type <i>type-value</i>	(Optional) OSPF metric type for default routes. It can be one of the following values: 1: Type 1 external route 2: Type 2 external route
no-summary	(Optional) Allows an area to be an NSSA but not have summary routes injected into it.

Example: NSSA Configuration




Example: NSSA Configuration

In the figure, router R1 is the ASBR that is redistributing Routing Information Protocol (RIP) routes into area 1, the NSSA. Router R2 is the NSSA ABR. This router converts LSA type 7 into type 5 for advertisement into the backbone area 0. Router R2 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 R2 (the NSSA ABR) to generate an O *N2 default route (O *N2 0.0.0.0/0) into the NSSA, use the **default-information-originate** option on the **area area-id nssa** command on router R2.

NSSA Totally Stubby Configuration



- NSSA totally stubby area is a Cisco proprietary feature.

Example: NSSA Totally Stubby Configuration

In the figure, the ABR (R2) is using the **area 1 nssa no-summary** command. This command works exactly 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 R2, 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 require the **area 1 nssa** command only. The NSSA totally stubby configuration is a Cisco proprietary feature like the totally stubby area feature.

Verifying All Stub Area Types

This topic describes how to verify all types of OSPF stub areas.

show Commands for Stub and NSSA

```
RouterA#  
show ip ospf
```

- Displays which areas are normal, stub, or NSSA

```
RouterA#  
show ip ospf database
```

- Displays details of LSAs

```
RouterA#  
show ip ospf database nssa-external
```

- Displays specific details of each LSA type 7 update in database

```
RouterA#  
show ip route
```

- Displays all routes

© 2006 Cisco Systems, Inc. All rights reserved. BSCI v3.0—3-18

The **show** commands in the figure are used to display information about all of 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: standard, backbone, stub, totally stubby, and 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 `show ip ospf`, `show ip ospf database`, `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.

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—3-19

Configuring OSPF Authentication

Overview

You can prevent your router from receiving fraudulent route updates by configuring neighbor router authentication. Open Shortest Path First Protocol (OSPF) neighbor authentication (also called neighbor router authentication or route authentication) can be configured such that routers can participate in routing based on predefined passwords.

This lesson describes the two types of OSPF authentication and how to configure and troubleshoot 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:

- Describe the two types of authentication used in OSPF
- Configure simple password authentication
- Configure MD5 authentication
- Troubleshoot simple password authentication
- Troubleshoot MD5 authentication

Types of Authentication

This topic describes two types of authentication used in OSPF:

- Simple password or plain-text authentication
- Message Digest 5 (MD5) authentication

OSPF Authentication Types

- **OSPF supports 2 types of authentication:**
 - Simple password (or plain text) authentication
 - MD5 authentication
- **Router generates and checks every OSPF packet. Router authenticates the source of each routing update packet that it receives.**
- **Configure a “key” (password); each participating neighbor must have same key configured.**

© 2006 Cisco Systems, Inc. All rights reserved. BSCI v3.0—3-2

OSPF neighbor authentication (also called neighbor router authentication or route authentication) can be configured such that routers can participate in routing based on predefined passwords.

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

By default, OSPF uses null authentication, which means that routing exchanges over a network are not authenticated. OSPF supports two other authentication methods: simple password authentication (also called plain-text authentication), and MD5 authentication.

OSPF MD5 authentication includes a nondecreasing sequence number in each OSPF packet to protect against replay attacks.

Configuring Simple Password Authentication

This topic describes how to configure simple password authentication.

Configuring OSPF Simple Password Authentication

```
Router(config-if)#
```

```
ip ospf authentication-key password
```

- Assigns a password to be used with neighboring routers

```
Router(config-if)#
```

```
ip ospf authentication [message-digest | null]
```

- Specifies the authentication type for an interface (since Cisco IOS software 12.0)

```
Router(config-router)#
```

```
area area-id authentication [message-digest]
```

- Specifies the authentication type for an area (was in Cisco IOS software before 12.0)

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—3.3

To configure OSPF simple password authentication, complete the following steps:

1. Assign a password to be used with neighboring routers that are using the OSPF simple password authentication using the **ip ospf authentication-key** command, as shown in the figure.

ip ospf authentication-key Parameter

The table describes the parameter of the **ip ospf authentication-key** command.

Parameter	Description
<i>password</i>	Any continuous string of characters that can be entered from the keyboard, up to 8 bytes in length

Note In Cisco IOS Software Release 12.4, the router will give a warning message if you try to configure a password longer than eight characters; only the first eight characters will be used. Some earlier Cisco IOS releases did not provide this warning.

The password 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 plain text 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 <i>encryption-type</i> of 7 specified before the encrypted key.

- Specify the authentication type using the **ip ospf authentication** command as shown in the figure.

ip ospf authentication Parameters

The table describes the parameters of the **ip ospf authentication** command.

Parameter	Description
message-digest	(Optional) Specifies that MD5 authentication will be used.
null	(Optional) No authentication is used. Useful for overriding password or MD5 authentication if configured for an area.

For simple password authentication, use the **ip ospf authentication** command with no parameters. Before using this command, configure a password for the interface using the **ip ospf authentication-key** command.


The **ip ospf authentication** command was introduced in Cisco IOS Software Release 12.0. For backward compatibility, 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 area-id authentication [message-digest]** router configuration command.

area authentication Parameters

The parameters for the **area authentication** command are described in the table.

Parameter	Description
<i>area-id</i>	Identifier of the area for which authentication is to be enabled. The identifier can be specified as either a decimal value or an IP address.
message-digest	(Optional) Enables MD5 authentication on the area specified by the <i>area-id</i> argument.

Example Simple Password Authentication Configuration



Example: Simple Password Authentication Configuration

The figure shows the network used to illustrate the configuration, verification, and troubleshooting of simple password authentication. The configuration of the R1 router is also shown in this figure.

Simple password authentication is configured on interface serial 0/0/1 with the **ip ospf authentication** command. The interface is configured with an authentication key of “plainpas.”

R2 Configuration for Simple Password Authentication

```
<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 authentication
 ip ospf authentication-key plainpas

<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
```

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0-3.5

Example: R2 Configuration for Simple Password Authentication

The configuration of the R2 router is shown in this figure.

Simple password authentication is configured on interface serial 0/0/1 with the **ip ospf authentication** command. The interface is configured with an authentication key of “plainpas.” Notice that the connecting interfaces on both R1 and R2 are configured for the same type of authentication with the same authentication key.

Verifying Simple Password Authentication

```
R1#sh ip ospf neighbor
Neighbor ID      Pri   State           Dead Time   Address        Interface
10.2.2.2         0     FULL/ -         00:00:32    192.168.1.102  Serial0/0/1

R1#show ip route
<output omitted>
Gateway of last resort is not set
  10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
O       10.2.2.2/32 [110/782] via 192.168.1.102, 00:01:17, Serial0/0/1
C       10.1.1.0/24 is directly connected, Loopback0
       192.168.1.0/27 is subnetted, 1 subnets
C       192.168.1.96 is directly connected, Serial0/0/1

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
```

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—3-6

Verifying Simple Password Authentication

The figure shows the output of the **show ip ospf neighbor** and **show ip route** commands.

Notice that the neighbor state is FULL, indicating that the two routers have successfully formed an OSPF adjacency. The routing table verifies that the 10.2.2.2 address has been learned via OSPF over the serial connection.

The results of a ping to the R2 loopback interface address are also displayed to illustrate that the link is working.

Configuring MD5 Authentication

This topic describes how to configure MD5 authentication.

Configuring OSPF MD5 Authentication

```
Router(config-if)#  
ip ospf message-digest-key key-id md5 key
```

- Assigns a key ID and key to be used with neighboring routers

```
Router(config-if)#  
ip ospf authentication [message-digest | null]
```

- Specifies the authentication type for an interface (since Cisco IOS software 12.0)

```
Router(config-router)#  
area area-id authentication [message-digest]
```

- Specifies the authentication type for an area (was in Cisco IOS software before 12.0)

© 2006 Cisco Systems, Inc. All rights reserved. BSCI v3.0-3.7

With OSPF MD5 authentication, a key and key ID are configured on each router.

To configure OSPF MD5 authentication, complete the following steps:

1. Assign a key ID and key to be used with neighboring routers that are using the OSPF MD5 authentication, using the **ip ospf message-digest-key** command, as shown in the figure.

ip ospf message-digest-key Parameters

The table describes the parameters in the **ip ospf message-digest-key** command.

Parameter	Description
<i>key-id</i>	An identifier in the range from 1 to 255
<i>key</i>	Alphanumeric password of up to 16 bytes

Note In Cisco IOS Software Release 12.4, the router will give a warning message if you try to configure a password longer than 16 characters; only the first 16 characters will be used. Some earlier Cisco IOS releases did not provide this warning.

The key and the key ID 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 of its neighbors have adopted the new key.

For the purpose of illustrating the process of changing keys, suppose the current configuration is as follows:

```
interface FastEthernet 0/0
  ip ospf message-digest-key 100 md5 OLD
```

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 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 is used for authentication on FastEthernet interface 0/0.


It is recommended that you not keep more than one key per interface. Every time 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.

Note	If the service password-encryption command is not used when implementing OSPF authentication, the key will be stored as plain text 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 <i>encryption-type</i> of 7 specified before the encrypted key.

2. Specify the authentication type using the **ip ospf authentication** command as shown in the figure. The parameters for this command are as described in the previous topic. For MD5 authentication, use 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 Software Release 12.0. As it is for simple password authentication, the MD5 authentication type for an area is still supported using the **area area-id authentication message-digest** router configuration command, for backward compatibility.

Example MD5 Authentication Configuration



```
<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 secretpass

<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
```

318P_377

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—3-8

Example: MD5 Authentication Configuration

The figure shows the network used to illustrate the configuration, verification, and troubleshooting of MD5 authentication. The configuration of the R1 router is also shown in this figure.

MD5 authentication is configured on interface serial 0/0/1 with the **ip ospf authentication message-digest** command. The interface is configured with an authentication key number 1 set to “secretpass.”

R2 Configuration for MD5 Authentication

```
<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 authentication message-digest
 ip ospf message-digest-key 1 md5 secretpass

<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
```

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0–3.9

Example: R2 Configuration for MD5 Authentication

The configuration of the R2 router is shown in this figure.

MD5 authentication is configured on interface serial 0/0/1 with the **ip ospf authentication message-digest** command. The interface is configured with an authentication key number 1 set to “secretpass.” Notice that the connecting interfaces on both R1 and R2 are configured for the same type of authentication with the same authentication key and key ID.

Verifying MD5 Authentication

```
R1#show ip ospf neighbor
Neighbor ID      Pri   State           Dead Time   Address        Interface
10.2.2.2          0    FULL/  -         00:00:31    192.168.1.102  Serial0/0/1

R1#show ip route
<output omitted>
Gateway of last resort is not set
  10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
O       10.2.2.2/32 [110/782] via 192.168.1.102, 00:00:37, Serial0/0/1
C       10.1.1.0/24 is directly connected, Loopback0
        192.168.1.0/27 is subnetted, 1 subnets
C       192.168.1.96 is directly connected, Serial0/0/1

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
```

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0--3-10

Verifying MD5 Authentication

The figure shows the output of the **show ip ospf neighbor** and **show ip route** commands.

Notice that the neighbor state is FULL, indicating that the two routers have successfully formed an OSPF adjacency. The routing table verifies that the 10.2.2.2 address has been learned via OSPF over the serial connection.

The results of a ping to the R2 loopback interface address is also displayed to illustrate that the link is working.

Troubleshooting Simple Password Authentication

This topic describes how to troubleshoot simple password authentication.

Troubleshooting Simple Password Authentication

```
Router#  
debug ip ospf adj
```

- Displays the OSPF adjacency-related events

```
R1#debug ip ospf adj  
OSPF adjacency events debugging is on  
R1#  
<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
```

© 2006 Cisco Systems, Inc. All rights reserved. BSCI v3.0—3-11

The **debug ip ospf adj** command is used to display OSPF adjacency-related events and is very useful when troubleshooting authentication.

Example: Successful Simple Password Authentication

The following output of the **debug ip ospf adj** command (part of which is shown in the figure) illustrates successful simple password authentication on R1 after the serial 0/0/1 interface, on which authentication has been configured, comes up:

```
*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 shown in the figure illustrates that R1 has successfully formed an adjacency with R2.

Troubleshooting Simple Password Authentication Problems

Simple authentication on R1, no authentication 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
```

Simple authentication on R1 and R2, but different passwords

```
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
```

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—3-12

Example: Troubleshooting Simple Password Authentication Problems

If simple password authentication is configured on the R1 serial 0/0/1 interface but no authentication is configured on the R2 serial 0/0/1 interface, the routers will not be able to form an adjacency over that link. The output of the **debug ip ospf adj** command shown in the upper portion of the figure illustrates 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, MD5—type 2.

If simple password authentication is configured on the R1 serial 0/0/1 interface and on the R2 serial 0/0/1 interface, 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 shown in the lower portion of the figure illustrates that the routers report a mismatch in authentication key; no OSPF packets will be sent between the neighbors.

Troubleshooting MD5 Authentication

This topic describes how to troubleshoot MD5 authentication.

Troubleshooting MD5 Authentication

```
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
```

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0-3.13

Example: Successful MD5 Authentication

The following output of the **debug ip ospf adj** command (part of which is shown in the figure) illustrates successful MD5 authentication on R1 after the serial 0/0/1 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 shown in the figure illustrates that R1 has successfully formed an adjacency with R2.

Troubleshooting MD5 Authentication Problems

MD5 authentication on both R1 and R2, but 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
```

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0--3-14

Example: Troubleshooting MD5 Authentication Problems

If MD5 authentication is configured on the R1 serial 0/0/1 interface and on the R2 serial 0/0/1 interface, 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 shown in the figure illustrates that the routers report a mismatch in authentication key. 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 plain text) 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.**

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—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 `show ip ospf neighbor`, `show ip route`, and `debug ip ospf adj` to verify and troubleshoot both types of authentication.**
- **With MD5 authentication, the `debug ip ospf adj` command output indicates the key ID sent.**

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—3-16

Module Summary

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

Module Summary

- **OSPF is an open-standard link-state routing protocol, offering quick convergence and the ability to scale large networks.**
- **There are five OSPF packet types: hello, DBD, LSU, LSR, and LSAck.**
- **Configuration of OSPF is a two-step process:**
 - **Enter OSPF configuration with the router ospf command.**
 - **Use the network command to describe which interfaces will run OSPF in which area.**
- **OSPF defines three types of networks: point-to-point, broadcast, and NBMA. On NBMA networks, OSPF mode options include nonbroadcast, broadcast, point-to-multipoint, point-to-multipoint nonbroadcast, and point-to-point.**

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—3-1

Module Summary (Cont.)

- **LSAs are the building blocks of the LSDB. There are 11 types of OSPF LSAs.**
- **Route summarization reduces OSPF LSA flooding and routing table size, which reduces memory and CPU utilization on routers.**
- **Stub area techniques improve OSPF performance by reducing the LSA flooding.**
- **OSPF supports two types of authentication:**
 - **Simple password (or plain text) authentication**
 - **MD5 authentication**

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—3-2

Open Shortest Path First Protocol (OSPF) is one of the most commonly used interior gateway protocols in IP networking. OSPF is a fairly complex, open-standard protocol 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) Which table is NOT maintained by a link-state routing protocol? (Source: Introducing the OSPF Protocol)
- A) routing
 - B) topology
 - C) update
 - D) neighbor
- Q2) The memory needed to maintain tables is one disadvantage of link-state protocols. (Source: Introducing the OSPF Protocol)
- A) true
 - B) false
- Q3) Match each table to its function. (Source: Introducing the OSPF 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: Introducing the OSPF 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: Introducing the OSPF Protocol)
- A) 50
 - B) 10
 - C) 200
 - D) 500
- Q6) Which OSPF packet helps form the neighbor adjacency? (Source: Introducing the OSPF 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: Introducing the OSPF 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 the SPF calculations? (Source: Introducing the OSPF 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: Introducing the OSPF Protocol)
- A) one
 - B) two
 - C) four
 - D) eight
- Q10) An area border router maintains _____. (Source: Introducing the OSPF 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 others
 - D) a separate routing table for each area
- Q11) In a multiarea network, any area can be the backbone area, although this is most often area 0. (Source: Introducing the OSPF Protocol)
- A) true
 - B) false
- Q12) When an LSA is received by an OSPF router, it is installed in the _____. (Source: Introducing the OSPF Protocol)
- A) neighbor table
 - B) topology table
 - C) routing table
- Q13) An OSPF router receives an LSA, the router checks its sequence number, and this number matches the sequence number of the LSA that the receiving router already has. What does the receiving router do with the LSA? (Source: Introducing the OSPF Protocol)
- A) ignores the LSA
 - B) adds it to the database
 - C) sends newer LSU update to source router
 - D) floods the LSA to the other routers

- Q14) An OSPF router receives an LSA. The router checks its sequence number and finds that this number is *higher* than the sequence number it already has. Which two tasks does the router perform with the LSA? (Choose two.) (Source: Introducing the OSPF Protocol)
- A) ignores the LSA
 - B) adds it to the database
 - C) sends newer LSU update to source router
 - D) floods the LSA to the other routers
- Q15) An OSPF router receives an LSA. The router checks its sequence number and finds that this number is *lower* than the sequence number it already has. What does the router do with the LSA? (Source: Introducing the OSPF Protocol)
- A) ignores the LSA
 - B) adds it to the database
 - C) sends newer LSU update to 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: Introducing the OSPF Protocol)
- A) 30 seconds
 - B) 1 minute
 - C) 30 minutes
 - D) 1 hour
- Q17) Distance vector protocols use the concept of split horizon, but link-state routing protocols, such as OSPF, do not. (Source: Introducing the OSPF Protocol)
- A) true
 - B) false
- Q18) The outcome of Dijkstra's calculation is used to populate the _____. (Source: Introducing the OSPF Protocol)
- A) topology table
 - B) routing table
 - C) neighbor table
 - D) adjacency table
- Q19) What is the IP protocol number for OSPF packets? (Source: OSPF Packet Types)
- A) 89
 - B) 86
 - C) 20
 - D) 76
- Q20) Which packet is NOT an OSPF packet type? (Source: OSPF Packet Types)
- A) LSU
 - B) LSR
 - C) DBD
 - D) LSAck
 - E) hello
 - F) query

- Q21) Which multicast address does the OSPF Hello protocol use? (Source: OSPF Packet Types)
- 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 updates to ensure that a neighbor relationship is maintained between adjacent routers. (Source: OSPF Packet Types)
- A) true
 - B) false
- Q23) Place the exchange protocol states in the correct order. (Source: OSPF Packet Types)
- A) _____ two-way
 - B) _____ loading
 - C) _____ down
 - D) _____ full
 - E) _____ exchange
 - F) _____ init
 - G) _____ exstart
- Q24) DBD packets are involved during which two states? (Choose two.) (Source: OSPF Packet Types)
- A) exstart
 - B) loading
 - C) exchange
 - D) two-way
- Q25) At which interval does OSPF refresh LSAs? (Source: OSPF Packet Types)
- A) 10 seconds
 - B) 30 seconds
 - C) 30 minutes
 - D) 1 hour
- Q26) Which field is NOT a field within an OSPF packet header? (Source: OSPF Packet Types)
- A) packet length
 - B) router ID
 - C) authentication type
 - D) maxage time
- Q27) Which two commands are required for basic OSPF configuration? (Choose two.) (Source: Configuring 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**

- Q28) Which OSPF **show** command describes a list of OSPF adjacencies? (Source: Configuring OSPF Routing)
- A) **show ip ospf interface**
 - B) **show ip ospf**
 - C) **show ip route**
 - D) **show ip ospf neighbor**
- Q29) Which technique is NOT used for router ID selection? (Source: Configuring 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
- Q30) When you are using the **router-id** command, the router ID immediately changes to the IP address that has been entered. (Source: Configuring OSPF Routing)
- A) true
 - B) false
- Q31) Which network statement is used to configure OSPF on an interface with IP address 172.16.1.1 in area 0? (Source: Configuring 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**
- Q32) Only one OSPF process can run on a Cisco router at one time. (Source: Configuring OSPF Routing)
- A) true
 - B) false
- Q33) Which mode is the **ip ospf process-id area area-id** command entered in? (Source: Configuring OSPF Routing)
- A) **(config)#**
 - B) **(config-if)#**
 - C) **(config-router)#**
 - D) **(config-if)# or (config-router)#**
- Q34) The OSPF **router-id** command should be used in global configuration mode. (Source: Configuring OSPF Routing)
- A) true
 - B) false

- Q35) A router has a FastEthernet interface with IP address 172.16.45.1, a loopback 0 interface with IP address 10.3.3.3, a loopback 1 interface with 10.2.2.2, and a **router-id** command with IP address 10.1.1.1. Which router ID will be selected? (Source: Configuring OSPF Routing)
- A) 172.16.45.1
 - B) 10.3.3.3
 - C) 10.2.2.2
 - D) 10.1.1.1
- Q36) The **show ip ospf neighbor** command shows a FULL state on one of the two neighbors in its table. Which neighbor or neighbors successfully exchange LSDB information? (Source: Configuring OSPF Routing)
- A) neighbor in FULL state
 - B) neighbor not in FULL state
 - C) both have exchanged databases
 - D) neither has exchanged databases
- Q37) Which two **show** commands can be used to verify the OSPF router ID of a router? (Choose two.) (Source: Configuring OSPF Routing)
- A) **show ip ospf interface**
 - B) **show ip ospf neighbor**
 - C) **show ip ospf**
 - D) **show ip route**
- Q38) 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 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
- Q39) Which statement describes the process ID on the **router ospf** command? (Source: Configuring 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 and does not need to match that on other routers.
 - C) The OSPF process ID is similar to an AS number.
 - D) There can be only one OSPF process ID in a router configuration.
- Q40) OSPF does not require a Hello protocol on point-to-point links because the adjacent router is directly connected. (Source: OSPF Network Types)
- A) true
 - B) false

- Q41) 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: OSPF Network Types)
- 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.
- Q42) When the DR fails, the BDR builds new adjacencies, exchanges databases, and takes over as DR automatically. (Source: OSPF Network Types)
- A) true
 - B) false
- Q43) What is the default hello interval for NBMA interfaces? (Source: OSPF Network Types)
- A) 10 seconds
 - B) 30 seconds
 - C) 120 seconds
 - D) 60 seconds
- Q44) An OSPF router automatically builds adjacencies with neighboring routers on an NBMA link. (Source: OSPF Network Types)
- A) true
 - B) false
- Q45) Which mode of OSPF operation is RFC-compliant? (Source: OSPF Network Types)
- A) point-to-multipoint nonbroadcast
 - B) point-to-multipoint
 - C) broadcast
 - D) point-to-point
- Q46) Match the OSPF over Frame Relay mode of operation with its description. (Source: OSPF Network Types)
- A) broadcast
 - B) point-to-multipoint
 - C) nonbroadcast
- _____ 1. does not discover neighbors automatically
- _____ 2. discovers neighbors automatically and requires DR and BDR election
- _____ 3. used in partial-mesh topologies, does not require DR and BDR election, automatically discovers neighbors
- Q47) Which two OSPF over Frame Relay modes elect a DR? (Choose two.) (Source: OSPF Network Types)
- A) broadcast
 - B) nonbroadcast
 - C) point-to-multipoint
 - D) point-to-point

- Q48) A point-to-point subinterface solves which two problems with OSPF over Frame Relay? (Choose two.) (Source: OSPF Network Types)
- A) works with multiple vendors
 - B) manual configuration of neighbors not required
 - C) DR and BDR not required
 - D) saves on subnets
- Q49) When troubleshooting a DR election problem, which is an excellent command to use? (Source: OSPF Network Types)
- A) **show ip ospf**
 - B) **show ip route**
 - C) **debug ip ospf neighbor**
 - D) **debug ip ospf adj**
- Q50) Which destination IP address does OSPF use when advertising to all SPF routers? (Source: OSPF Network Types)
- A) 224.0.0.6
 - B) 224.0.0.5
 - C) 255.255.255.255
 - D) IP address of output interface
- Q51) What are the three types of networks defined by OSPF? (Choose three.) (Source: OSPF Network Types)
- A) point-to-point
 - B) broadcast
 - C) point-to-multipoint
 - D) point-to-multipoint nonbroadcast
 - E) nonbroadcast multiaccess
- Q52) With a hello interval of 10 seconds, what does the dead interval default to? (Source: OSPF Network Types)
- A) 10 seconds
 - B) 20 seconds
 - C) 40 seconds
 - D) 60 seconds
- Q53) The BDR, like the DR, maintains a full set of adjacencies on a broadcast link. (Source: OSPF Network Types)
- A) true
 - B) false
- Q54) Which two protocols use the OSPF nonbroadcast mode by default? (Choose two.) (Source: OSPF Network Types)
- A) PPP
 - B) HDLC
 - C) X.25
 - D) ATM
 - E) SLIP

- Q55) Which two modes require a DR? (Choose two.) (Source: OSPF Network Types)
- A) point-to-point
 - B) broadcast
 - C) point-to-multipoint
 - D) nonbroadcast
- Q56) Which two statements regarding the OSPF nonbroadcast mode are correct? (Choose two.) (Source: OSPF Network Types)
- A) requires manual **neighbor** commands
 - B) does not use a DR and BDR
 - C) uses a DR and BDR
 - D) requires multiple subnets
- Q57) When you are using an OSPF **neighbor** command, you must configure it under an interface. (Source: OSPF Network Types)
- A) true
 - B) false
- Q58) Which OSPF mode requires **neighbor** commands? (Source: OSPF Network Types)
- A) broadcast
 - B) point-to-point
 - C) point-to-multipoint
 - D) point-to-multipoint nonbroadcast
- Q59) You have a partially meshed hub-and-spoke Frame Relay network. You consider using the **ip ospf network broadcast** command on the Frame Relay interface because you do not want to configure neighbor router statements. Is this a good idea? (Source: OSPF Network Types)
- A) yes
 - B) no
- Q60) Which three benefits are derived from a multiarea design in OSPF? (Choose three.) (Source: Link-State Advertisements)
- A) reduced LSA flooding
 - B) reduced SPF calculations
 - C) reduced size of the neighbor table
 - D) reduced size of the routing table
- Q61) Which router is not an OSPF router type? (Source: Link-State Advertisements)
- A) backbone
 - B) ABR
 - C) ASBR
 - D) core

- Q62) List the four link types that a type 1 LSA defines. (Source: Link-State Advertisements)
- A) _____
 - B) _____
 - C) _____
 - D) _____
- Q63) Match each LSA name with the number that corresponds to its LSA type. (Source: Link-State Advertisements)
- A) external
 - B) network
 - C) summary
 - D) multicast
 - E) router
 - F) opaque
 - G) NSSA
- _____ 5
- _____ 2
- _____ 3 and 4
- _____ 6
- _____ 1
- _____ 9–11
- _____ 7
- Q64) Which of the following is NOT described in the output of the **show ip ospf database** command? (Source: Link-State Advertisements)
- A) advertising router
 - B) maximum age counter
 - C) link count
 - D) LSA type
 - E) link type
- Q65) If the OSPF routing table shows an O E1 route, what does this mean? (Source: Link-State Advertisements)
- 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.

- Q66) Which two LSAs describe intra-area routing information? (Choose two.) (Source: Link-State Advertisements)
- A) summary
 - B) external 1
 - C) external 2
 - D) router
 - E) network
- Q67) Where will a type 3 LSA be sent? (Source: Link-State Advertisements)
- A) only within the area it originates from
 - B) within the area it originates from plus the backbone area
 - C) within the area it originates from plus all other areas
 - D) within the backbone area plus all other areas
- Q68) 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: Link-State Advertisements)
- A) true
 - B) false
- Q69) 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: Link-State Advertisements)
- 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
- Q70) Looking at the routing table, you notice “[110/55].” What does this mean? (Source: Link-State Advertisements)
- 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.
- Q71) What does it mean if a route in the routing table has an indicator of O? (Source: Link-State Advertisements)
- A) It is intra-area.
 - B) It is interarea.
 - C) It is external.
 - D) It is stub.
- Q72) What is the difference between an LSA 3 and an LSA 4? (Source: Link-State Advertisements)
- 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.

- Q73) By default, OSPF assigns a cost of 1 to a bandwidth of _____. (Source: Link-State Advertisements)
- A) T1
 - B) 1 Gbps (gigabits per second)
 - C) 100 Mbps (megabits per second)
 - D) 10 Gbps
- Q74) The OSPF LSDB shows an LSA with an age of 1799. What does this mean? (Source: Link-State Advertisements)
- 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.
- Q75) Summary LSAs are not automatically summarized. (Source: Link-State Advertisements)
- A) true
 - B) false
- Q76) What are the two reasons why route summarization is important? (Choose two.) (Source: OSPF Route Summarization)
- A) reduces LSA type 1 flooding
 - B) reduces LSA type 3 flooding
 - C) reduces the size of the routing table
 - D) reduces the size of the neighbor table
- Q77) Which two features play a key role in route summarization? (Choose two.) (Source: OSPF Route Summarization)
- A) network numbers in areas should be assigned contiguously
 - B) network numbers in areas should be assigned discontinuously
 - C) FLSM
 - D) VLSM
- Q78) Which command would you use to summarize routes into area 0 from the ABR? (Source: OSPF Route Summarization)
- A) **summary-address**
 - B) **area x range**
 - C) **network**
 - D) **area x summary**
- Q79) Which command would you use to summarize routes into OSPF from the ASBR? (Source: OSPF Route Summarization)
- A) **summary-address**
 - B) **area x range**
 - C) **network**
 - D) **area x summary**

- Q80) A default route is identified in the OSPF database as an _____. (Source: OSPF Route Summarization)
- A) LSA type 1
 - B) LSA type 2
 - C) LSA type 3
 - D) LSA type 4
 - E) LSA type 5
- Q81) The primary purpose of a default route is to reduce the routing table and LSDB size. A default route avoids detailed updating of routes by inserting a single 0.0.0.0 into the routing table, making this 0.0.0.0 route act as a gateway of last resort. (Source: OSPF Route Summarization)
- A) true
 - B) false
- Q82) When should you use the **always** keyword with the **default-information originate** command? (Source: 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
- Q83) Default routes must always be O E2 routes; there is no other choice. (Source: OSPF Route Summarization)
- A) true
 - B) false
- Q84) A summary LSA (type 3 LSA) is designed to automatically summarize a network into blocks. (Source: OSPF Route Summarization)
- A) true
 - B) false
- Q85) Route summarization reduces the flooding of which two of the following LSA types? (Choose two.) (Source: OSPF Route Summarization)
- A) router
 - B) network
 - C) summary
 - D) external
 - E) NSSA
- Q86) 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: OSPF Route Summarization)
-

- Q87) 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: OSPF Route Summarization)
-
- Q88) It is important to always summarize the routes from area 0 into other areas. Suboptimal path selection can occur if you do not. (Source: OSPF Route Summarization)
- A) true
 - B) false
- Q89) The **area range** command has an optional **not-advertise** parameter, which is used to prevent advertising _____. (Source: 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
- Q90) Generally, a default route is described in the routing table as an _____. (Source: OSPF Route Summarization)
- A) O route
 - B) O IA route
 - C) O *E1 route
 - D) O *E2 route
- Q91) 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: 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**
- Q92) The **area x range** and **network** commands are similar because both use inverse masks for configuration purposes. (Source: OSPF Route Summarization)
- A) true
 - B) false
- Q93) A default route is a form of route summarization. (Source: OSPF Route Summarization)
- A) true
 - B) false
- Q94) Which is NOT permitted in a stub area? (Source: Configuring OSPF Special Area Types)
- A) an ABR
 - B) an ASBR
 - C) summary routes
 - D) summary LSAs

- Q95) Which type of router advertises the default into a stub area? (Source: Configuring OSPF Special Area Types)
- A) ASBR
 - B) backbone router
 - C) ABR
 - D) internal router
- Q96) What is the correct configuration for stub area 10? (Source: Configuring OSPF Special Area Types)
- A) **area 10 stub-area**
 - B) **router ospf 10 stub**
 - C) **area 10 stub**
 - D) **area 10 stub no-summary**
- Q97) What is the meaning of the **no-summary** parameter of the **area x stub** command? (Source: Configuring 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 in the stub area.
 - D) There are no external LSAs in the stub area.
- Q98) The default route has a cost of 1 from the stub area ABR if no **area default-cost** command is used. (Source: Configuring OSPF Special Area Types)
- A) true
 - B) false
- Q99) Which characteristic relates to NSSA? (Source: Configuring OSPF Special Area Types)
- A) allows stub area benefits without meeting stub area requirements
 - B) is a Cisco proprietary technique
 - C) allows ASBRs but not virtual links
 - D) floods LSA type 7 into the backbone area
- Q100) A disadvantage of NSSA is that it does not have a totally stubby feature like a normal stub area. (Source: Configuring OSPF Special Area Types)
- A) true
 - B) false
- Q101) Which characteristic is not a prerequisite for stub areas? (Source: Configuring 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
- Q102) Stub area design will not improve _____. (Source: Configuring 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

- Q103) An LSA type 7 appears in the routing table as an _____. (Source: Configuring OSPF Special Area Types)
- A) O E1 route
 - B) O E2 route
 - C) O N2 route
 - D) O I/A route
- Q104) What is the difference between stub area and totally stubby area configuration? (Source: Configuring OSPF Special Area Types)
- A) **no-summary** option at the ABR
 - B) **area area-id totally-stubby** command at the internal routers
 - C) **area area-id nssa** command at the internal routers
 - D) **default-cost** command at the ABR
- Q105) A stub area blocks summary LSAs (type 3 and 4 LSAs). (Source: Configuring OSPF Special Area Types)
- A) true
 - B) false
- Q106) Where should you configure the **area area-id stub** command when you are configuring a stub area? (Source: Configuring 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
- Q107) Which two features are specific to Cisco? (Choose two.) (Source: Configuring OSPF Special Area Types)
- A) stub areas
 - B) totally stubby areas
 - C) NSSA
 - D) totally stubby NSSA
- Q108) In NSSA, the NSSA ABR translates type 7 LSAs into type 5 LSAs. (Source: Configuring OSPF Special Area Types)
- A) true
 - B) false
- Q109) The ABR injects a default route into which three types of areas? (Choose three.) (Source: Configuring OSPF Special Area Types)
- A) stub
 - B) totally stubby NSSA
 - C) totally stubby
 - D) area 0
- Q110) Which two types of authentication are used in OSPF? (Choose two.) (Source: Configuring OSPF Authentication)
- A) MD5
 - B) encrypted password
 - C) simple password
 - D) MD6

- Q111) When OSPF authentication is configured between two routers, each router has its own unique password. (Source: Configuring OSPF Authentication)
- A) true
 - B) false
- Q112) Which three of the following are used to generate the message digest when OSPF MD5 authentication is configured? (Choose three.) (Source: Configuring OSPF Authentication)
- A) packet
 - B) sequence number
 - C) key ID
 - D) key
 - E) router ID
- Q113) Which command is used to specify that OSPF simple password authentication is to be used? (Source: Configuring 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**
- Q114) Which command is used to specify that OSPF MD5 authentication is to be used? (Source: Configuring 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**
- Q115) When a new MD5 key is configured on a router for OSPF authentication, it will use both the old and new key until the new key is configured on neighboring routers. (Source: Configuring OSPF Authentication)
- A) true
 - B) false
- Q116) Which command is used to troubleshoot OSPF authentication? (Source: Configuring 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) 1 = B, 2 = C, 3 = A
- 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) B
- 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, C
- Q28) D
- Q29) C
- Q30) B
- Q31) B
- Q32) B
- Q33) B
- Q34) B
- Q35) D

- Q36) A
- Q37) A, C
- Q38) B
- Q39) B
- Q40) B
- Q41) B
- Q42) B
- Q43) B
- Q44) B
- Q45) B
- Q46) 1 = C, 2 = A, 3 = B
- Q47) A, B
- Q48) B, C
- Q49) D
- Q50) B
- Q51) A, B, E
- Q52) C
- Q53) A
- Q54) C, D
- Q55) B, D
- Q56) A, C
- Q57) B
- Q58) D
- Q59) B
- Q60) A, B, D
- Q61) D
- Q62) A = point-to-point
B = transit
C = stub
D = virtual link
- Q63) A = 5, B = 2, C = 3 and 4, D = 6, E = 1, F = 9 and 11, G = 7
- Q64) E
- Q65) D
- Q66) D, E
- Q67) D
- Q68) B
- Q69) B
- Q70) B

- Q71) A
- Q72) C
- Q73) C
- Q74) C
- Q75) A
- Q76) B, C
- Q77) A, D
- Q78) B
- Q79) A
- Q80) E
- Q81) A
- Q82) D
- Q83) B
- Q84) B
- Q85) C, D
- Q86) area 1 range 172.16.0.0 255.255.0.0
- Q87) summary-address 172.16.32.0 255.255.224.0
- Q88) B
- Q89) B
- Q90) D
- Q91) C
- Q92) B
- Q93) A
- Q94) B
- Q95) C
- Q96) C
- Q97) B
- Q98) A
- Q99) A
- Q100) B
- Q101) C
- Q102) C
- Q103) C
- Q104) A
- Q105) B
- Q106) A
- Q107) B, D

- Q108) A
- Q109) A, B, C
- Q110) A, C
- Q111) B
- Q112) A, C, D
- Q113) B
- Q114) E
- Q115) A
- Q116) A

The IS-IS Protocol

Overview

This module provides an overview of Intermediate System-to-Intermediate System Protocol (IS-IS), as well as basic configuration examples. IS-IS is a part of the Open Systems Interconnection (OSI) suite of protocols.

The OSI suite uses Connectionless Network Service (CLNS) to provide connectionless delivery of data, and the actual Layer 3 protocol is Connectionless Network Protocol (CLNP). CLNP is the solution for “unreliable” (connectionless) delivery of data, similar to IP. IS-IS uses CLNS addresses to identify the routers and to build the link-state database (LSDB).

IS-IS operates in strictly CLNS terms; however, Integrated IS-IS supports IP routing as well as CLNS. CLNS addresses are required to configure and troubleshoot IS-IS, even when it is used only for IP. IS-IS supports different data-link environments, such as Ethernet and Frame Relay.

IS-IS supports the most important characteristics of Open Shortest Path First (OSPF) and Enhanced Interior Gateway Routing Protocol (EIGRP), because it supports variable-length subnet masking (VLSM) and converges quickly. Each protocol has advantages and disadvantages, but this commonality makes any of the three scalable and appropriate for supporting the large-scale networks of today.

Module Objectives

Upon completing this module, you will be able to configure integrated IS-IS in a single area. This ability includes being able to meet these objectives:

- Describe the features and benefits of IS-IS as a routing protocol for large networks
- Describe IS-IS operation
- Implement Integrated IS-IS in an enterprise network

Lesson 1

Introducing IS-IS and Integrated IS-IS Routing

Overview

Intermediate System-to-System Protocol (IS-IS) is a proven and extensible IP routing protocol that converges quickly and supports variable length subnet masking (VLSM). IS-IS is a public standard, published as International Organization for Standardization (ISO) 9542 and republished as RFC 995. Integrated IS-IS (or dual IS-IS) is specified in RFC 1195 and offers support for IP and OSI protocols.

Although not as common, IS-IS is comparable to, and in some cases preferable to, Open Shortest Path First Protocol (OSPF).

This lesson describes IS-IS and Integrated IS-IS routing and compares Integrated IS-IS with OSPF. It also addresses some of the concepts necessary to develop an understanding of Integrated IS-IS.

Objectives

Upon completing this lesson, you will be able to describe the features and benefits of IS-IS as a routing protocol for large networks. This ability includes being able to meet these objectives:

- Describe IS-IS routing and some of the ways in which IS-IS is used
- Describe the features of Integrated IS-IS routing
- Explain the principles and issues of Integrated IS-IS design
- Describe the features of the ES-IS protocol
- Describe how to differentiate among the four OSI routing levels
- Explain the similarities and differences between IS-IS and OSPF

IS-IS Routing

This topic describes IS-IS routing and some of the ways in which IS-IS is used.

Uses for IS-IS Routing

Large ISPs

- **Stable protocol**
- **Originally deployed by ISPs because U.S. government mandated Internet support of OSI and IP**

© 2006 Cisco Systems, Inc. All rights reserved. BSCI v3.0—4.2

IS-IS is a popular IP routing protocol in the Internet service provider (ISP) industry. The simplicity and stability of IS-IS make it robust in large internetworks. IS-IS is found in large ISPs and in some networks that support Open Systems Interconnection (OSI) protocols.

IS-IS development began before development of OSPF. Large ISPs chose IS-IS because of their unique requirement for scalability, convergence, and stability. The U.S. government also required support for OSI protocols in the early Internet. Although this requirement was later dropped, IS-IS met both constraints.

Later, businesses typically chose OSPF because it was a more widely supported native IP protocol. Today it is harder to find information and expertise on IS-IS than on OSPF. Nevertheless, some of the largest networks in the world persist in using IS-IS, which is a tribute to its capabilities.

IS-IS Routing

- **IS = router.**
- **IS-IS was originally designed as the IGP for the Connectionless Network Service (CLNS), part of the OSI protocol suite.**
- **The OSI protocol suite layer 3 protocol is the Connectionless Network Protocol (CLNP).**
- **IS-IS uses CLNS addresses to identify routers and build the LSDB.**

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—4-3

ISO specifications refer to routers as intermediate systems (ISs). Thus, IS-IS is a protocol that allows routers to communicate with other routers.

The OSI suite uses Connectionless Network Service (CLNS) to provide connectionless delivery of data, and the actual Layer 3 protocol is Connectionless Network Protocol (CLNP).

IS-IS uses CLNS addresses to identify the routers and to build the link-state database (LSDB). IS-IS serves as an interior gateway protocol (IGP) for the CLNS.

CLNP is the solution for “unreliable” (connectionless) delivery of data, similar to IP.

IS-IS Features

- **Link-state routing protocol**
- **Supports VLSM**
- **Uses Dijkstra's SPF algorithm; has fast convergence**
- **Uses hellos to establish adjacencies and LSPs to exchange link-state information**
- **Efficient use of bandwidth, memory, and processor**
- **Supports two routing levels:**
 - **Level 1: Builds common topology of system IDs in local area and routes within area using lowest cost path.**
 - **Level 2: Exchanges prefix information (area addresses) between areas. Routes traffic to area using lowest-cost path.**

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—4-4

IS-IS is the dynamic link-state routing protocol for the OSI protocol stack. It distributes routing information for routing CLNP data for the ISO CLNS environment.

IS-IS operates similarly to OSPF. IS-IS allows the routing domain to be partitioned into areas. IS-IS routers establish adjacencies using a Hello protocol and exchange link-state information, using link-state packets (LSPs), throughout an area to build the LSDB.


Each router then runs Dijkstra's SPF algorithm against its LSDB to pick the best paths. There is a minimal amount of information communicated between areas, which reduces the burden on routers supporting the protocol.

IS-IS routing takes place at two levels within an AS: Level 1 and Level 2.

Level 1 routing occurs within an IS-IS area. It recognizes the location of the end systems (ESs) and ISs, and then builds a routing table to reach each system. All devices in a Level 1 routing area have the same area address. Routing within an area is accomplished by looking at the locally significant address portion (known as the system ID) and choosing the lowest-cost path.

Level 2 routers learn the locations of Level 1 routing areas and build an interarea routing table. All ISs in a Level 2 routing area use the destination area address to route traffic using the lowest-cost path.

IS-IS Link-State Operation



Routers are identified as Level 1, Level 2, or Level 1-2:

- **Level 1 routers** use LSPs to build topology for local area.
- **Level 2 routers** use LSPs to build topology between different areas.
- **Level 1-2 routers** act as border routers between Level 1 and Level 2 routing domains.

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—4-5

To support the two routing levels, IS-IS defines three types of routers:

- **Level 1:** Level 1 routers learn about paths within the areas they connect to (intra-area).
- **Level 2:** Level 2 routers learn about paths between areas (interarea).
- **Level 1-2:** Level 1-2 routers learn about paths both within and between areas. Level 1-2 routers are equivalent to ABRs in OSPF.

The path of connected Level 2 and Level 1-2 routers is called the backbone. All areas and the backbone must be contiguous.

Note	Area boundaries fall on the links. Each IS-IS router belongs to exactly one area. Neighboring routers learn whether they are in the same area or different areas and negotiate appropriate adjacencies: Level 1, Level 2, or both.

Integrated IS-IS Routing

This topic describes the features of Integrated IS-IS routing.

Integrated (or Dual) IS-IS Routing

- **Integrated IS-IS is IS-IS for multiple protocols:**
 - For IP, CLNS, or both
- **Uses its own PDUs to transport IP routing information; updates not sent in IP packets**
- **Requires CLNS addresses, even if only routing for IP**

© 2006 Cisco Systems, Inc. All rights reserved. BSCI v3.0—4-8

Integrated IS-IS or dual IS-IS is an implementation of the IS-IS protocol for routing multiple network protocols, IP and CLNS (and is specified in RFC 1195 and ISO 10589).

Integrated IS-IS tags CLNP routes with information about IP networks and subnets. Integrated IS-IS provides IP with an alternative to OSPF and combines ISO CLNS and IP routing in one protocol. Integrated IS-IS can be used for IP routing, CLNS routing, or for a combination of the two.


Integrated IS-IS uses its own protocol data units (PDUs) to transport information between routers, including IP reachability information. IS-IS information is not carried within a network-layer protocol but is instead carried directly within data link layer frames.

Note This protocol independence makes IS-IS easily extensible; there is also a version of Integrated IS-IS that supports IP version 6 (IPv6).

Since IS-IS uses CLNS addresses to identify the routers and to build the LSDB, an understanding of CLNS addresses is required to configure and troubleshoot IS-IS, even when it is used only for routing IP.

Principles and Issues of Integrated IS-IS Design

This topic explains the principles and issues of Integrated IS-IS design.




Effective networks are well planned. The first and most important step in building a scalable network is developing a good addressing plan that allows for route summarization. Route summarization is possible only when using a hierarchical addressing structure.

Effective address planning presents opportunities to group devices into areas. Using areas confines the scope of LSP propagation and saves bandwidth. Level 1-2 routers, on the border between a Level 1 area and the Level 2 backbone, are logical places to implement route summarization.

Route summarization saves memory because each IS is no longer responsible for the LSPs of the entire routing domain. Route summarization also saves CPU usage because a smaller routing table is easier to maintain.

Issues with Integrated IS-IS



- **Default narrow metrics are limited to 6-bit interface and 10-bit path metric**
 - In Cisco IOS Software Release 12.0, wide metrics allow 24-bit interface and 32-bit path metric.
- **Cisco IOS software has default metric of 10 on all interfaces.**

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0-4-8

One issue with IS-IS is that older implementations, those using the narrow metrics, are limited to a maximum interface metric of 63 (6 bits) and a maximum total path metric of 1,023 (10 bits). There is little room to distinguish between paths.


Cisco IOS software, beginning in Software Release 12.0, supports wide metrics that allow a 24-bit interface and 32-bit path metrics. The default, however, is still the narrow metrics.

Note Complications can occur if you use wide metrics along with narrow metrics (for example, on older routers or in a multivendor environment).

IS-IS as implemented on Cisco routers does not automatically scale the interface metric. Instead, all IS-IS interfaces have a default metric of 10; this setting can be changed manually. If the default metric is not adjusted on each interface, the IS-IS metric becomes similar to the hop count metric used by the Routing Information Protocol (RIP).

The ES-IS Protocol

The End System-to-Intermediate System (ES-IS) protocol permits ESs (hosts) and ISs (routers) to discover one another. ES-IS also allows ESs to learn their network-layer addresses. This topic describes the features of the ES-IS protocol.



Hosts in the OSI terminology are called *end systems*. ES-IS handles topology information discovery and exchange between ESs (hosts) and ISs (routers).

ES-IS performs the following tasks:

- Identifies the area (prefix) to the ESs
- Creates adjacencies between ESs and ISs
- Creates data link-to-network address mappings

ESs send End System Hellos (ESHs) to well-known addresses that announce their presence to ISs. Routers listen to ESHs to find the ESs on a segment. Routers include information about ESs in LSPs.

Routers transmit Intermediate System Hellos (ISHs) to well-known addresses, announcing their presence to ESs. ESs listen for these ISHs and randomly pick an IS to which they forward all their packets. When an ES needs to send a packet to another ES, it sends the packet to one of the ISs (routers) on its directly attached network.


Routers use IS-IS Hellos (IIHs) for establishing and maintaining adjacencies between ISs.

IP systems do not use ES-IS. IP has its own processes and applications to handle the same functions as ES-IS, such as Internet Control Message Protocol (ICMP), Address Resolution Protocol (ARP), and DHCP.

Although Integrated IS-IS can support IP exclusively, IS-IS still uses CLNS to transmit reachability information and still forms adjacencies using IIHs.

OSI Routing Levels

This topic describes how to differentiate among the four OSI routing levels.



The OSI specifications discuss four unique types of routing operations, numbered 0 to 3. As discussed earlier, IS-IS is responsible for Level 1 and Level 2 OSI routing.

Level 0 Routing

OSI routing begins with ES-IS, when the ESs discover the nearest IS by listening to ISH packets.

When an ES needs to send a packet to another ES, it sends the packet to an IS on an attached network. This process is known as Level 0 routing.

IS-IS Level 1 Routing

Each ES and IS resides in a particular area. To pass traffic, the router looks up the destination address and forwards the packet along the best route. If the destination is on the same subnetwork, the IS is aware of the location (from listening to the ESH) and forwards the packet appropriately.

The IS can also provide a redirect message to the source that tells it that a more direct route is available. If the destination is on a different subnetwork but within the same area, the router identifies the best path using the system ID and forwards the traffic appropriately.

Note Level 1 routing is also called intra-area routing.

IS-IS Level 2 Routing

If a destination address is in another area, the Level 1 IS sends the packet to the nearest Level 1-2 IS; this process is called Level 2 routing. Packet forwarding continues through Level 2 ISs until the packet reaches a Level 1-2 or Level 2 IS in the destination area. Within the destination area, ISs forward the packet along the best path, based on system ID, until the packet reaches the destination.

Note Level 2 routing is also called interarea routing.

Level 3 Routing

Routing between separate domains is called Level 3 routing. Level 3 routing is comparable to Border Gateway Protocol (BGP) interdomain routing in IP. Level 3 routing passes traffic between different autonomous systems, which might have different routing logic and so might not have metrics that can be directly compared. Level 3 OSI routing is not implemented on Cisco routers but is specified as being accomplished through the Interdomain Routing Protocol (IDRP).

Summary of Routing Levels

Routing levels can be summarized as follows:

- Level 0 routing is conducted by ES-IS.
- Level 1 and Level 2 routing are functions of IS-IS.
- IDRP conducts Level 3 routing. IDRP is similar in purpose to BGP. Cisco Systems routers do not support IDRP.

Comparing IS-IS to OSPF

This topic explains the similarities and differences between IS-IS and OSPF in terms of features, topologies, and the advantages and disadvantages of each.

Similarities Between IS-IS and OSPF

- **Integrated IS-IS and OSPF are both open standard link-state protocols with the following similar features:**
 - **Link-state representation, aging timers, and LSDB synchronization**
 - **SPF algorithms**
 - **Update, decision, and flooding processes**
 - **VLSM support**
- **Scalability of link-state protocols has been proven (used in ISP backbones).**
- **They both converge quickly after changes.**

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—4-11

IS-IS and OSPF are more similar than dissimilar. Both routing protocols have the following characteristics:

- They are open-standard link-state routing protocols.
- They support VLSM.
- They use similar mechanisms (link-state advertisements [LSAs], link-state aging timers, and LSDB synchronization) to maintain the health of the LSDB.
- They use the SPF algorithm, with similar update, decision, and flooding processes.
- They are successful in the largest and most demanding deployments (ISP networks).
- They converge quickly after network changes.

Most of the development of these two protocols was done concurrently. The work of the development groups produced two protocols that are very similar, and each is better because of the other. The practical differences between the two protocols deal with perceived issues of resource usage and ability to customize.

Most debates of the merits of these protocols are colored by their mutual history; different groups with different cultures developed the protocols.

Digital Equipment Corporation originally developed IS-IS for DECnet Phase V. In 1987, it was selected by the American National Standards Institute (ANSI) to be the OSI IGP. At that time, it was capable of routing CLNP only.


The ISO process is an international standards development process. According to an account given by Christian Huitema in his book *Routing in the Internet*, groups within ISO and outside the United States did not approve of TCP/IP because of its origin (it was also called the U.S. Department of Defense protocol).

From the perspective of ISO, IP development was chaotic and imprecise, based on the famous maxim of “loose consensus and running code.” From the perspective of the early Internet engineers, the ISO process was slow, irritating, and disenfranchising.

In 1988, the U.S. National Science Foundation Network (NSFnet) was created. The IGP used was based on an early draft of IS-IS. The extensions to IS-IS for handling IP were developed in 1988. OSPF development began during this time and was loosely based on IS-IS.

In 1989, OSPF version 1 (OSPF v1) was published, and conflict ensued between the proponents of IS-IS and OSPF. The Internet Engineering Task Force (IETF) eventually supported both, although it continued to favor OSPF. With the unofficial endorsement of the IETF, OSPF eventually became the more popular protocol.


By the mid-1990s, large ISPs in need of an IGP selected IS-IS for two reasons. First, IS-IS supported both CLNS and IP, which solved two problems at once. Second, OSPF was seen as immature at the time.



With OSPF, network design is constrained by the fact that OSPF is based on a central backbone, area 0, with all other areas being physically attached to area 0. The border between areas is inside the ABRs; each link is in only one area.

When you use this type of hierarchical model, a consistent IP addressing structure is necessary to summarize addresses into the backbone. Summarization also reduces the amount of information carried in the backbone and advertised across the network.

Integrated IS-IS vs. OSPF: Area Design (Cont.)




In IS-IS the area borders lie on links

- Each IS-IS router belongs to exactly one area.
- IS-IS is more flexible when extending the backbone.

In comparison, IS-IS has a hierarchy of Level 1 and Level 2 or Level 1-2 routers, and the area borders lie on links.

IS-IS permits a more flexible approach to extending the backbone. The backbone can be extended by simply adding more Level 2 and Level 1-2 routers, a less complex process than with OSPF.

Advantages of Integrated IS-IS



The differences between OSPF and IS-IS are small, but they do exist.

OSPF produces many small LSAs. IS-IS updates are grouped by the router and sent as one LSP. Thus, as network complexity increases, the number of IS-IS updates is not an issue. Each update packet must be routed, though, and routing takes network resources, so more packets represent a larger impact on the network. Since IS-IS uses significantly fewer LSPs, more routers, at least 1,000, can reside in a single area, making IS-IS more scalable than OSPF.

OSPF runs on top of IP, whereas IS-IS runs through CLNS.

IS-IS is also more efficient than OSPF in the use of CPU resources and in the way it processes routing updates. Not only are there fewer LSPs to process (LSAs, in OSPF terminology) but also the mechanism by which IS-IS installs and withdraws prefixes is less intensive because it uses network entity title (NET) addresses, which are already summarized.

Both OSPF and IS-IS are link-state protocols and thus provide fast convergence. The convergence time depends on a number of factors, such as timers, number of nodes, and type of router. Based on the default timers, IS-IS detects a failure faster than OSPF; therefore, convergence occurs more rapidly. If there are many neighboring routers and adjacencies, the convergence time may also depend on the processing power of the router. IS-IS is less CPU-intensive than OSPF.

New ideas are not easily expressed in OSPF packets; they require the creation of a new LSA. The OSPF description schema is difficult to extend, because of compatibility issues and because it was developed exclusively for IPv4. IS-IS is easy to extend through the Type, Length, Value (TLV) mechanism. TLV strings, called *tuples*, encode all IS-IS updates. IS-IS can easily grow to cover IPv6 or any other protocol, because extending IS-IS consists simply of creating new type codes.

Advantages of OSPF

- **OSPF has more features, including:**
 - **Has three area types: normal, stub, and NSSA**
 - **Defaults to scaled metric (IS-IS always 10)**
- **OSPF is supported by many vendors.**
- **Information, examples, and experienced engineers are easier to find.**

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—4-19

A company may choose OSPF over IS-IS because OSPF is more optimized and was designed exclusively as an IP routing protocol. For example, OSPF defines different area types (normal, stub, and not-so-stubby [NSSA]). The default OSPF metric is related to the interface bandwidth, while IS-IS defaults to a metric of 10 on all interfaces.

If a company does choose OSPF, it will require networking equipment that supports OSPF and network engineers that are familiar with OSPF theory and operation. It is relatively easy to find both equipment and personnel to support an OSPF infrastructure. Furthermore, OSPF documentation is much more readily available than documentation for IS-IS.

Summary of Differences between OSPF and Integrated IS-IS

The table summarizes the differences between OSPF and Integrated IS-IS.

OSPF	Integrated IS-IS
Area border inside routers (ABRs)	Area border on links
Each link in only one area	Each router in only one area
More complex to extend backbone	Simple extension of backbone
Many small LSAs sent	Fewer LSPs sent
Runs on top of IP	Runs on top of data link layer
Requires IP addresses	Requires IP and CLNS addresses
Default metric is scaled by interface bandwidth	Default metric is 10 for all interfaces
Not easy to extend	Easy to support new protocols with new TLV tuples
Equipment, personnel, and information more readily available	Equipment, personnel, and information not as easily available

Summary

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

Summary

- **IS-IS is a popular routing protocol in the ISP industry.**
- **IS-IS is a stable, fast-converging IGP that is positioned to route IPv4, CLNS, or IPv6.**
- **All IS-IS interfaces have a default metric of 10.**
- **ES-IS (for CLNS routing only) provides discovery between host and routers using hello packets to form adjacencies. Hosts send ESHs, while routers send ISHs.**
- **OSI defines routing levels 0 through 3. Level 0 is between ES and IS. Levels 1 and 2 are between IS and IS to support intradomain routing. Level 3 supports interdomain routing.**
 - **Level 1 is intra-area**
 - **Level 2 is interarea.**
- **IS-IS and OSPF are both open-standard link-state routing protocols that support VLSM, scalability, and quick convergence.**

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—4-16

Performing IS-IS Routing Operations

Overview

Unlike IP addresses, Connectionless Network Service (CLNS) addresses apply to entire nodes and not to interfaces. Because Intermediate System-to-Intermediate System Protocol (IS-IS) was originally designed for CLNS, IS-IS requires CLNS addresses, even if the router is used only for routing IP. CLNS addresses that are used by routers are called network service access points (NSAPs). One part of an NSAP address is the NSAP selector (NSEL) byte. When an NSAP is specified with an NSEL of 0, then the NSAP is called the network entity title (NET). This lesson starts by describing NSAP and NET addresses for use with Integrated IS-IS.

The lesson then describes how CLNS addressing affects IS-IS operation and how the IS-IS protocol learns topology, makes routing decisions, and handles different data links.

Objectives

Upon completing this lesson, you will be able to describe IS-IS operation. This ability includes being able to meet these objectives:

- Describe the features of and applications for NSAP addresses
- Describe the features of and applications for NET addresses
- Describe routing levels that are associated with IS-IS
- Describe the features of intra-area and interarea addressing and routing
- Describe the four types of IS-IS PDUs
- Describe how routers use LSPs
- Explain the different IS-IS network types and the considerations for selecting broadcast or point-to-point modes when implementing IS-IS in ATM and Frame Relay networks
- Describe the features of a broadcast networks
- Describe the levels of LSPs and IIHs
- Describe the types of LSDB synchronization

NSAP Addresses

This topic describes the features of and applications for NSAP addresses.

OSI Addresses

- **OSI network layer addressing is implemented with NSAP addresses.**
- **An NSAP address identifies a system in the OSI network; an address represents an entire node, not an interface.**
- **Various NSAP formats are used in various systems, because different protocols may use different representations of NSAP.**
- **NSAP addresses are a maximum of 20 bytes:**
 - **Higher-order bits identify the interarea structure.**
 - **Lower-order bits identify systems within area.**

© 2006 Cisco Systems, Inc. All rights reserved. BSCI v3.0—4-2

CLNS addresses that are used by routers are called NSAP addresses. Unlike IP addresses, NSAP addresses apply to entire nodes and not to interfaces.

IS-IS link-state packets (LSPs) use NSAP addresses to identify the router and build the topology table and the underlying IS-IS routing tree; therefore IS-IS requires NSAP addresses to function properly, even if it is used only for routing IP. NSAP addresses contain the following:


- Open Systems Interconnection (OSI) address of the device
- Link to the higher-layer process

The NSAP address is equivalent to the combination of the IP address and upper-layer protocol in an IP header.

NSAP addresses have a maximum size of 20 bytes. The high-order bits identify the interarea structure, and the low-order bits identify unique systems within an area.

There are a variety of NSAP address formats.

Integrated IS-IS NSAP Address Structure



- The Cisco implementation of Integrated IS-IS distinguishes only the following three fields in the NSAP address:
 - Area address: Variable-length field (1 to 13 octets) composed of the higher-order NSAP octets, excluding system ID and NSEL.
 - System ID: ES or IS identifier in an area; fixed length of six octets in Cisco IOS software.
 - NSEL: One octet NSAP selector, service identifier.
- Total length of NSAP is from 8 (minimum) to 20 octets (maximum).

The Cisco implementation of Integrated IS-IS divides the NSAP address into three fields: the area address, the system ID, and the NSEL. Cisco routers routing CLNS use addressing that conforms to the ISO 10589 standard. ISO NSAP addresses consist of the following:

- The authority and format identifier (AFI) and the initial domain identifier (IDI) make up the initial domain part (IDP) of the NSAP address. The IDP corresponds roughly to an IP classful major network.
 - The AFI byte specifies the format of the address and the authority that assigned that address. Some valid values are shown in the table.

AFI Value	Address Domain
39	ISO Data Country Code (DCC)
45	E.164
47	ISO 6523 International Code Designator (ICD)
49	Locally administered (private)

- Addresses starting with the AFI value of 49 are private addresses, analogous to RFC 1918 for IP addresses. IS-IS routes these addresses; however, this group of addresses should not be advertised to other CLNS networks because they are ad hoc addresses. Other companies that use a value of 49 may have created different numbering schemes that, when used together, could create confusion.
- The IDI identifies a subdomain under the AFI. For instance, 47.0005 is assigned to civilian departments of the U.S. government, and 47.0006 is assigned to the U.S. Department of Defense.

- The domain specific part (DSP) contributes to routing within an IS-IS routing domain. The DSP comprises the high-order domain specific part (HO-DSP), the system ID, and the NSEL.
 - The HO-DSP subdivides the domain into areas. The HO-DSP is more or less the OSI equivalent of a subnet in IP.
 - The system ID identifies an individual OSI device. In OSI, a device has an address just as it does in DECnet, while in IP each interface has an address.
 - The NSEL identifies a process on the device and corresponds roughly to a port or socket in IP. The NSEL is not used in routing decisions.

Typical NSAP Address Structure

The simplest NSAP format used by most companies running IS-IS as their IGP is as follows:

- **Area address (must be at least 1 byte)**
 - **AFI set to 49**
 - **Locally administered; thus, you can assign your own addresses.**
 - **Area ID**
 - **The octets of the area address after the AFI.**
- **System ID**
 - **Cisco routers require a 6-byte system ID.**
- **NSEL**
 - **Always set to 0 for a router.**

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—4-4

The simplest NSAP format, used by most companies running IS-IS as their interior gateway protocol (IGP), comprises the following:

- **The area address:** It must be at least 1 byte, separated into two parts:
 - The AFI, set to 49, which signifies that the AFI is locally administered and thus individual addresses can be assigned by the company
 - The area identifier (ID), the octets of the area address after the AFI
- **A system ID:** Cisco routers compliant with the U.S. Government OSI Profile (GOSIP) version 2.0 standards require a 6-byte system ID.
- **The NSEL:** It must always be set to 0 for a router.

Note The NSAP is called the NET when it has an NSEL of 0. Routers use the NET to identify themselves in the IS-IS protocol data units (PDUs).


For example, you might assign 49.0001.0000.0c12.3456.00, which represents the following:

- AFI of 49
- Area ID of 0001
- System ID of 0000.0c12.3456, the MAC address of a LAN interface
- NSEL of 0

Note The area address is also referred to as the prefix.

Note Some IS-IS documentation uses the terms “area ID” and “area address” as synonyms.

Identifying Systems in IS-IS: Area Address



The area address uniquely identifies the routing area, and the system ID identifies each node.

- All routers within an area must use the same area address.
- An ES may be adjacent to a router only if they share a common area address.
- Area address is used in Level 2 routing.

© 2006 Cisco Systems, Inc. All rights reserved.


BSCI v3.0—4-5

The first part of an NSAP is the area address and is associated with the IS-IS routing process. Unlike Open Shortest Path First Protocol (OSPF), an IS-IS router can be a member of only one area.

All routers in an area must use the same area address, which actually defines the area. The area address is used in Level 2 routing.

End systems (ESs) recognize only intermediate systems (ISs) and other ESs on the same subnetwork that share the same area address.

Identifying Systems in IS-IS: System ID



- **System ID in the address used to identify the IS; it is not just an interface. Cisco supports only a 6-byte system ID.**
- **System ID is used in Level 1 routing and has to be unique within an area.**
- **System ID has to be unique within Level 2 routers that form the routing domain.**
- **General recommendation: use domain-wide unique system ID.**
 - This may be **MAC** (for example, 0000.0c12.3456) or **IP address** (for example, 1921.6800.0001) taken from an interface.

The 6-byte NSAP system ID must be unique within an area. It is customary to use a MAC address from the router, or, for Integrated IS-IS, to encode an IP address into the system ID. All of the system IDs in a domain must be of equal length. Cisco enforces this OSI directive by fixing the length of the system ID at 6 bytes in all cases.

Level 1 intra-area routing is based on system IDs; therefore, each ES and IS must have a unique system ID within the area.

All Level 2 ISs eventually recognize all other ISs in the Level 2 backbone; therefore, they must also have unique system IDs.

Thus, system IDs should remain unique across the domain. If the system IDs remain unique, there can never be a conflict at Level 1 or Level 2 if, for example, a device moves into a different area.

NET Addresses

This topic identifies the features of and applications for NET addresses.

OSI Addressing: NET Addresses

- **NSAP address includes NSEL field (process or port number)**
- **NET: NSAP with a NSEL field of 0**
 - **Refers to the device itself (equivalent to the Layer 3 OSI address of the device)**
 - **Used in routers because they implement the network layer only (base for SPF calculation)**

© 2006 Cisco Systems, Inc. All rights reserved. BSCI v3.0—4-7


As discussed earlier, NSAP address have a one-octet NSEL field that identifies a process on the device, corresponding roughly to a port number in IP. NET addresses are NSAP addresses with an NSEL value of 0.

A NET address is used to uniquely identify an OSI host within an IS-IS routing domain. Because IS-IS originates from the OSI world, NET addresses are required even if the only routed protocol is IP.

The NET refers to the device itself; that is, it is the equivalent of the Layer 3 OSI address of that device..

Routers use the NET to identify themselves in the LSPs and to form the basis for the OSI routing calculation.

Subnetwork Point of Attachment (SNPA) and Circuit



SNPA is equivalent to Layer 2 address; for example:

- Virtual circuit ID (DLCI on Frame Relay)
- MAC address on LAN interfaces

Interfaces uniquely identified by circuit ID:

- On point-to-point interfaces, SNPA is used.
- On LANs, circuit ID concatenated with six-octet system ID of a designated IS to form seven-octet LAN ID (for example, 1921.6800.0001.01) is used.
- Cisco routers use host name instead of system ID (for example, "R1.01").

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—4-6

Three additional IS-IS terms are subnetwork point of attachment (SNPA), circuit, and link.

The SNPA is the point that provides subnetwork services. The SNPA is the equivalent of the Layer 2 address corresponding to the NET or NSAP address.

The SNPA is assigned by using one of the following:

- The MAC address on a LAN interface.
- The virtual circuit ID from X.25 or ATM connections, or the data-link connection identifier (DLCI) from Frame Relay connections.
- For High-Level Data Link Control (HDLC) interfaces, simply "HDLC"

A circuit is the IS-IS term for an interface. Since the NSAP and NET refer to the entire device, a circuit ID is used to distinguish a particular interface. The router assigns a circuit ID (one octet) to each interface on the router as follows:

- In the case of point-to-point interfaces, the SNPA is the sole identifier for the circuit. For example, on an HDLC point-to-point link, the circuit ID is 0x00.
- In the case of LAN interfaces, the circuit ID is tagged to the end of the system ID of the Designated IS (DIS) to form a 7-byte LAN ID, for example, 1921.6800.0001.01. On Cisco routers, the router host name is used instead of the system ID; therefore the circuit ID of a LAN interface may look like "R1.01."

A link is the path between two neighbor ISs and is defined as being up when communication is possible between the two neighbor SNPAs.

IS-IS Routing Levels

This topic describes the different routing levels that are associated with IS-IS.

Level 1, Level 2, and Level 1-2 Routers

Level 1 (like OSPF internal nonbackbone routers):

- Intra-area routing enables ESs to communicate.
- Level 1 area is a collection of Level 1 and Level 1-2 routers.
- Level 1 IS keeps copy of the Level 1 area LSDB.

Level 1-2 (like OSPF ABR):

- Intra-area and interarea routing.
- Level 1-2 IS keeps separate Level 1 and Level 2 LSDBs and advertises default route to Level 1 routers.

Level 2 (like OSPF backbone routers):

- Interarea routing.
- Level 2 (backbone) area is a contiguous set of Level 1-2 and Level 2 routers.
- Level 2 IS keeps a copy of the Level 2 area LSDB.

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—4-9

Recall that IS-IS defines three types of routers as follows:

- **Level 1:** Level 1 routers learn about paths within the areas they connect to (intra-area).
- **Level 2:** Level 2 routers learn about paths between areas (interarea).
- **Level 1-2:** Level 1-2 routers learn about paths both within and between areas. Level 1-2 routers are equivalent to ABRs in OSPF.

The path of connected Level 2 and Level 1-2 routers is called the backbone. All areas and the backbone must be contiguous.

Area boundaries fall on the links. Each IS-IS router belongs to exactly one area. Neighboring routers learn whether they are in the same area or different areas and negotiate appropriate adjacencies, Level 1, Level 2, or both. Each router keeps a copy of the link-state databases (LSDBs) for the levels that it is responsible for.

A Level 1-2 router automatically advertises to all Level 1 routers (within its area) that it is a potential exit point of the area. Level 1 routers default to the nearest attached Level 1-2 router.

Intra-Area and Interarea Addressing and Routing

This topic describes the features of intra-area and interarea addressing and routing.

Addressing and Routing

- **Area address is used to route between areas; system ID is not considered.**
- **System ID is used to route within an area; area address is not considered.**

© 2006 Cisco Systems, Inc. All rights reserved. BSCI v3.0—4-10

IS-IS routing flows naturally from the OSI address plan in which areas are identified and unique system IDs are given to each device.

The area address portion of the NSAP address can range from 1 to 13 bytes in length, as specified by the ISO standard. Therefore, an NSAP for an IS-IS network can be as little as 8 bytes in length. The NSAP is usually longer to permit some granularity in the allocation of areas. The area address prefix is common to all devices in an area and unique for each area. ISs and ESs are in the same area if they share the same area address.

Routing within an area involves collecting system IDs and adjacencies for all ISs and ESs in an area and using Dijkstra's algorithm to compute best paths between devices. Level 1 routers are aware only of the local area topology. They pass the traffic destined outside the area to the closest Level 1-2 router.

Routing between areas is based on area address. Level 2 routers in different areas exchange area address information and use Dijkstra's algorithm to compute best paths between areas. They pass traffic between areas to the closest Level 1-2 router.

OSI IS-IS Routing Logic

Level 1 router: For a destination address, compare the area address to this area.

- If not equal, pass to nearest Level 1-2 router.
- If equal, use Level 1 database to route by system ID.

Level 1-2 router: For a destination address, compare the area address to this area.

- If not equal, use Level 2 database to route by area address.
- If equal, use Level 1 database to route by system ID.

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—4-11

When an ES is required to send a packet to another ES, the packet goes to one of the ISs on a network directly attached to the ES. The router then searches for the destination address and forwards the packet along the best route.


If the destination ES is in the same area, the local IS recognizes the location by listening to End System Hello (ESH) packets and forwards the packet appropriately.

If the destination address is an ES in another area, the Level 1 IS sends the packet to the nearest Level 1-2 IS. Forwarding through Level 2 ISs continues until the packet reaches a Level 2 IS in the destination area.

Within the destination area, ISs forward the packet along the best path until the destination ES is reached.

Because each router makes its own best-path decisions at every hop along the way, there is a significant chance that paths will not be reciprocal. That is, return traffic can take a different path than the outgoing traffic. For this reason, it is important to know the traffic patterns within your network and tune IS-IS for optimal path selection if necessary.

Example: Identifying Systems: OSI Addressing in Networks




Example: Identifying Systems—OSI Addressing in Networks

Consider traffic from router R7 to router R9.

1. R7 recognizes that the prefix (49.00CC) of R9 is not the same as the prefix (49.00BB) of R7. R7 therefore passes the traffic to the closest Level 1-2 router, R5. R7 uses its Level 1 topology database to find the best path to R5.
2. R5 uses its Level 2 topology database to pick the best next hop to reach the prefix 49.00CC: R3. R5 does not use the destination system ID in this decision.
3. R3 uses its Level 2 topology database to pick the best next hop to reach the prefix 49.00CC: R1. R3 does not use the destination system ID in this decision.
4. R1 uses its Level 2 topology database to pick the best next hop to reach the prefix 49.00CC: R8. R1 does not use the destination system ID in this decision.
5. R8 recognizes that the prefix (49.00CC) of R9 is the same as the prefix (49.00CC) of R8. R8 therefore passes the traffic to R9 using its Level 1 topology database to find the best path.

Example: OSI Area Routing



© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—4-13

Example: OSI Area Routing

In the figure, area 1 contains two routers:

- One router borders area 2 and is a Level 1-2 IS.
- The other router is contained within the area and is a Level 1 only.

Area 2 has many routers:

- A selection of routers is specified as Level 1. The routers route either internally to that area or to the exit points (the Level 2 routers).
- Level 1-2 routers form a chain across the area linking to the neighbor areas. Although the middle router of the three Level 1-2 routers does not link directly to another area, the middle router must support Level 2 routing to ensure that the backbone is contiguous. If the middle router fails, the other Level 1-only routers cannot perform the Level 2 function (despite providing a physical path across the area), and the backbone is broken.

Area 3 contains one router that borders areas 2 and 4, yet it has no intra-area neighbors and is performing Level 2 functions only. If you add another router to area 3, the border router reverts to Level 1-2 functions.

As the figure shows, the border between the areas in an IS-IS network is the link between Level 2 routers. (This is in contrast to OSPF, where the border exists inside the ABR itself.)

In the figure, symmetric routing does not occur because Level 2 details are hidden from Level 1 routers, which recognize only a default route to the nearest Level 1-2 router. Traffic from router X to router Y flows from router X to its closest Level 1-2 router. The Level 1-2 router then forwards the traffic along the shortest path to the destination area (area 2). When the traffic flows into area 2, the traffic is routed along the shortest intra-area path to router Y.

Router Y routes return packets to router X via its nearest Level 1-2 router. The Level 1-2 router recognizes the best route to area 1 via area 4, based on the lowest-cost Level 2 path.

Because Level 1 and Level 2 computations are separate, the path taken from router Y back to router X is not necessarily the least-cost path from router Y to router X.

Asymmetric routing (packets taking different paths in different directions) is not detrimental to the network; however, this type of routing can make troubleshooting difficult and is sometimes a symptom of suboptimal design. Like Enhanced Interior Gateway Routing Protocol (EIGRP) and OSPF, a good IS-IS design is generally hierarchical.

Route Leaking

- **Available since Cisco IOS Software Release 12.0**
- **Helps reduce suboptimal routing by allowing Level 2 information to be leaked into Level 1**
- **Uses up/down bit in Type, Length, and Value (TLV) field**

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—4-14

A feature available since Cisco IOS Software Release 12.0 allows selected Level 2 routes to leak in a controlled manner to Level 1 routers, which helps avoid asymmetric routing.

Route leaking helps reduce suboptimal routing by providing a mechanism for leaking, or redistributing, Level 2 information into Level 1 areas. By having more detail about interarea routes, a Level 1 router is able to make a better choice about which Level 1-2 router to forward the packet.

Route leaking is defined in RFC 2966, *Domain-wide Prefix Distribution with Two-Level IS-IS*, for use with the narrow metric Type, Length and Value (TLV) types 128 and 130. The Internet Engineering Task Force (IETF) has also defined route leaking for use with the wide metric (using TLV type 135).


To implement route leaking, an up/down bit in the TLV is used to indicate whether or not the route identified in the TLV has been leaked. If the up/down bit is set to 0 the route was originated within that Level 1 area.

If the up/down bit is set to 1 the route has been redistributed into the area from Level 2. The up/down bit is used to prevent routing loops: a Level 1-2 router does not readvertise into Level 2 any Level 1 routes that have the up/down bit set.

Route leaking should be planned and deployed carefully to avoid the situation where any topology change in one area results in having to recompute many routes in all other areas.

IS-IS PDUs

This topic describes the OSI PDUs and four types of IS-IS PDUs.



The OSI stack defines a unit of data as a PDU. OSI recognizes a frame as a data-link PDU and a packet (or datagram, in the IP environment) as a network PDU.

The figure shows examples of three types of PDUs (all with IEEE 802.2 Logical Link Control [LLC] encapsulation). IS-IS and ES-IS PDUs are encapsulated directly in a data-link PDU (frame); there is no Connectionless Network Protocol (CLNP) header and no IP header. (In other words, IS-IS and ES-IS do not put routing information in IP or CLNP packets; rather, they put routing information directly in a data link layer frame.)

True CLNP (data) packets contain a full CLNP header between the data-link header and any higher-layer CLNS information.

The IS-IS and ES-IS PDUs contain variable-length fields, depending on the function of the PDU. Each field contains a type code, a length, and the appropriate values; this information is known as the TLVs.

IS-IS PDUs

- **IS-IS PDUs are encapsulated directly into a data-link frame. There is no CLNP or IP header in a PDU.**
- **IS-IS PDUs are as follows:**
 - **Hello (ESH, ISH, IIH)**
 - **LSP**
 - **PSNP (partial sequence number PDU)**
 - **CSNP (complete sequence number PDU)**

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—4-16


IS-IS PDUs are encapsulated directly into an OSI data-link frame. There is no CLNP or IP header.

IS-IS defines four types of PDUs:

- **Hello PDU (ESH, Intermediate System Hello [ISH], IS-IS Hello [IIH]):** Used to establish and maintain adjacencies
- **LSP:** Used to distribute link-state information
- **Partial sequence number PDU (PSNP):** Used to acknowledge and request missing pieces of link-state information
- **Complete sequence number PDU (CSNP):** Used to describe the complete list of LSPs in the LSDB of a router

Link-State Packets

This topic describes how routers use LSPs.



In IS-IS, characteristics of a router are defined by an LSP. The router's LSP contains an LSP header and TLV fields.

- An LSP header includes the following:
 - The PDU type and length
 - The LSP ID
 - The LSP sequence number, used to identify duplicate LSPs and to ensure that the latest LSP information is stored in the topology table
 - The remaining lifetime for the LSP, which is used to age out LSPs
- TLV variable-length fields contain elements including:
 - The neighbor ISs of the router, which are used to build the map of the network
 - The neighbor ESs of the router
 - Authentication information, which is used to secure routing updates
 - Attached IP subnets (optional for Integrated IS-IS)

LSP Header

LSPs are sequenced to prevent duplication of LSPs.

- **Assists with synchronization.**
- **Sequence numbers begin at 1.**
- **Sequence numbers are increased to indicate the newest LSP.**

LSPs in LSDB have a remaining lifetime.

- **Allows synchronization.**
- **Decreasing timer.**

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—4-18

LSPs are given sequence numbers that allow receiving routers to:

- Ensure that they use the latest LSPs in their route calculations.
- Avoid entering duplicate LSPs in the topology tables.

If a router reloads, the sequence number is set to 1. The router then receives its previous LSPs from its neighbors. These LSPs have the last valid sequence number before the router reloaded. The router records this number and reissues its own LSPs with the next-highest sequence number.

Each LSP has a remaining lifetime that is used by the LSP aging process to ensure the removal of outdated and invalid LSPs from the topology table after a suitable time. This process is known as the count to zero operation; 1200 seconds is the default start value.

LSP TLV Examples

TLV	Type Code	Length Field	Value Variable Length
Area address	1	Area ID length + 1	Areas
Intermediate system neighbors	2	Neighbor count + 1	IS neighbors
IP internal reachability	128	Number of connected prefixes	Connected IP prefixes — 4-byte metric, 4-byte prefix, 4-byte mask
IP external reachability	130	Number of redistributed prefixes	Redistributed IP prefixes — 4-byte metric, 4-byte prefix, 4-byte mask

Each set of information, called a “tuple,” includes a type code, length field, and value.

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—4-19

Each LSP includes specific information about networks and stations attached to a router. This information is found in multiple TLV fields that follow the common header of the LSP. TLV is sometimes also referred to as Code, Length, Value (CLV). The TLV structure is a flexible way to add data to the LSP and an easy mechanism for adding new data fields that may be required in the future.

Example: LSP TLV Examples

The figure shows examples of TLVs.

You can find documentation on important TLVs in ISO 10589 and RFC 1195.

Implementing IS-IS in NBMA Networks

This topic explains the different IS-IS network types and the considerations for selecting broadcast or point-to-point modes when implementing IS-IS in ATM and Frame Relay networks.

IS-IS Network Representation

- **Generally, physical links can be placed in the following two groups:**
 - **Broadcast:** Multiaccess subnetworks that support addressing of a group of attached systems
 - **Point-to-point:** Permanent or dynamically established links
- **Only two link-state representations are available in IS-IS:**
 - **Broadcast for LANs and multipoint WANs**
 - **Point-to-point for all other topologies**
- **IS-IS has no concept of NBMA networks.**

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—4-20

Network topologies can be divided into two general types as follows:

- **Point-to-point networks:** Point-to-point links that are either permanently established (leased line, permanent virtual circuit [PVC]) or dynamically established (ISDN, switched virtual circuit [SVC])
- **Broadcast networks:** Multipoint WAN links or LAN links such as Ethernet, Token Ring, or FDDI

IS-IS supports the following two media representations for its link states:

- Broadcast for LANs and multipoint WAN links
- Point-to-point for all other media

Note IS-IS has no concept of nonbroadcast multiaccess (NBMA) networks. It is recommended that you use point-to-point links, such as point-to-point subinterfaces, over NBMA networks such as ATM, Frame Relay, or X.25.

Implementing Network Types in NBMA

When implementing IS-IS in NBMA (such as Frame Relay or ATM):

- **Broadcast mode assumes fully meshed connectivity.**
- **In broadcast mode, you must enable CLNS mapping and include the broadcast keyword, in addition to creating IP maps with the broadcast keyword.**
- **Point-to-point mode is highly recommended (using subinterfaces).**

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0--4-21

Cisco IOS software automatically uses broadcast mode for LAN links and multipoint WAN links. It uses point-to-point mode for point-to-point links, such as point-to-point subinterfaces and dialer interfaces.

There is no specific support for NBMA networks in IS-IS. When implemented in broadcast mode, Cisco IOS software assumes that the NBMA environment features a full mesh of PVCs.

You should use the **broadcast** keyword when creating static maps to map the remote IP address to the local DLCI on a Frame Relay interface because broadcast mode uses multicast updates that will not be sent without this keyword set.

When you use multipoint WAN links like multipoint Frame Relay interfaces, you must also allow CLNS broadcasts and multicasts. This can be done using the command **frame-relay map clns dlci-number broadcast** (in addition to creating the IP maps).

It is highly recommended that you implement NBMA environments, such as Frame Relay, as point-to-point links (using subinterfaces) instead of multipoint links.

Implementing IS-IS in Broadcast Networks

This topic describes the features of broadcast networks.

Broadcast Mode

- **Used for LAN and multipoint WAN interfaces.**
- **Adjacency is recognized through hellos; separate adjacencies for Level 1 and Level 2.**
- **Designated IS (DIS) creates a pseudonode and represents LAN.**
- **DIS for Level 1 and Level 2 may be different.**
- **DIS is elected based on these criteria:**
 - **Only routers with adjacencies are eligible.**
 - **Highest interface priority.**
 - **Highest SNPA (MAC) breaks ties.**
- **There is no backup DIS.**

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—4-22

Broadcast networks are LAN interfaces or multipoint WAN interfaces.

Note	Broadcast mode is recommended for use only on LAN interfaces, although it is also the default for multipoint WANs.

Separate adjacencies are established for Level 1 and Level 2. If two neighboring routers in the same area run both Level 1 and Level 2, they establish two adjacencies, one for each level. The router stores the Level 1 and Level 2 adjacencies in separate Level 1 and Level 2 adjacency tables.

On LANs, routers establish the two adjacencies with specific Layer 1 and Layer 2 IIH PDUs. Routers on a LAN establish adjacencies with all other routers on the LAN (unlike OSPF, where routers establish adjacencies only with the designated router [DR] and backup designated router [BDR]).


IIH PDUs announce the area address. Separate IIH packets announce the Level 1 and Level 2 neighbors. Adjacencies form based on the area address communicated in the incoming IIH and the type of router (Level 1 or Level 2). Level 1 routers accept Level 1 IIH PDUs from their own area and establish adjacencies with other routers in their own area. Level 2 routers (or the Level 2 process within any Level 1-2 router) accept only Level 2 IIH PDUs and establish only Level 2 adjacencies.

Dijkstra's algorithm requires a virtual router (a pseudonode), represented by the DIS, to build a directed graph for broadcast media. Criteria for DIS selection are, first, highest priority (the priority value is configurable) and, second, highest SNPA (on LANs the SNPA is the MAC address).

Cisco router interfaces have a default Level 1 and Level 2 priority of 64. You can configure the priority from 0 to 127 using the **isis priority number-value [level-1 | level-2]** command. The Level 1 DIS and the Level 2 DIS on a LAN may or may not be the same router because an interface can have different Level 1 and Level 2 priorities.

A selected router is not guaranteed to remain the DIS. Any adjacent IS with a higher priority automatically takes over the DIS role. This behavior is called preemptive. Because the IS-IS LSDB is synchronized frequently on a LAN, giving priority to another IS over the DIS is not a significant issue. Unlike OSPF, IS-IS does not use a backup DIS, and routers on a LAN establish adjacencies both with the DIS and with all other routers.

LSP Representing Routers: LAN Representation



In IS-IS, a broadcast link itself is modeled as a pseudonode that connects all attached routers to a star-shaped topology. The pseudonode is represented by the DIS.

Rather than having each router connected to the LAN advertise an adjacency with every other router on the LAN, each router (including the DIS) just advertises a single adjacency to the pseudonode. Otherwise, each IS on a broadcast network with n connected ISs would require $(n)(n - 1) / 2$ adjacency advertisements. Generating LSPs for each adjacency uses considerable overhead in terms of LSDB synchronization.

The DIS generates the pseudonode LSPs. A pseudonode LSP details only the adjacent ISs (for example, the ISs connected to that LAN). The pseudonode LSP is used to build the map of the network and to calculate the shortest path first (SPF) tree. The pseudonode LSP is the equivalent of a network link-state advertisement (LSA) in OSPF.

In IS-IS, all routers on the LAN establish adjacencies with all other routers and with the DIS; therefore, if the DIS fails, another router takes over immediately with little or no impact on the topology of the network. There is no backup DIS. Contrast this with OSPF, where the DR and BDR are selected, and the other routers on the LAN establish full adjacencies with the DR and BDR. In case of DR failure, the BDR is promoted to DR and a new BDR is elected.

LSP and IIH Levels

This topic describes the levels of LSPs and IIHs.

Level 1 and Level 2 LSPs and IIHs

- **The two-level nature of IS-IS requires separate types of LSPs: Level 1 and Level 2 LSPs.**
- **DIS is representative of LAN:**
 - **DIS sends pseudo-Level 1 and pseudo-Level 2 LSPs for LAN.**
 - **Separate DIS for Level 1 and Level 2.**
- **LSPs are sent as unicast on point-to-point networks.**
- **LSPs are sent as multicast on broadcast networks.**
- **LAN uses separate Level 1 and Level 2 IIHs; sent as multicast.**
- **Point-to-point uses a common IIH format; sent as unicast.**

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—4-24

Level 1 and Level 2 LSP

IS-IS uses a two-level area hierarchy. The link-state information for these two levels is distributed separately, which results in Level 1 LSPs and Level 2 LSPs. Each IS originates its own LSPs (one for Level 1 and one for Level 2).

On a LAN, one router (the DIS) sends out LSP information on behalf of the LAN. The DIS represents a pseudonode. The DIS sends out the separate Level 1 or Level 2 LSPs for the pseudonode. The Level 1 DIS and the Level 2 DIS on a LAN may or may not be the same router because an interface can have different Level 1 and Level 2 priorities.

LSPs on point-to-point links are sent as unicast, whereas on broadcast media (LANs) LSPs are sent as multicast.

Level 1 and Level 2 IIH

IIHs are used to establish and maintain neighbor adjacency between ISs. The default hello interval is every 10 seconds; however, the hello interval timer is adjustable.

On a LAN, separate Level 1 and Level 2 IIHs are sent periodically as multicasts to a multicast MAC address. Level 1 announcements are sent to the A11L1IS multicast MAC address 0180.C200.0014, and Level 2 announcements are sent to the A11L2IS multicast MAC address 0180.C200.0015.

The default hello interval for the DIS is three times faster (that is, three times smaller) than the interval for the other routers so that DIS failures can be quickly detected. Unlike the DR and BDR in OSPF, there is no backup DIS in IS-IS.

A neighbor is declared dead if hellos are not received within the hold time. The hold time is calculated as the product of the hello multiplier and hello time. The default hello time is 10 seconds, and the default multiplier is three; therefore, the default hold time is 30 seconds.

Unlike LAN interfaces with separate Level 1 and Level 2 IIHs, point-to-point links have a common point-to-point IIH format that specifies whether the hello relates to Level 1 or Level 2 or both. Point-to-point hellos are sent to the unicast address of the connected router.

Comparing Broadcast and Point-to-Point Topologies

	Broadcast	Point-to-Point
Usage	LAN, full-mesh WAN	PPP, HDLC, partial-mesh WAN
Hello timer	3.3 sec for DIS else 10 sec	10 sec
Adjacencies	$n(n-1) / 2$	$n-1$
Uses DIS	Yes	No
IIH type	Level 1 IIH, Level 2 IIH	Point-to-point IIH

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—4-25

Example: Comparing Broadcast and Point-to-Point Topologies

This figure summarizes the differences between broadcast and point-to-point links.

LSDB Synchronization

This topic describes the types of LSDB synchronization.

LSP Flooding

- **Single procedure for flooding, aging, and updating of LSPs.**
- **Level 1 LSPs are flooded within an area.**
- **Level 2 LSPs are flooded throughout the Level 2 backbone.**
- **Large PDUs are divided into fragments that are independently flooded.**
 - **Each PDU is assigned an LSP fragment number, starting at 0 and incrementing by 1.**
- **Separate LSDBs are maintained for Level 1 and Level 2 LSPs.**

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—4-26

An IS-IS update process is responsible for flooding the LSPs throughout the IS-IS domain. An LSP is typically flooded to all adjacent neighbors except the neighbor from which it was received. Level 1 LSPs are flooded within their local areas. Level 2 LSPs are flooded throughout the backbone.

Each IS originates its own LSP (one for Level 1 and one for Level 2). These LSPs are identified by the system ID of the originator and an LSP fragment number starting at 0. If an LSP exceeds the maximum transmission unit (MTU), it is fragmented into several LSPs, numbered 1, 2, 3, and so on.

IS-IS maintains the Level 1 and Level 2 LSPs in separate LSDBs.

When an IS receives an LSP, it examines the checksum and discards any invalid LSPs, flooding them with an expired lifetime age. If the LSP is valid and newer than what is currently in the LSDB, it is retained, acknowledged, and given a lifetime of 1200 seconds.

The age is decremented every second until it reaches 0, at which point the LSP is considered to have expired. When the LSP has expired, it is kept for an additional 60 seconds before it is flooded as an expired LSP.

LSDB Synchronization

- **SNP packets are used to ensure synchronization and reliability.**
 - **Contents are LSP descriptions**
- **PSNP is used for the following:**
 - **For acknowledgment of LSPs on point-to-point links**
 - **To request missing pieces of LSDB**
- **CSNP is used for the following:**
 - **Periodically by DIS on LAN to ensure LSDB accuracy**
 - **On point-to-point link when the link comes up**

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—4-27

Sequence number PDUs (SNPs) are used to acknowledge the receipt of LSPs and to maintain LSDB synchronization. There are two types of SNPs: CSNP and PSNP. The use of SNPs differs between point-to-point and broadcast media.

CSNPs and PSNPs share the same format; that is, each carries summarized LSP information. The main difference is that CSNPs contain summaries of all LSPs in the LSDB, while PSNPs contain only a subset of LSP entries.

Separate CSNPs and PSNPs are used for Level 1 and Level 2 adjacencies.

Adjacent IS-IS routers exchange CSNPs to compare their LSDB. In broadcast subnetworks, only the DIS transmits CSNPs. All adjacent neighbors compare the LSP summaries received in the CSNP with the contents of their local LSDBs to determine if their LSDBs are synchronized (in other words, if they have the same copies of LSPs as other routers for the appropriate levels and area of routing).


CSNPs are periodically multicast (every 10 seconds) by the DIS on a LAN to ensure LSDB accuracy.

If there are too many LSPs to include in one CSNP, the LSPs are sent in ranges. The CSNP header indicates the starting and ending LSP ID in the range. If all LSPs fit in the CSNP, the range is set to default values.

Adjacent IS-IS routers use PSNPs to acknowledge the receipt of LSPs and to request transmission of missing or newer LSPs.

On point-to-point networks, CSNPs are sent when the link comes up to synchronize the LSDBs.

LSDB Synchronization: LAN



© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—4-28


On a LAN, the DIS periodically (every 10 seconds) sends CSNPs that list the LSPs that it holds in its LSDB. This update is a multicast to all Level 1 or Level 2 IS-IS routers on the LAN.

Example: LSDB Synchronization—LAN

In the example, router R1 compares this list of LSPs with its topology table and realizes that it is missing one LSP. Therefore, it sends a PSNP to the DIS (router R2) to request the missing LSP.

The DIS reissues only that missing LSP (LSP 77), and router R1 acknowledges it with a PSNP.

LSDB Synchronization: Point-to-Point



In contrast to broadcast links such as LAN links, CSNPs are not periodically sent on point-to-point links. A CSNP is sent only once, when the point-to-point link first becomes active. After that, LSPs are sent to describe topology changes, and they are acknowledged with a PSNP.


Example: LSDB Synchronization: Point-to-Point

This figure shows what happens on a point-to-point link when a link failure is detected. The sequence is as follows:

1. A link fails.
2. Router R2 notices this failure and issues a new LSP noting the change.
3. Router R1 receives the LSP, stores it in its topology table, and sends a PSNP back to R2 to acknowledge receipt of the LSP.

LAN Adjacencies

Adjacencies are established based on the area address announced in the incoming IIHs and the type of the router.



© 2006 Cisco Systems, Inc. All rights reserved.


BSCI v3.0—4-30

IIH PDUs announce the area address. On LANs, separate IIH packets announce the Level 1 and Level 2 neighbors.

For example, where a LAN has routers from two areas attached, the following processes apply:

- The routers from one area accept Level 1 IIH PDUs only from their own area and therefore establish adjacencies only with their own area routers.
- The routers from a second area similarly accept Level 1 IIH PDUs only from their own area.
- The Level 2 routers (or the Level 2 process within any Level 1-2 router) accept only Level 2 IIH PDUs and establish only Level 2 adjacencies.

Example: WAN Adjacencies



On point-to-point links (that is, on point-to-point WAN links), the IIH PDUs are common to both levels but announce the level type and the area address in the hellos as follows:

- Level 1 routers in the same area (which includes links between Level 1 and Level 1-2 routers) exchange IIH PDUs that specify Level 1 and establish a Level 1 adjacency.
- Level 2 routers (in the same area or between areas, and including links between Level 2 only and Level 1-2 routers) exchange IIH PDUs that specify Level 2 and establish a Level 2 adjacency.
- Two Level 1-2 routers in the same area establish both Level 1 and Level 2 adjacencies and maintain these with a common IIH PDU that specifies the Level 1 and Level 2 information.

Two Level 1 routers that are physically connected, but that are not in the same area, can exchange IIHs, but they do not establish adjacency because the area addresses do not match.

Example: WAN Adjacencies

The figure illustrates a variety of IS-IS connections. For each pair of routers, indicate the type of adjacency, if any, which will be established, in the space between the two routers. The answers are provided at the end of the Module Self-Check Answer Key.

Summary

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

Summary

- An NSAP is the OSI network-layer address. Cisco uses the area address (comprising the AFI and area ID), system ID, and NSEL fields. The system ID must be 6 bytes.
- A NET address is an NSAP with an NSEL value of 0 and is used to identify the device itself.
- IS-IS defines three types of routers: Level 1, Level 2, and Level 1-2.
- The area address is used to route between areas; the system ID is used to route within an area.
- The four types of IS-IS PDUs are hello, LSP, PSNP, and CSNP.

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—4-32

Summary (Cont.)

- LSPs are used by routers to describe their characteristics. LSPs contain a header and TLV fields. The header ensures unique sequenced packets; the TLV fields include information about the network and stations attached to the router.
- IS-IS recognizes two topology types: point-to-point and broadcast.
- Broadcast networks are LAN interfaces or multipoint WAN interfaces.
- The two-level nature of IS-IS requires separate types of LSPs: Level 1 and Level 2.
- Level 1 LSPs are flooded within an area; Level 2 LSPs are flooded throughout the Level 2 backbone.

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—4-33

Configuring Basic Integrated IS-IS

Overview

Even when Integrated Intermediate System-to-Intermediate System Protocol (IS-IS) is used to support IP exclusively, network devices must still be configured to use the Open Systems Interconnection (OSI) Connectionless Network Service (CLNS) protocol. Each IS-IS router requires a network entity title (NET), and IS-IS packets are directly encapsulated onto the data-link layer instead of traveling inside IP packets.

The commands to configure Integrated IS-IS are slightly different from those of the other IP routing protocols that you have studied, so it is important to understand how to enable IS-IS processes.

Additionally, the default settings for IS-IS can result in the inefficient use of router and network resources and suboptimal routing; therefore, a network administrator also needs to know how to effectively tune IS-IS for optimum performance.

This lesson discusses the mechanics of Integrated IS-IS operation in an IP and CLNS environment and outlines specific commands necessary to implement Integrated IS-IS on a Cisco router.

Objectives

Upon completing this lesson, you will be able to implement Integrated IS-IS in an enterprise network. This ability includes being able to meet these objectives:

- Describe the requirement for CLNS addressing even when using IP in an IS-IS environment
- Describe the configuration process for Integrated IS-IS in an IP environment
- Describe how to optimize IS-IS operation
- Describe how to configure route summarization in IS-IS
- Describe how to verify the IS-IS configuration
- Describe how to verify the CLNS IS-IS structures

Integrated IS-IS in a CLNS Environment

This topic describes the requirement for CLNS addressing even when using IP in an IS-IS environment.

Integrated IS-IS: Requires NET Addresses

- **Common CLNS parameters (NET) and area planning are still required even in an IP environment.**
- **Even when Integrated IS-IS is used for IP routing only, routers still establish CLNS adjacencies and use CLNS packets.**

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—4-2

A NET address identifies a device (an intermediate system [IS] or end system [ES]) and not an interface. This is a critical difference between a NET address and an IP address.

Even if you use Integrated IS-IS only for IP routing, each IS-IS router must have a NET address configured because Integrated IS-IS depends on the support of CLNS routing.

The OSI protocols (hello protocol data units [PDUs]) are used to form the neighbor relationship between routers, and the shortest path first (SPF) calculations rely on a configured NET address to identify the routers.

A device identifies other devices within its own area based on matching area addresses in their NET. It then knows that it can communicate with these other devices without using a default route. A default route is injected into the area by the Level 1-2 router. If the area addresses do not match, then the device knows that it must forward that traffic to its nearest Level 1-2 router.

When you are using IS-IS to route IP traffic, IP subnets are treated as leaf objects associated with IS-IS areas. When routing IP traffic, the router looks up the destination network in its routing table. If the network belongs to a different area, then that traffic must also be forwarded to the nearest Level 1-2 router.

OSI Area Routing: Building an OSI Forwarding Database (Routing Table)

- When databases are synchronized, Dijkstra's algorithm (SPF) is run on the LSDB to calculate the SPF tree.
- The shortest path to the destination is the lowest total sum of metrics.
- Separate route calculations are made for Level 1 and Level 2 routes in Level 1-2 routers.
- Best paths are placed in the OSI forwarding database (CLNS routing table).

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—4-3

IS-IS uses an OSI forwarding database (routing table) to select the best path to a destination. When the databases are synchronized, routers use the link-state database (LSDB) to calculate the SPF tree to OSI destinations, the NETs. The total of the link metrics along each path determines the shortest path to any given destination.

Level 1 and Level 2 routes have separate LSDBs; therefore, routers may run the SPF algorithm twice, once for each level, and create separate SPF trees for each level.

Routers insert the best paths in the CLNS routing table (the OSI forwarding database).

Building an IP Routing Table

Partial route calculation (PRC) is run to calculate IP reachability.

- Because IP and ES are represented as leaf objects, they do not participate in SPF.

Best paths are placed in the IP routing table following IP preferential rules.

- They appear as Level 1 or Level 2 IP routes.

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—4-4

Integrated IS-IS includes IP information in the link-state packets (LSPs), treating it as if it were ES information, leaf connections to the SPF tree. Therefore, updating IP reachability requires only a partial route calculation (PRC), similar to ES reachability in an OSI network.

The PRC generates best-path choices for IP routes and offers the routes to the IP routing table, where they are accepted based on normal IP routing table rules. For example, if there is more than one routing protocol running, then the router compares administrative distance. When the IP IS-IS routes are entered in the routing table, they are shown as via Level 1 or Level 2, as appropriate.

The separation of IP reachability from the core IS-IS network architecture provides Integrated IS-IS better scalability than, for example, Open Shortest Path First Protocol (OSPF)s:

- OSPF sends link-state advertisement (LSAs) for individual IP subnets. If an IP subnet fails, the LSA floods through the network and all routers must run a full SPF calculation, which is extremely CPU-intensive.
- Integrated IS-IS builds the SPF tree from CLNS information. If an IP subnet fails, the IS-IS LSP floods through the network, which is the same for OSPF. However, if this is a leaf (stub) IP subnet (that is, if the loss of the subnet does not affect the underlying CLNS architecture), the SPF tree is unaffected; therefore, only a PRC occurs.

Configuring Integrated IS-IS

This topic describes the configuration process for Integrated IS-IS in an IP environment.

Integrated IS-IS Configuration Steps

1. **Define areas, prepare addressing plan (NETs) for routers, and determine interfaces.**
2. **Enable IS-IS on the router.**
3. **Configure the NET.**
4. **Enable Integrated IS-IS on the appropriate interfaces. Do not forget interfaces to stub IP networks, such as loopback interfaces (although there are no CLNS neighbors there).**

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—4-6

Four steps are required for the basic setup of IS-IS. Additional commands are available for fine-tuning the configuration.

Before you configure Integrated IS-IS, you must map out the areas and plan the addressing. After that is done, you need three commands to enable Integrated IS-IS on a router for IP routing. You can then use additional commands to fine-tune the IS-IS processes.

The table describes the three basic commands used to enable Integrated IS-IS.

Command	Description
router isis [area-tag]	Enables IS-IS as an IP routing protocol and assigns a tag to the process (optional). Given in global configuration mode.
net network-entity-title	Identifies the router for IS-IS by assigning a NET to the router. Given in router configuration mode.
ip router isis [area-tag]	Enables IS-IS on the interfaces that run IS-IS. (This approach is slightly different from most other IP routing protocols, where the interfaces are defined by network statements; there is no network statement under the IS-IS process.) Given in interface configuration mode.

Step 1: Define Area and Addressing

- **Area determined by NET prefix**
 - **Assign to support two-level hierarchy.**
- **Addressing**
 - **IP: Plan to support summarization.**
 - **CLNS: Prefix denotes area. System ID must be unique.**

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—4-6

Recall that all intra-area traffic in IS-IS must traverse the Level 2 backbone area. Thus, CLNS addresses must be planned to execute a two-level hierarchy.

You must decide which routers will be backbone (Level 2) routers, which routers will be Level 1-2, and which will be internal area (Level 1) routers. If some routers must do both Level 1 and Level 2 routing, then you should identify the specific interfaces that will participate in each type of routing.

Remember that the CLNS address of a router is called the NET, and it consists of three main parts:

- The prefix, which identifies the area that the router is a part of
- The system ID, which uniquely identifies each device
- The network service access point (NSAP) selector (NSEL), which must be 0

It is not enough to plan the IS-IS area addressing. You must also plan IP addressing to have a scalable network, and the IP addresses must be planned to allow for summarization of addresses.

Route summarization is the key idea that enables all the benefits of the hierarchical addressing design. Route summarization minimizes routing update traffic and resource utilization.

Be particularly careful when you configure the IP addressing on the router, because it is more difficult to troubleshoot IP address misconfigurations with IS-IS. The IS-IS neighbor relationships are established over OSI CLNS, not over IP. Because of this approach, two ends of a CLNS adjacency can have IP addresses on different subnets, with no impact to the operation of IS-IS.

Step 2: Enable IS-IS on the Router

```
router(config)#
```

```
router isis [area-tag]
```

- **Enables the IS-IS routing protocol**
 - **area-tag**—name for a process
- **When routing of CLNS packets is also needed, use the `clns routing` command.**

© 2006 Cisco Systems, Inc. All rights reserved.BSCI v3.0—4-7

The **router is-is** global configuration command enables Integrated IS-IS on the router. Optionally, you can apply a tag to identify multiple IS-IS processes. Just as multiple OSPF processes can be present on the same router, multiple IS-IS processes are possible.

The process name is significant only to the local router. If it is omitted, Cisco IOS software assumes a tag of 0. If more than one IS-IS process is used, then the network plan should indicate which interfaces would participate in which IS-IS process.

CLNS routing is disabled by default. To enable CLNS routing in addition to IP routing, use the **clns routing** global configuration command. Additionally, you must enable CLNS routing at each interface.

Note	By default, Cisco IOS software makes the router a Level 1-2 router.

Step 3: Configure the NET

```
Router(config-router)#
```

```
net network-entity-title
```

- Configures an IS-IS NET address for the routing process

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—4-8

After the IS-IS process is enabled, the router must be identified for IS-IS by assigning a NET to the router with the **net** command given in router configuration mode.

Even when you use IS-IS for IP routing only (no CLNS routing enabled), you must still configure a NET. The NET is a combination of area number, a unique system identification number for each particular router, and the NSEL of 00 at the end.

The area number must be at least 1 byte in length and can be as long as 13 bytes. The system ID has a fixed length of 6 bytes in Cisco routers. The system ID must be unique throughout each area (Level 1) and throughout the backbone (Level 2).

Step 4: Enable Integrated IS-IS

```
router(config-if)#  
ip router isis [area-tag]
```

- Includes an interface in an IS-IS routing process

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—4-9

The final step is to select which interfaces participate in IS-IS routing. Interfaces that use IS-IS to route IP (and thus must establish IS-IS adjacencies) must be configured using the **ip router isis [area-tag]** interface configuration command. Enable Integrated IS-IS on the appropriate interfaces.

Do not forget interfaces to stub IP networks, such as loopback interfaces (even though there are no CLNS neighbors on those interfaces).

If there is more than one IS-IS process, the IS-IS process to which the interface belongs must be specified using the appropriate process name in the optional area tag field. If no area tag is listed, Cisco IOS software assumes an *area-tag* value of 0. If there is only one IS-IS process active on the router, no *area-tag* value is needed.

Use the **clns router isis [area-tag]** interface command to enable the IS-IS routing process on an interface to support CLNS routing.

Simple Integrated IS-IS Example

The configured router acts as an IP-only Level 1-2 router.

```
interface FastEthernet0/0
ip address 10.1.1.2 255.255.255.0
ip router isis
!
interface Serial 0/0/1
ip address 10.2.2.2 255.255.255.0
ip router isis
!

<output omitted>

router isis
net 49.0001.0000.0000.0002.00
```

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—4-10

Example: Simple Integrated IS-IS Configuration

The figure shows an example of a simple Integrated IS-IS configuration for IP routing only; CLNS routing is not enabled. This configuration specifies only one IS-IS process, thus the optional tag is not used.

The **net** command configures the router to be in area 49.0001 and assigns a system ID of 0000.0000.0002. IS-IS has been enabled on the FastEthernet 0/0 and Serial 0/0/1 interfaces.

Because no level has been configured under the IS-IS routing process, this router acts as a Level 1-2 router by default.

Optimizing IS-IS

This topic describes how to optimize IS-IS operation.

Change IS-IS Router Level

```
Router(config-router)#  
is-type {level-1 | level-1-2 | level-2-only}
```

- **Configures the IS-IS level globally on a router; the default is Level 1-2.**

© 2006 Cisco Systems, Inc. All rights reserved. BSCI v3.0—4-11

Optimizing IS-IS facilitates its smooth functioning and maximizes its efficiency. Remember that the default configuration of IS-IS results in the router having an IS type of Level 1-2. Although this configuration has the advantage of allowing all routers to learn of each other and pass routes without too much administrative oversight, it is not the most efficient way to build an IS-IS network.

Routers with the default configuration send out both Level 1 and Level 2 hellos and maintain both Level 1 and Level 2 LSDBs. Each router should be configured to support the minimum level of routing required, which does the following:

- **Saves memory:** If a router does not need the LSDB for one of the levels, it will not maintain one.
- **Saves bandwidth:** Hellos and LSPs will be sent only for the necessary level.

If a router is to operate only as an internal area router or a backbone router, then specify this configuration by entering the **is-type** router configuration command. To specify that the router act as an internal area (Level 1)-only router, use the **is-type level-1** command. To specify that the router act as a backbone (Level 2)-only router, use the **is-type level-2-only** command.

If the level type has been changed from the default, and the router needs to return to acting as a Level 1-2 router, then use the **is-type level-1-2** command.

Change IS-IS Interface Level

```
Router(config-if)#
```

```
isis circuit-type {level-1 | level-1-2 | level-2-only}
```

- Configures the type of adjacency on an interface; the default is Level 1-2.

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—4-12

Although the router can be a Level 1-2 router, establishing both types of adjacencies over all interfaces may not be required.

If a particular interface has only Level 1 routers connected to it, there is no need for the router to send Level 2 hellos out that interface. Similarly, if an interface has only Level 2 routers connected to it, there is no need for the router to send Level 1 hellos out that interface. Doing so would waste bandwidth and router resources trying to establish adjacencies that do not exist.

To make IS-IS more efficient in these types of situations, configure the interface to send only the needed type of hellos by using the interface command **isis circuit-type** and specifying either the **level-1** or **level-2-only** keyword.

If the circuit type is not configured, by default, Cisco IOS software attempts to establish both types of adjacencies over the interface (Level 1-2).

Change IS-IS Metric

Router(config-if)#

```
isis metric metric [delay-metric [expense-metric [error-  
metric]]] {level-1 | level-2}
```

- Configures the metric for an interface; the default is 10.
- Metric value is from 1 to 63.

Router(config-router)#

```
metric default-value {level-1 | level-2}
```

- Alternately, configures the metric globally for all interfaces

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—4-13

Unlike most other IP protocols, IS-IS on a Cisco router does not take into account line speed or bandwidth when it sets its link metrics. All interfaces are assigned a metric value of 10 by default. In a network with links of varying types and speeds, this assignment can result in suboptimal routing.


To change the metric value, use the **isis metric metric {level-1 | level-2}** interface command. The metric can have different values for Level 1 and Level 2 over the same interface.

Note The IS-IS specification defines four different types of metrics. Cost, being the default metric, is supported by all routers. Delay, expense, and error are optional metrics. The delay metric measures transit delay, the expense metric measures the monetary cost of link utilization, and the error metric measures the residual error probability associated with a link. The default Cisco implementation uses cost only. However, Cisco IOS software does now allow all four metrics to be set with the optional parameters in the **isis metric** command.

If the metric value for all interfaces needs to be changed from the default value of 10, then the change needs to be performed one by one on all IS-IS interfaces. This change can be time-consuming and error-prone, especially for routers with many IS-IS interfaces.

The **metric** command changes the metric value for all IS-IS interfaces. If the keyword **level-1** or **level-2** is not entered, the metric will be applied to both Level 1 and Level 2 IS-IS interfaces. This command is available only in Cisco IOS Software Release 12.3(4)T and later; it supports only the cost metric.

Example: Tuning IS-IS Configuration



Example: Tuning IS-IS Configuration

In the figure, there are two different areas. Area 49.0002 contains only one router (router R3) and needs to do only Level 2 routing. It is appropriate to change the IS type of router R3 to Level 2 only.

Area 49.0001 has two routers. Router R1 is strictly an internal area router; it does not connect to routers in any other area. It is appropriate to configure this router as IS type Level 1. Router R2 connects to the internal area routers and also to router R3, in a different area.

R2 must do both Level 1 and Level 2 routing, so it is left at the default setting; however, there is no need for R2 to send Level 2 hellos out the interface connected to R1. It is appropriate to set the IS-IS circuit type of the FastEthernet 0/0 of R2 to Level 1. Similarly, because the Serial 0/0/1 interface of R2 connects only to a Level 2 router, the IS-IS circuit type should be set to Level 2 only.

Remember that the IS-IS metric for all interfaces is 10. In the topology shown, the serial link is slower than the Fast Ethernet link. Using the default metric does not give the routers a true picture of the value of each link, so the routers cannot make truly informed routing decisions. As shown in the sample configuration, you should change the IS-IS metric at each serial interface to reflect your preference for a link.

Configuring Route Summarization in IS-IS

This topic describes how to configure route summarization in IS-IS.

IP Summarization

```
Router(config-router)#  
summary-address address mask [level-1 | level-2 | level-1-2]  
[tag tag-number] [metric metric-value]
```

- Creates summary
- Default is Level 2

Example

```
P3R1(config-router)# summary-address 10.3.2.0 255.255.254.0 level-1-2
```

- Summarizes 10.3.2.0/23 into Level 1-2

© 2006 Cisco Systems, Inc. All rights reserved. BSCI v3.0—4-15

Routing protocol scalability is a function of the appropriate use of route summarization. An IS can be configured to aggregate a range of IP addresses into a summary address, using the **summary-address** command as shown in the figure.

This command can be used on any router in an IS-IS network. The router summarizes IP routes into Level 1, Level 2, or both. The optional tag number is used to tag the summary route. The optional metric value is applied to the summary route.

The benefits of summarization are as follows:

- Reduced routing table size
- Reduced LSP traffic and protection from flapping routes
- Reduced memory requirements
- Reduced CPU usage

Remove route summarization with the **no** form of the command.

Verifying IS-IS Configuration

This topic describes how to verify the IS-IS configuration.

Example: Is Integrated IS-IS Running?

```
Router#  
show ip protocols
```

```
R2#show ip protocols  
Routing Protocol is "isis"  
  Invalid after 0 seconds, hold down 0, flushed after 0  
  Outgoing update filter list for all interfaces is not set  
  Incoming update filter list for all interfaces is not set  
  Redistributing: isis  
  Address Summarization:  
    None  
  Maximum path: 4  
  Routing for Networks:  
    FastEthernet0/0  
    Loopback0  
    Serial0/0/1  
  Routing Information Sources:  
    Gateway         Distance      Last Update  
    10.10.10.10      115          00:00:02  
    10.30.30.30      115          00:00:03  
  Distance: (default is 115)  
R2#
```

Displays the parameters and current state of the active routing protocol processes

© 2006 Cisco Systems, Inc. All rights reserved. BSCI v3.0—4-16

To verify the IS-IS configuration and IP functionality of the Integrated IS-IS network, use the following commands (these commands can also be useful for troubleshooting problems with the IS-IS network):

- **show ip protocols:** Displays the active IP routing protocols, the interfaces on which they are active, and the networks for which they are routing.
- **show ip route:** Displays the IP routing table. The detail for a particular route or a list of all routes in the routing table from a particular process can be specified.

Example: Is Integrated IS-IS Running?

This sample output from the **show ip protocols** command shows information about IP routing being done over Integrated IS-IS. In this example, IS-IS is running, it is not redistributing any other protocols, and address summarization has not been configured.

The example also shows that interfaces FastEthernet 0/0, Serial 0/0/1, and Loopback 0 are taking part in Integrated IS-IS, that there are two sources of routing information (the neighboring routers), and that the default administrative distance of Integrated IS-IS is 115.

Example: Are There Any IP Routes?

router#

```
show ip route [address [mask]] | [protocol [process-id]]
```

```
R2#show ip route isis
      10.0.0.0/24 is subnetted, 5 subnets
i L2   10.30.30.0 [115/45] via 10.2.2.3, Serial0/0/1
i L1   10.10.10.0 [115/20] via 10.1.1.1, FastEthernet0/0
R2#
```

Displays the current state of the routing table

This sample output from the **show ip route isis** command shows only the IS-IS routes. One route is from Level 1, as indicated by the **i L1** tag, and the other is from Level 2, as indicated by the **i L2** tag.

Integrated IS-IS uses, by default, an administrative distance of 115. The metric shown for each route is taken from the IS-IS cost to the destination. In the figure, for the value of [115/20], 115 is the Integrated IS-IS administrative distance and 20 is the IS-IS metric.

Verifying CLNS IS-IS Structures

This topic describes how to verify CLNS IS-IS structures.

Troubleshooting Commands: CLNS

```
Router#  
show clns
```

- Displays information about the CLNS network

```
Router#  
show clns [area-tag] protocol
```

- Lists the protocol-specific information

```
Router#  
show clns interface [type number]
```

- Lists the CLNS-specific information about each interface

```
Router#  
show clns [area-tag] neighbors [type number] [detail]
```

- Displays both ES and IS neighbors

© 2006 Cisco Systems, Inc. All rights reserved. BSCI v3.0—4-18

You can use the following **show clns** commands to verify the router configuration and to troubleshoot the Integrated IS-IS network:

- **show clns:** This command displays general information about the CLNS network.
- **show clns protocol:** This command displays information for the specific IS-IS processes in the router.
- **show clns interface:** This command displays information about the interfaces that currently run CLNS.
- **show clns neighbors:** This command displays IS and ES neighbors, if there are any. The neighbors are the routers with which this router has IS-IS adjacencies. You can reduce the list of neighbors displayed to those across a particular interface if you specify the interface type and number in the command. The optional keyword **detail** displays the area addresses advertised by the neighbor in the hello messages.

Troubleshooting Commands: CLNS and IS-IS

Router#

```
show isis [area-tag] route
```

- **Displays IS-IS Level 1 routing table (system IDs)**
(requires that CLNS routing be enabled)

Router#

```
show clns route [nsap]
```

- **Displays IS-IS routing table (areas)**

Router#

```
show isis [area-tag] database
```

- **Displays the IS-IS LSDB**

Router#

```
show isis [area-tag] topology
```

- **Displays IS-IS least-cost paths to destinations**

© 2006 Cisco Systems, Inc. All rights reserved.


BSCI v3.0—4-19

You can use the following **show** commands to verify the router configuration and to troubleshoot the Integrated IS-IS network:

- **show isis route:** This command displays the IS-IS Level 1 routing table, which includes all other system IDs in the area. This command is available only if CLNS routing is enabled both globally and at the interface level.
- **show clns route:** This command displays the IS-IS Level 2 routing table, which includes the areas known to this router and the routes to them. Specify a specific address with the optional **nsap** parameter.
- **show isis database:** This command displays the contents of the IS-IS LSDB. To force IS-IS to refresh its LSDB and recalculate all routes, issue the **clear isis** command; an asterisk (*) can be used to clear all IS-IS processes.
- **show isis topology:** This command displays the Level 1 and Level 2 topology tables, which show the least-cost IS-IS paths to the ISs.

Example: OSI Intra-Area and Interarea Routing

Routing in a Two-Level Area Structure



© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—4-20

Example: OSI Intra-Area and Interarea Routing

This figure shows three routers in two areas. Routers R1 and R2 belong to area 49.0001. Router R3 belongs to area 49.0002. R1 is a Level 1 router doing only Level 1 routing. R2 is a Level 1-2 router doing both Level 1 and Level 2 routing. R3 is a Level 2 router doing only Level 2 routing.

This figure forms the basis for the following **show** command examples.

Level 1 and Level 2 Topology Table

```
R1#show isis topology
IS-IS paths to level-1 routers
System Id      Metric  Next-Hop  Interface  SNPA
R1              --
R2             10      R2        Fa0/0      0016.4650.c470

R2#show isis topology
IS-IS paths to level-1 routers
System Id      Metric  Next-Hop  Interface  SNPA
R1             10      R1        Fa0/0      0016.4610.fdb0
R2              --
IS-IS paths to level-2 routers
System Id      Metric  Next-Hop  Interface  SNPA
R1              **
R2              --
R3             35      R3        Se0/0/1    *HDLC*
```

The **show isis topology** command displays the topology databases with the least-cost paths to destination ISs.

Notice that the output for R1 (a Level 1 router) shows the topology database for Level 1 only, and the output for R2 (a Level 1-2 router) shows that separate topology databases exist for Level 1 and Level 2.

The fields in the topology database are common for both levels of routing. They are as follows:

- The system ID shows the system ID from the NET of the destination IS. Cisco IOS software uses dynamic host-name mapping (RFC 2763) to map the system ID to a host name that is available to the router.
- The metric displays the sum of the metrics on the least-cost path to the destination.
- The next-hop column displays the next IS along the path to a destination.
- The interface column shows the output interface that leads to the system listed in the next-hop column.
- The subnetwork point of attachment (SNPA) column contains the OSI Layer 2 address of the next hop. High-Level Data Link Control (HDLC) is shown as the SNPA across an HDLC serial interface. The SNPA on a Fast Ethernet interface is the MAC address. The SNPA would be the data-link connection identifier (DLCI) if it is on a Frame Relay interface.

The topology database on router R1 (a Level 1 router) shows only routers within the local area. R1 is doing only Level 1 routing, and thus does not know of any routers outside its area. Traffic bound for other areas would be forwarded to the nearest router doing Level 2 routing, in this case, router R2.

R2 is doing both levels of routing. It thus maintains two topology databases. The Level 1 database looks very much like the R1 database; only routers within the local area are listed. The Level 2 database is where the external router, R3, shows up.

Simple Troubleshooting: What About CLNS Protocol?

```
R2# show clns protocol

IS-IS Router: <Null Tag>
System Id: 0000.0000.0002.00  IS-Type: level-1-2
Manual area address(es):
  49.0001
Routing for area address(es):
  49.0001
Interfaces supported by IS-IS:
  Loopback0 - IP
  Serial0/0/1 - IP
  FastEthernet0/0 - IP
Redistribute:
  static (on by default)
Distance for L2 CLNS routes: 110
RRR level: none
Generate narrow metrics: level-1-2
Accept narrow metrics: level-1-2
Generate wide metrics: none
Accept wide metrics: none
```

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—4-22

In the figure, the example output from the **show clns protocol** command shows this information:

- The Integrated IS-IS process, its tag, if present
- The system ID and area address for this router
- The IS level types for the router
- The interfaces using Integrated IS-IS for routing, including whether they are routing for IP, CLNS, or both
- Any redistribution of other route sources
- Information about the distances for Level 2 CLNS routes and the acceptance and generation of metrics

Are Adjacencies Established?

```
R2# show clns neighbors
System Id      Interface  SNPA          State  Holdtime  Type  Protocol
R3             Se0/0/1    *HDLC*        Up     28        L2    IS-IS
R1             Fa0/0      0016.4610.fdb0 Up     23        L1    IS-IS

R2#show clns interface s0/0/1
Serial0/0/1 is up, line protocol is up
Checksums enabled, MTU 1500, Encapsulation HDLC
ERPDUs enabled, min. interval 10 msec.
CLNS fast switching enabled
CLNS SSE switching disabled
DEC compatibility mode OFF for this interface
Next ESH/ISH in 45 seconds
Routing Protocol: IS-IS
  Circuit Type: level-2
  Interface number 0x1, local circuit ID 0x100
  Neighbor System-ID: R3
  Level-2 Metric: 35, Priority: 64, Circuit ID: R2.00
  Level-2 IPv6 Metric: 10
  Number of active level-2 adjacencies: 1
  Next IS-IS Hello in 5 seconds
  if state UP
```

In the figure, the example output from the **show clns neighbors** command shows this information:

- The IS-IS neighbors
- The SNPAs and state
- The hold time, which is the timeout for receipt of no hellos, after which the neighbor is declared down
- The neighbor level type

Also in this figure, the example of output from the **show clns interface** command shows this information:

- The interface runs IS-IS and attempts to establish Level 2 adjacencies.
- The interface numbers and circuit ID for IS-IS purposes.
- The ID of the neighbor.
- The metric or metrics for the interface.
- The priority for Designated IS (DIS) negotiation. Priority is not relevant in this case because it is a serial HDLC interface.
- The information regarding hello timers and the number of established adjacencies.
- The state of the interface.

Summary

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

Summary

- **Integrated IS-IS routing for IP uses CLNS and therefore requires CLNS addresses, that is, NET addresses.**
- **Integrated IS-IS requires planning the addresses, enabling the router, defining the router NET, and enabling the appropriate interfaces.**
- **IS-IS can be optimized by adjusting adjacency levels and changing the default metric cost.**
- **IS-IS summarization can be configured with the summary-address command.**
- **The show ip protocols and show ip route commands verify the IS-IS configuration and IP functionality.**
- **Various show commands are used to troubleshoot CLNS IS-IS structures and Integrated IS-IS networks**

© 2006 Cisco Systems, Inc. All rights reserved.

BSCI v3.0—4-24

Module Summary

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

Module Summary

- **IS-IS is a proven and extensible IP routing protocol that converges quickly and supports VLSM.**
- **Unlike IP addresses, CLNS addresses apply to entire nodes and not to interfaces. IS-IS runs directly on the data-link layer and does not use IP or CLNS as a network protocol.**
- **Even when IS-IS is installed to support IP exclusively, network devices must also be configured with NET addresses. Default settings for IS-IS may result in the inefficient use of router and network resources and suboptimal routing.**

© 2006 Cisco Systems, Inc. All rights reserved. BSCI v3.0—4-1

Connectionless Network Service (CLNS) provides connectionless delivery of data. As a result, CLNS is the solution for unreliable delivery of data, similar to IP. Intermediate System to Intermediate System (IS-IS) operates in strictly CLNS terms, although Integrated IS-IS supports IP routing as well as CLNS. IS-IS supports various data-link environments, such as Ethernet and Frame Relay.

References

For additional information, refer to these resources:

- RFC 995, *End System to Intermediate System Routing Exchange Protocol for Use in Conjunction with ISO 8437*.
- RFC 1195, *Use of OSI IS-IS for Routing in TCP/IP and Dual Environments*.
- Martey, A. *IS-IS Network Design Solutions*. Indianapolis, Indiana: Cisco Press; 2002.

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 protocols does Integrated IS-IS support? (Choose two.) (Source: Introducing IS-IS and Integrated IS-IS Routing)
- A) IP
 - B) IPX
 - C) OSI CLNS
 - D) AppleTalk
- Q2) Which two characteristics describe IS-IS? (Choose two.) (Source: Introducing IS-IS and Integrated IS-IS Routing)
- A) It is an IGP.
 - B) It is an EGP.
 - C) It is efficient in its use of network resources.
 - D) It is an advanced distance vector routing protocol.
- Q3) Which two routing levels does IS-IS support? (Choose two.) (Source: Introducing IS-IS and Integrated IS-IS Routing)
- A) Level 1 for IP
 - B) Level 1 for interarea routing
 - C) Level 2 for interarea routing
 - D) Level 2 for CLNS
 - E) Level 0 for interdomain routing
 - F) Level 1 for intra-area routing
 - G) Level 2 for intra-area routing
- Q4) Which of the following is ES-IS responsible for? (Source: Introducing IS-IS and Integrated IS-IS Routing)
- A) Level 0 routing
 - B) Level 1 routing
 - C) Level 2 routing
 - D) Level 3 routing
- Q5) Which two of the following are IS-IS responsible for? (Choose two.) (Source: Introducing IS-IS and Integrated IS-IS Routing)
- A) Level 0 routing
 - B) Level 1 routing
 - C) Level 2 routing
 - D) Level 3 routing
- Q6) Routing between areas is described as _____. (Source: Introducing IS-IS Routing and Integrated IS-IS Routing)
- A) Level 0 routing
 - B) Level 1 routing
 - C) Level 2 routing
 - D) Level 3 routing

- Q7) Check the characteristics that can be attributed to IS-IS and those that can be attributed to OSPF. Some characteristics may apply to both. (Source: Introducing IS-IS and Integrated IS-IS Routing)

	IS-IS	OSPF
Link-state protocol		
Fast convergence		
Supports VLSM		
More extensible		
Documentation and experienced engineers easy to find		
Most customized to IP		
Metrics scale automatically		

- Q8) Level 1 routing is responsible for which task? (Source: Introducing IS-IS and Integrated IS-IS Routing)

- A) exchanging information about paths between areas
- B) building topology of ESs and ISs in areas
- C) using ES-IS to learn prefix information
- D) CLNP routing

- Q9) Level 2 routing is responsible for which task? (Source: Introducing IS-IS and Integrated IS-IS Routing)

- A) exchanging information about paths between areas
- B) building topology of ESs and ISs in areas
- C) using ES-IS to learn prefix information
- D) CLNP routing

- Q10) A good IS-IS design features which two properties? (Choose two.) (Source: Introducing IS-IS and Integrated IS-IS Routing)

- A) CLNS addresses confined to Level 1
- B) a summarizable address plan
- C) two-level hierarchy
- D) routers with minimal memory and CPU

- Q11) Which two characteristics describe IS-IS advantages over OSPF? (Choose two.) (Source: Introducing IS-IS and Integrated IS-IS Routing)

- A) better support from the IETF
- B) more documentation and support
- C) ubiquitously implemented
- D) support for CLNS
- E) more extensible
- F) faster convergence

- Q12) Put the parts of an NSAP address in the correct order. (Source: Performing IS-IS Routing Operations)
- A) system ID
 - B) HO-DSP
 - C) AFI
 - D) NSEL
 - E) IDI
- Q13) Which component does a NET address identify? (Source: Performing IS-IS Routing Operations)
- A) interface
 - B) device
 - C) process
 - D) protocol
- Q14) Which characteristic describes a NET? (Source: Performing IS-IS Routing Operations)
- A) always has an AFI of 49
 - B) always has a 3-byte HO-DSP
 - C) always has an NSEL of 00
 - D) is never assigned to a router
- Q15) Which prefix does private AFI use? (Source: Performing IS-IS Routing Operations)
- A) 49.0001
 - B) 37
 - C) 47
 - D) 49
 - E) 37.1921
- Q16) Which term describes an NSAP address? (Source: Performing IS-IS Routing Operations)
- F) NSCAR address
 - G) CLNS address
 - H) protocol-specific port
 - I) replacement for BGP
- Q17) Which two statements are true with reference to CLNS? (Choose two.) (Source: Performing IS-IS Routing Operations)
- A) Within an area, all system IDs must be unique.
 - B) Within an area, all area addresses (AFI, IDI, HO-DSP) must be identical.
 - C) Within an area, AFI and IDI must be identical, but HO-DSP may vary.
 - D) Within an area, all IP addresses must be in the same classful network.
- Q18) Which two terms are SNPAs? (Choose two.) (Source: Performing IS-IS Routing Operations)
- A) MAC address
 - B) Frame Relay DLCI
 - C) IP address
 - D) CLNS address including NSEL
 - E) delivery point for traffic

- Q19) Which three conditions are system ID requirements? (Choose three.) (Source: Performing IS-IS Routing Operations)
- A) must be 8 bytes
 - B) must be unique in an area
 - C) must be classless
 - D) must be unique within Level 2 routers in a routing domain
 - E) must be 6 bytes on a Cisco router
- Q20) Which two statements are parts of OSI routing logic?
A Level 1-2 router should compare the destination area address to its own area address and if they are _____ (Choose two.) (Source: Performing IS-IS Routing Operations)
- A) the same, route at Level 1.
 - B) different, route at Level 1.
 - C) the same, route at Level 2.
 - D) different, route at Level 2.
- Q21) Route leaking allows Level 1 information to be leaked into Level 2. (Source: Performing IS-IS Routing Operations)
- A) true
 - B) false
- Q22) Adjacencies are formed in IS-IS between _____. (Choose two.) (Source: Performing IS-IS Routing Operations)
- A) ISs in the same area
 - B) pseudonodes
 - C) two ISs doing routing at the same level (Level 1 or Level 2)
 - D) interfaces on a broadcast network
- Q23) How are IS-IS PDUs transported? (Source: Performing IS-IS Routing Operations)
- A) within IP packets
 - B) within CLNS packets
 - C) directly within frames
 - D) reliably using TCP
- Q24) What are the four types of IS-IS PDUs? (Choose four.) (Source: Performing IS-IS Routing Operations)
- A) hello
 - B) SNAP
 - C) LSP
 - D) PSNP
 - E) CSNP
 - F) IP/IPX
 - G) OSPF
- Q25) What does TLV stand for? (Source: Performing IS-IS Routing Operations)
- A) Time, Level, Value
 - B) Text, Level, Volume
 - C) Time, Length, Volume
 - D) Type, Length, Value

- Q26) Which three topology types are supported by IS-IS? (Choose three.) (Source: Performing IS-IS Routing Operations)
- A) broadcast for LAN links
 - B) broadcast for multipoint WAN links
 - C) stub for networks with a single exit
 - D) point-to-point for LAN links
 - E) point-to-point for point-to-point WAN links
 - F) NBMA for WAN links
- Q27) How does IS-IS deal with DIS failures? (Source: Performing IS-IS Routing Operations)
- A) The Backup IS (BIS) asserts itself as the new DIS.
 - B) The failure is recognized quickly and a new DIS is elected. Because IS-IS synchronizes LSDBs frequently on a LAN, this is rarely a problem.
 - C) The network reverts to general topology, and every IS forms an adjacency with every other IS.
 - D) It does not. The DIS must be restored quickly to prevent communication problems.
- Q28) How does IS-IS adapt the directed graph approach (implicit in Dijkstra's algorithm) for use in a multipoint environment such as Ethernet? (Source: Performing IS-IS Routing Operations)
- A) It uses a pseudonode.
 - B) It uses a system ID.
 - C) The algorithm is amended.
 - D) Dijkstra's algorithm is used only on point-to-point links.
- Q29) Which four technologies are appropriate for point-to-point topology? (Choose four.) (Source: Performing IS-IS Routing Operations)
- A) ATM PVC
 - B) leased line
 - C) Ethernet
 - D) Frame Relay
 - E) Token Ring
 - F) ISDN Dial-on-Demand
- Q30) What is the function of CSNPs? (Source: Performing IS-IS Routing Operations)
- A) to request a missing LSP
 - B) to acknowledge an LSP
 - C) to provide alerts
 - D) to maintain LSDB synchronization
- Q31) Which address component is used to identify the router in an IS-IS environment? (Source: Performing IS-IS Routing Operations)
- A) SNPA
 - B) NRLI
 - C) NET
 - D) system number

- Q32) What is an SPNA? (Source: Performing IS-IS Routing Operations)
- A) the OSI specification for SPF
 - B) a data-link address
 - C) level 2 routing
 - D) the interdomain component of IS-IS, not supported by Cisco routers
- Q33) Which two network representations are supported by IS-IS? (Choose two.) (Source: Performing IS-IS Routing Operations)
- A) broadcast
 - B) NBMA
 - C) stub
 - D) point-to-point
 - E) NSSA
- Q34) Which communication method is used by devices in two Level 1 areas? (Source: Performing IS-IS Routing Operations)
- A) tunnel
 - B) Level 1 runs on broadcast multiaccess networks
 - C) through Level 2
 - D) using PDUs
- Q35) What does the network entity title identify in IP networks? (Source: Configuring Basic Integrated IS-IS)
- A) a router interface
 - B) a router process
 - C) a router
 - D) a subnet
- Q36) When does Dijkstra's algorithm run to determine the best path? (Source: Configuring Basic Integrated IS-IS)
- A) after the PRC process is completed
 - B) when maximum age is reached
 - C) after databases are synchronized
 - D) at IS-IS startup
- Q37) Which two criteria does a router running Integrated IS-IS use to determine which routes to place in the IP routing (forwarding) table? (Choose two.) (Source: Configuring Basic Integrated IS-IS)
- A) area addresses
 - B) administrative distance
 - C) information from a PRC
 - D) Level 2 routes only
- Q38) Which command displays the IS-IS Level 1 routing table? (Source: Configuring Basic Integrated IS-IS)
- A) **show clns route**
 - B) **show ip neighbor**
 - C) **show isis route**
 - D) **which-route**
 - E) **show clns neighbor**

- Q39) Which command provides a list of all IS-IS areas known to a router? (Source: Configuring Basic Integrated IS-IS)
- A) **show clns route**
 - B) **show ip route**
 - C) **which-route**
 - D) **show isis route**
- Q40) A Level 1-2 IS with the NET 49.000A.0000.0C12.3456.00 receives traffic going to 49.001A.0000.0C78.9AB.00. Which table does it use to route the traffic? (Source: Configuring Basic Integrated IS-IS)
- A) IS-IS topology
 - B) Level 1 routing table
 - C) Level 2 routing table
 - D) CLNS routing table
- Q41) Which command reveals the redistribution processes in IS-IS? (Source: Configuring Basic Integrated IS-IS)
- A) **show clns protocol**
 - B) **show clns interface**
 - C) **show isis route**
 - D) **show isis database**
- Q42) Which command would you use to view the router interfaces that are currently running CLNS? (Source: Configuring Basic Integrated IS-IS)
- A) **show clns protocol**
 - B) **show clns interface**
 - C) **show isis route**
 - D) **show isis database**
- Q43) Which address does Integrated IS-IS require? (Source: Configuring Basic Integrated IS-IS)
- A) IP address to use as router ID
 - B) IP addresses on interfaces
 - C) CLNS address to identify the device
 - D) CLNS addresses on interfaces
- Q44) Which two characteristics are true about Dijkstra calculations in IS-IS? (Choose two.) (Source: Configuring Basic Integrated IS-IS)
- A) performs on separate Level 1 and Level 2 databases
 - B) performs on TLV
 - C) uses shortest path—lowest sum of link metrics
 - D) can be disabled
- Q45) Routing from 49.BAD1.1921.6800.0111.00 to 49.BAD7.1921.6800.0112.00 takes place at what level? (Source: Configuring Basic Integrated IS-IS)
- A) Level 1
 - B) Level 2
 - C) not enough information to determine
 - D) routing not possible between these networks

- Q46) Place the IS-IS configuration steps in the correct order. (Source: Configuring Basic Integrated IS-IS)
- A) _____ configure NET
 - B) _____ enable IS-IS on the router
 - C) _____ enable IS-IS on the interfaces
 - D) _____ define areas and addressing
- Q47) The router identifies which interfaces participate in IS-IS routing by the _____. (Source: Configuring Basic Integrated IS-IS)
- A) **network** command
 - B) NET that is configured for each interface
 - C) **ip router isis** command
 - D) NET that is configured for each router
- Q48) IS-IS summarization allows you to _____. (Source: Configuring Basic Integrated IS-IS)
- A) summarize the list of NETs
 - B) summarize the area address
 - C) summarize a set of IP addresses into a less specific address
 - D) enumerate the IS and ES neighbors
- Q49) Which command shows the sources of Integrated IS-IS routing information? (Source: Configuring Basic Integrated IS-IS)
- A) **show ip protocols**
 - B) **show clns protocols**
 - C) **show ip route isis**
 - D) **show ip route**
- Q50) What is the default IS-IS routing level of a Cisco router? (Source: Configuring Basic Integrated IS-IS)
- A) 0
 - B) Level 1
 - C) Level 2
 - D) 3
 - E) Level 1-2
- Q51) A NET address is required to configure Integrated IS-IS for routing IP only. (Source: Configuring Basic Integrated IS-IS)
- A) true
 - B) false
- Q52) To configure IS-IS on an interface, which command must be executed from interface configuration mode? (Source: Configuring Basic Integrated IS-IS)
- A) **router isis**
 - B) **isis interface**
 - C) **ip router isis**
 - D) **is-type level-1**

- Q53) What is the default IS-IS metric for Fast Ethernet interfaces? (Source: Configuring Basic Integrated IS-IS)
- A) 10
 - B) 16
 - C) 83
 - D) 100
- Q54) IP routing scalability is achieved by _____. (Source: Configuring Basic Integrated IS-IS)
- A) NET assignment
 - B) route summarization
 - C) controlling ES-IS
 - D) limiting IS-IS resynchronization issues

Module Self-Check Answer Key


- Q1) A, C
Q2) A, C
Q3) C, F
Q4) A
Q5) B, C
Q6) C
Q7)

	IS-IS	OSPF
Link-state protocol	X	X
Fast convergence	X	X
Supports VLSM	X	X
More extensible	X	
Documentation and experienced engineers easy to find		X
Most customized to IP		X
Metrics scale automatically		X

- Q8) B
Q9) A
Q10) B, C
Q11) D, E
Q12) 1 = C, 2 = E, 3 = B, 4 = A, 5 = D
Q13) B
Q14) C
Q15) D
Q16) B
Q17) A, B
Q18) A, B
Q19) B, D, E
Q20) A, D
Q21) B
Q22) A, C
Q23) C
Q24) A, C, D, E

- Q25) D
- Q26) A, B, E
- Q27) B
- Q28) A
- Q29) A, B, D, F
- Q30) D
- Q31) C
- Q32) B
- Q33) A, D
- Q34) C
- Q35) C
- Q36) C
- Q37) B, C
- Q38) C
- Q39) A
- Q40) C
- Q41) A
- Q42) B
- Q43) C
- Q44) A, C
- Q45) B
- Q46) 1 = D, 2 = B, 3 = A, 4 = C
- Q47) C
- Q48) C
- Q49) A
- Q50) E
- Q51) A
- Q52) C
- Q53) A
- Q54) B

Example: WAN Adjacencies



The figure shows the answers to the WAN adjacencies example at the end of the “Performing IS-IS Routing Operations” lesson of this module.