

Network Troubleshooting Video Mentor

Kevin Wallace, CCIE® No. 7945

Cisco Press

800 East 96th Street

Indianapolis, Indiana 46240 USA

Network Troubleshooting Video Mentor

Kevin Wallace

Copyright© 2010 Pearson Education, Inc.

Published by:

Cisco Press

800 East 96th Street

Indianapolis, IN 46240 USA

All rights reserved. No part of this book may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or by any information storage and retrieval system, without written permission from the publisher, except for the inclusion of brief quotations in a review.

Printed in the United States of America

1 2 3 4 5 6 7 8 9 0

First Printing April 2010

ISBN-13: 978-1-58720-296-4

ISBN-10: 1-58720-296-4

Warning and Disclaimer

This book and video product is designed to provide information to help you prepare for the CCNP Exams. Every effort has been made to make this book as complete and as accurate as possible, but no warranty or fitness is implied.

The information is provided on an “as is” basis. The authors, Cisco Press, and Cisco Systems, Inc. shall have neither liability nor responsibility to any person or entity with respect to any loss or damages arising from the information contained in this book or from the use of the discs or programs that may accompany it.

The opinions expressed in this book belong to the author and are not necessarily those of Cisco Systems, Inc.

Publisher

Paul Boger

Associate Publisher

Dave Dusthimer

Business Operation Manager

Anand Sundaram

Cisco Press Manager Global Certification

Erik Ullanderson

Executive Editor

Brett Bartow

Managing Editor

Patrick Kanouse

Development Editor

Dayna Isley

Project Editor

Jennifer Gallant

Copy Editor

Water Crest Publishing, Inc.

Technical Editor

Michelle Plumb

Team Coordinator

Vanessa Evans

Cover Designer

Gary Adair

Composition

Mark Shirar

Trademark Acknowledgments

All terms mentioned in this book that are known to be trademarks or service marks have been appropriately capitalized. Cisco Press or Cisco Systems, Inc., cannot attest to the accuracy of this information. Use of a term in this book should not be regarded as affecting the validity of any trademark or service mark.

Corporate and Government Sales

The publisher offers excellent discounts on this book when ordered in quantity for bulk purchases or special sales, which may include electronic versions and/or custom covers and content particular to your business, training goals, marketing focus, and branding interests. For more information, please contact:

U.S. Corporate and Government Sales

1-800-382-3419

corpsales@pearsontechgroup.com

For sales outside the United States, please contact:

International Sales

international@pearsoned.com

Feedback Information

At Cisco Press, our goal is to create in-depth technical books of the highest quality and value. Each book is crafted with care and precision, undergoing rigorous development that involves the unique expertise of members from the professional technical community.

Readers' feedback is a natural continuation of this process. If you have any comments regarding how we could improve the quality of this book, or otherwise alter it to better suit your needs, you can contact us through email at feedback@ciscopress.com. Please make sure to include the book title and ISBN in your message.

We greatly appreciate your assistance.



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

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

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

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

CCDE, CCENT, Cisco Eos, Cisco HealthPresence, the Cisco logo, Cisco Lumin, Cisco Nexus, Cisco StadiumVision, Cisco TelePresence, Cisco WebEx, DCE, and Welcome to the Human Network are trademarks. Changing the Way We Work, Live, Play, and Learn and Cisco Store are service marks; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQuick Study, IronPort, the IronPort logo, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0812R)

About the Author

Kevin Wallace, CCIE No. 7945, is a certified Cisco instructor, and he holds multiple Cisco certifications including the CCSP, CCVP, CCNP, and CCDP, in addition to multiple security and voice specializations. With Cisco experience dating back to 1989 (beginning with a Cisco AGS+ running Cisco IOS 7.x), Kevin has been a network design specialist for the Walt Disney World Resort, a senior technical instructor for SkillsSoft/Thomson NETg/KnowledgeNet, and a network manager for Eastern Kentucky University. Kevin holds a bachelor's of science degree in electrical engineering from the University of Kentucky. Also, Kevin has authored multiple books and video products for Cisco Press, including the *Routing Video Mentor* and *CCNP TSHOOT 642-832 Official Certification Guide*, both of which target the current CCNP certification. Kevin lives in central Kentucky with his wife (Vivian) and two daughters (Stacie and Sabrina).

About the Technical Reviewer

Michelle Plumb is a full-time Cisco-certified instructor for Skillssoft. Michelle has more than 19 years experience in the field as an IT professional and telephony specialist. She maintains a high level of Cisco and Microsoft certifications. Michelle has been a technical reviewer for numerous books related to the Cisco CCNP and CCVP course material track. Michelle currently lives in Scottsdale, Arizona, with her husband and two dogs.

Dedication

I dedicate this *Video Mentor* product to my aunt and uncle, Louise and Lawrence Pierce. They enthusiastically support my writing efforts, and they think of my daughters as their own grandchildren. Thanks for all your love and support!

Acknowledgments

My thanks go out to the team of professionals at Cisco Press. I'm honored to be associated with you. Also, since this *Video Mentor* addresses troubleshooting, I'm reminded of many network problems I've faced over the past two decades. Without working through those challenges, I would not have developed my troubleshooting muscle. So, I'm actually grateful for those issues (well, at least some of them) and the user interactions I had along the way.

On a personal note, I want to acknowledge God for His blessings and guidance in my life. And of course, this *Video Mentor* would not have been possible without my family's support. My wife, Vivian, is always understanding of my demanding writing schedule, and my daughters, Sabrina and Stacie, are continually encouraging.

Contents at a Glance

Lab 1	Spanning Tree Troubleshooting	1
Lab 2	Router Redundancy Troubleshooting	9
Lab 3	EIGRP Troubleshooting	17
Lab 4	OSPF Troubleshooting	27
Lab 5	Route Redistribution Troubleshooting	39
Lab 6	BGP Troubleshooting	49
Lab 7	IPv6 and OSPFv3 Troubleshooting	59
Lab 8	IPv6 and RIPng Troubleshooting	71
Lab 9	Cisco IOS Security Troubleshooting	85
Lab 10	DHCP Troubleshooting	93
Lab 11	NAT Troubleshooting	101
Lab 12	VoIP Troubleshooting	109

Icons Used in This Book



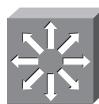
Router



Switch



Hub

Multilayer
SwitchCisco
IP Phone
(CallManager)

Phone



PC



Modem

File
ServerVoice-Enabled
RouterVoice-Enabled
Workgroup Switch

Call Agent



Cisco ASA

Command Syntax Conventions

The conventions used to present command syntax in this book are the same conventions used in the IOS Command Reference. The Command Reference describes these conventions as follows:

- **Boldface** indicates commands and keywords that are entered literally as shown. In actual configuration examples and output (not general command syntax), boldface indicates commands that are manually input by the user (such as a show command).
- *Italics* indicate arguments for which you supply actual values.
- Vertical bars (|) separate alternative, mutually exclusive elements.
- Square brackets [] indicate optional elements.
- Braces { } indicate a required choice.
- Braces within brackets [{ }] indicate a required choice within an optional element.

Introduction

The *Network Troubleshooting Video Mentor* helps those who troubleshoot or maintain Cisco routers or switches. Additionally, this *Video Mentor* product helps prepare candidates of the TSHOOT Cisco exam (642-832). Each *Network Troubleshooting Video Mentor* video presents a unique lab scenario, with both visual references and audio explanations of what you should expect to happen in a particular lab. The videos also show the command-line interface (CLI) commands used to implement the features described in each lab video, along with running commentary. The end result is a set of lab videos that explain some of the most important troubleshooting topics, with thorough explanations from a trusted mentor.

The *Network Troubleshooting Video Mentor* product was created out of a need for something more than just the static written word of a typical book. Cisco Press already offers many books that cover the wide breadth of troubleshooting strategies. However, many people learn better in a classroom setting, with an instructor explaining the concepts while showing details as projected on the wall. Many customers of Cisco Press book products asked for a product closer to what you might get in a class, and the *Video Mentor* series of products is the result.

Goals and Methods

The *Network Troubleshooting Video Mentor* has a very specific set of goals. First, this product seeks to help its viewers more effectively troubleshoot a wide range of router and switch configurations. Although you might have already read about troubleshooting scenarios in other books, or heard about them in classes, the *Network Troubleshooting Video Mentor* hopes to help you master these particular troubleshooting concepts. Using the *Network Troubleshooting Video Mentor* in addition to a book or attending a course should help solidify your knowledge, help you see how to apply the knowledge, and better prepare you for the application of knowledge in real-world deployments or on a Cisco exam.

Note that the *Network Troubleshooting Video Mentor* does not attempt to cover all possible troubleshooting scenarios of each technology discussed. Rather, it seeks to introduce the fundamentals of each technology and then address one or more troubleshooting issues commonly encountered with that technology.

Each of the 12 lab videos follows the same basic approach, including these basic steps:

1. The lab scenario steps are listed, giving a general outline of what the viewer should expect to see and hear during the video.
2. The lab reviews the fundamental concepts of the technology presented in each lab's trouble ticket.
3. A syntax reference is provided.
4. The lab topology used in the video is described.
5. The video shows the lab topology along with a router/switch CLI as Kevin walks you through the resolution of the identified issue.

Network Troubleshooting Video Mentor Contents

The *Network Troubleshooting Video Mentor* package contains a DVD with 12 lab videos and PDF instructions for each lab. The DVD has been optimized for viewing on a computer with a 1024×768 minimum pixel grid. When the DVD starts, it will display a menu, from which you can start one of the 12 lab videos.

The PDF files are intended to be used for reference when watching the videos, as opposed to being a standalone tool. The PDF files have a section corresponding to each of the 12 *Network Troubleshooting Video Mentor* labs and include the following:

- The list of objectives for the lab
- A review of the technology being addressed in the lab
- A list of Cisco IOS commands useful for configuring and troubleshooting the lab scenario
- The topology used in the lab
- Common troubleshooting targets encountered with a specific technology
- Troubleshooting steps (including verification steps) performed during the lab
- A summary of the lab

Who Should Use the Network Troubleshooting Video Mentor?

The *Network Troubleshooting Video Mentor* is primarily intended for people using self-study books as their primary method of preparing to troubleshoot Cisco routers, or to pass the TSHOOT Cisco exam (642-832). Additionally, this product should be useful to anyone who is studying troubleshooting topics, either by reading books or when taking classes.

How the Network Troubleshooting Video Mentor Is Organized

The *Network Troubleshooting Video Mentor* DVD menu enables you to have access to each of the 12 labs. The menu also gives you access to a PDF of the booklet included with the DVD.

The booklet itself simply contains 12 sections, each referencing one of the 12 lab videos. The 12 lab videos are as follows:

- **Lab 1: Spanning Tree Troubleshooting**—This lab reviews the operation and configuration of Spanning Tree Protocol (STP). You are also challenged to resolve an STP-related performance issue.
- **Lab 2: Router Redundancy Troubleshooting**—This lab contrasts three approaches to providing first-hop router redundancy: HSRP, VRRP, and GLBP. However, the primary focus of this lab is HSRP. You are given a collection of HSRP configuration and troubleshooting commands. The lab then walks you through the resolution of an HSRP trouble ticket.
- **Lab 3: EIGRP Troubleshooting**—This lab discusses common EIGRP troubleshooting targets and reviews a collection of common EIGRP commands. The trouble ticket presented in this lab involves EIGRP's load-balancing behavior.

-
- **Lab 4: OSPF Troubleshooting**—This lab introduces OSPF, describes the function of a designated router, and contrasts OSPF network types. The lab topology is configured for three OSPF areas, and the trouble ticket indicates that OSPF adjacencies are not being properly formed. This lab then walks you through the resolution of multiple OSPF configuration issues.
 - **Lab 5: Route Redistribution Troubleshooting**—This lab discusses the theory and configuration of route redistribution. Common route redistribution troubleshooting issues are considered, and you are presented with a route redistribution trouble ticket where OSPF and EIGRP need to be mutually redistributed.
 - **Lab 6: BGP Troubleshooting**—This lab highlights the characteristics of BGP and discusses BGP’s path selection criteria. This lab then explains how BGP routing decisions can be influenced by manipulating the Local Preference and ASPTH BGP attributes. The trouble ticket addresses a dual-homed BGP topology, where BGP is selecting a suboptimal path for its communication between an enterprise network and the Internet.
 - **Lab 7: IPv6 and OSPFv3 Troubleshooting**—This lab introduces IP version 6 (IPv6) addressing and discusses options for routing IPv6 traffic. One option for routing IPv6 traffic is OSPF version 3 (OSPFv3). The characteristics of OSPFv3 are discussed, and you are then introduced to configuration and troubleshooting commands that are useful when working with IPv6 and OSPFv3. Potential causes for an OSPFv3 adjacency issue are reviewed, and you are challenged to resolve a trouble ticket reporting that the lab topology does not have full IPv6 reachability throughout the network.
 - **Lab 8: IPv6 and RIPng Troubleshooting**—This lab describes the characteristics of RIPng, which is a routing protocol offering IPv6 support. RIPv6 configuration and troubleshooting syntax is presented, along with a trouble ticket that describes a RIPng route summarization problem.
 - **Lab 9: Cisco IOS Security Troubleshooting**—This lab discusses a collection of security features available on Cisco IOS routers. Also, several examples of Cisco IOS security issues are presented, along with a trouble ticket. The trouble ticket then addresses three separate security issues.
 - **Lab 10: DHCP Troubleshooting**—This lab reviews DHCP operation, in addition to configuration and troubleshooting syntax. Common symptoms of a DHCP issue are discussed. Finally, a DHCP trouble ticket needs to be resolved. Specifically, the trouble ticket addresses an issue where a Cisco IP Phone fails to obtain an IP address via DHCP.
 - **Lab 11: NAT Troubleshooting**—This lab explains NAT operation and reviews NAT address types. NAT troubleshooting syntax is presented, along with a listing of reasons that a NAT translation might fail. The presented trouble ticket describes a situation where NAT is partially working, but is still experiencing an issue.
 - **Lab 12: VoIP Troubleshooting**—This lab introduces several components and protocols used in voice over IP (VoIP) networks. The VoIP troubleshooting focus is on quality of service (QoS). A series of QoS mechanisms is discussed, and you are presented with configuration and troubleshooting syntax. This lab’s trouble ticket presents an issue where analog phones are experiencing voice quality issues, whereas a Cisco IP Phone is not experiencing this issue.

Spanning Tree Troubleshooting

This *Network Troubleshooting Video Mentor* demonstrates approaches to troubleshooting a *Spanning Tree Protocol (STP)* network issue. STP allows a topology to have redundant Layer 2 links, while preventing a Layer 2 loop where frames would endlessly circulate through the network and consume bandwidth.

Scenario

This lab includes the following steps:

- Step 1** Review Spanning Tree Protocol (STP) theory.
- Step 2** Examine lab topology.
- Step 3** Identify STP verification and troubleshooting commands.
- Step 4** Verify operation of STP in lab topology.
- Step 5** Discuss common symptoms of an STP issue.
- Step 6** Interpret a trouble ticket.
- Step 7** Troubleshoot and resolve the identified STP issue.
- Step 8** Summarize key elements of the STP troubleshooting process.

Review Spanning Tree Protocol (STP) Theory

STP prevents Layer 2 loops from occurring in a network, which could result in a broadcast storm or a corruption of a switch's MAC address table. Switches in an STP topology are classified as one of the following:

- **Root bridge:** The root bridge is a switch elected to act as a reference point for the spanning tree. The switch with the lowest bridge ID (BID) is elected as the root bridge. The BID is made up of a priority value and a MAC address.
- **Non-root bridge:** All other switches in the STP topology are considered to be non-root bridges.

Figure 1-1 illustrates the root bridge election in a network. Notice that because the bridge priorities are equal, the switch with the lowest MAC address is elected as the root bridge.

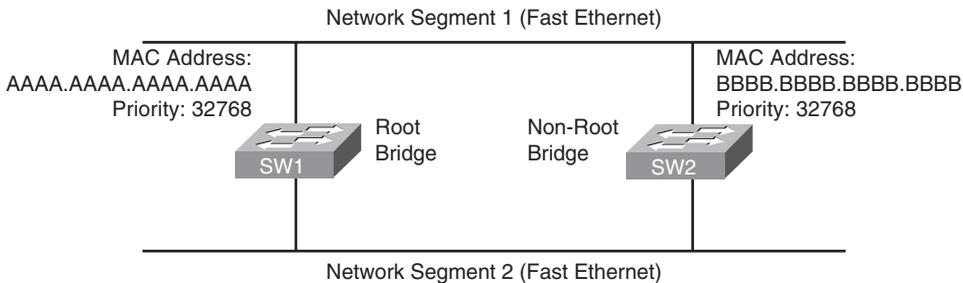


Figure 1-1 Root Bridge Election

Ports that interconnect switches in an STP topology are categorized as one of the following:

- **Root port:** Every non-root bridge has a single root port, which is the port on that switch that is closest to the root bridge, in terms of cost.
- **Designated port:** Every network segment has a single designated port, which is the port on that segment that is closest to the root bridge, in terms of cost. Therefore, all ports on a root bridge are designated ports.
- **Non-designated port:** Non-designated ports block traffic, in order to create a loop-free topology.

These ports are shown in Figure 1-2. Notice that the root port for switch SW2 was selected based on the lowest port ID, since the costs of both links are equal. Specifically, each link had a cost of 19, because the links are both Fast Ethernet links.

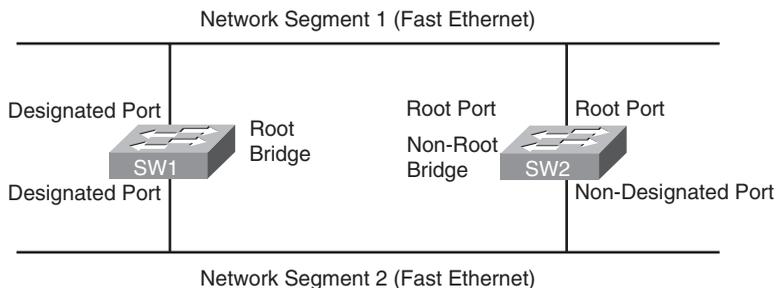


Figure 1-2 STP Port Assignments

The non-designated ports do not forward traffic during normal operation but do receive bridge protocol data units (BPDUs). If a link in the topology goes down, the non-designated port detects the link failure and determines if it needs to transition to the forwarding state.

If a non-designated port does need to transition to the forwarding state, it does not do so immediately. Rather, it transitions through the following states:

- **Blocking:** The port remains in the blocking state for 20 seconds. During this time, the non-designated port evaluates BPDUs in an attempt to determine its role in the spanning tree.

- **Listening:** The port moves from the blocking state to the listening state, and remains in this state for 15 seconds. During this time, the port sources BPDUs, which inform adjacent switches of the port's intent to forward data.
- **Learning:** The port moves from the listening state to the learning state, and remains in this state for 15 seconds. During this time, the port begins to add entries to its MAC address table.
- **Forwarding:** The port moves from the learning state to the forwarding state, and begins to forward frames.

Examine Lab Topology

The topology used in this lab is presented in Figure 1-3. SW1 is a Cisco Catalyst 3500 switch, and SW2 is a Cisco Catalyst 6506 switch.

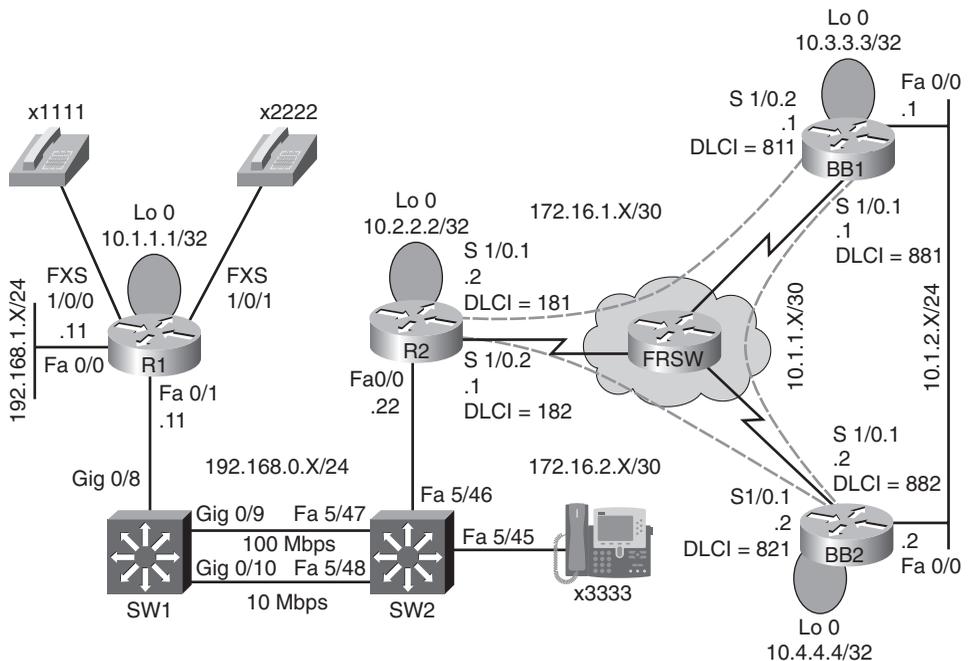


Figure 1-3 Lab Topology

Identify STP Verification and Troubleshooting Commands

This lab also illustrates the use of various **show** commands to verify and troubleshoot STP. The syntax for basic STP verification and configuration commands is presented in Table 1-1.

Table 1-1 STP Verification and Troubleshooting Syntax

Command	Description
Router# show spanning-tree summary	Displays a summary of port states.
Router# show spanning-tree	Provides information about all STP instances running on a switch.
Router# show spanning-tree interface <i>interface-id</i> detail	Shows interface status and configuration.
Router# spanning-tree cost <i>cost</i>	Specifies the path cost for STP calculation.
Router# spanning-tree vlan <i>vlan</i>	Enables STP for the specified VLAN.
Router# spanning-tree vlan <i>vlan</i> priority <i>priority</i>	Configures a switch's bridge priority for the specified VLAN.

Verify Operation of STP in Lab Topology

This lab issues various **show** commands to confirm the normal operation of STP between switches SW1 and Sw2. Examples 1-1 and 1-2 provide output from these commands.

Example 1-1 Baseline Output for Switch SW1

```

SW1#show spanning-tree vlan 1

VLAN0001
  Spanning tree enabled protocol ieee
  Root ID    Priority    32768
             Address    0009.122e.4181
             Cost      19
             Port      9 (GigabitEthernet0/9)
             Hello Time 2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    32769 (priority 32768 sys-id-ext 1)
             Address    000d.28e4.7c80
             Hello Time 2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time 300

Interface          Role Sts Cost          Prio.Nbr Type
-----
Gi0/8              Desg FWD 19           128.8   P2p
Gi0/9              Root FWD 19           128.9   P2p
Gi0/10             Altn BLK 100        128.10  Shr

SW1#show spanning-tree summary
Switch is in pvst mode

```

```

Root bridge for: none
Extended system ID          is enabled
Portfast Default            is disabled
PortFast BPDU Guard Default is disabled
Portfast BPDU Filter Default is disabled
Loopguard Default          is disabled
EtherChannel misconfig guard is enabled
UplinkFast                  is disabled
BackboneFast                is disabled
Configured Pathcost method used is short

```

Name	Blocking	Listening	Learning	Forwarding	STP Active
VLAN0001	1	0	0	2	3
1 vlan	1	0	0	2	3

SW1#show spanning-tree interface gig 0/10 detail

```

Port 10 (GigabitEthernet0/10) of VLAN0001 is alternate blocking
  Port path cost 100, Port priority 128, Port Identifier 128.10.
  Designated root has priority 32768, address 0009.122e.4181
  Designated bridge has priority 32768, address 0009.122e.4181
  Designated port id is 128.304, designated path cost 0
  Timers: message age 1, forward delay 0, hold 0
  Number of transitions to forwarding state: 0
  Link type is shared by default
  BPDU: sent 1, received 646

```

Example 1-2 Baseline Output for Switch SW2

```
SW2#show spanning-tree vlan 1
```

```

VLAN0001
  Spanning tree enabled protocol ieee
  Root ID    Priority 32768
             Address 0009.122e.4181
             This bridge is the root
             Hello Time 2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority 32768
             Address 0009.122e.4181
             Hello Time 2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time 300

```

Interface	Role	Sts	Cost	Prio.	Nbr	Type
Fa5/46	Desg	FWD	19	128.302		Shr
Fa5/47	Desg	FWD	19	128.303		P2p
Fa5/48	Desg	FWD	100	128.304		Shr

Discuss Common Symptoms of an STP Issue

Although performance and connectivity issues can be related to issues other than STP, this lab identifies the following as common symptoms of an STP issue:

- All users in a VLAN are simultaneously having performance or connectivity issues.
- Switch ports are experiencing unusually high port utilization.

Interpret a Trouble Ticket

The following trouble ticket is presented in this lab:

Users on network 192.168.1.X/24 are experiencing latency or no connectivity when attempting to reach network 10.1.2.X/24.

In this scenario, it is assumed that you gather additional information about this trouble ticket and discover that the port utilization LEDs on the Cisco Catalyst switches appear unusually high. Considering the information gathered, you suspect STP is not operating correctly.

Troubleshoot and Resolve the Identified STP Issue

To begin the troubleshooting process, this lab issues the **show spanning-tree vlan 1** command on each switch. As evidenced in Examples 1-3 and 1-4, STP is not enabled on either switch. Therefore, STP is enabled for VLAN 1 on each switch.

Example 1-3 Configuring STP on Switch SW1

```
SW1#
00:15:45: %SW_MATM-4-MACFLAP_NOTIF: Host 0009.b7fa.d1e1 in vlan 1 is flapping be
tween port Gi0/8 and port Gi0/9
SW1#
00:16:35: %SW_MATM-4-MACFLAP_NOTIF: Host 0009.b7fa.d1e1 in vlan 1 is flapping be
tween port Gi0/8 and port Gi0/9
SW1#
00:16:37: %SW_MATM-4-MACFLAP_NOTIF: Host c001.0e8c.0000 in vlan 1 is flapping be
tween port Gi0/9 and port Gi0/10
SW1#
00:16:41: %SW_MATM-4-MACFLAP_NOTIF: Host 0009.b7fa.d1e1 in vlan 1 is flapping be
tween port Gi0/8 and port Gi0/9
```

```
SW1#show spanning-tree vlan 1
```

```
Spanning tree instance(s) for vlan 1 does not exist.
```

```
SW1#conf term
```

```
Enter configuration commands, one per line. End with CNTL/Z.
```

```
SW1(config)#spanning-tree vlan 1
```

Example 1-4 Configuring STP on Switch SW2

```
SW2#show spanning-tree vlan 1
```

```
Spanning tree instance(s) for vlan 1 does not exist.
```

```
SW2#conf term
```

```
Enter configuration commands, one per line. End with CNTL/Z.
```

```
SW2(config)#spanning-tree vlan 1
```

This lab enables STP on both switches, which resolves the issue. However, other STP issues to keep in mind for the real world could stem from a duplex mismatch between switches. Also, bridge priority settings might be configured inappropriately, or port cost settings might be set incorrectly.

Summarize Key Elements of the STP Troubleshooting Process

This lab concludes by recapping best practices for troubleshooting STP. These best practices include verifying that STP is enabled, verifying that the switches have matching duplex settings, and checking port utilization.

Router Redundancy Troubleshooting

This *Network Troubleshooting Video Mentor* lab demonstrates approaches to troubleshooting a Hot Standby Router Protocol (HSRP) network issue. HSRP is one approach to providing next-hop router redundancy. Other approaches reviewed in this lab include Virtual Router Redundancy Protocol (VRRP) and Gateway Load Balancing Protocol (GLBP).

Scenario

This lab includes the following steps:

- Step 1** Describe three approaches to next-hop router redundancy.
- Step 2** Identify how HSRP is configured in the lab topology.
- Step 3** Review HSRP verification and troubleshooting commands.
- Step 4** Verify operation of HSRP in lab topology.
- Step 5** Discuss common symptoms of an HSRP issue.
- Step 6** Interpret a trouble ticket.
- Step 7** Troubleshoot and resolve the identified HSRP issue.
- Step 8** Summarize key elements of the HSRP troubleshooting process.

Describe Three Approaches to Next-Hop Router Redundancy

Many devices, such as PCs, are configured with a *default gateway*. The default gateway parameter identifies the IP address of a next-hop router. As a result, if that router were to become unavailable, devices that relied on the default gateway's IP address would be unable to send traffic off their local subnet.

Fortunately, Cisco offers technologies that provide next-hop gateway redundancy. These technologies include HSRP, VRRP, and GLBP. This troubleshooting lab begins by reviewing the basic operations of these technologies.

HSRP

HSRP uses virtual IP and MAC addresses. One router, known as the *active router*, services requests destined for the virtual IP and MAC addresses. Another router, known as the *standby router*, can service such requests in the event the active router becomes unavailable. Figure 2-1 illustrates a basic HSRP topology.

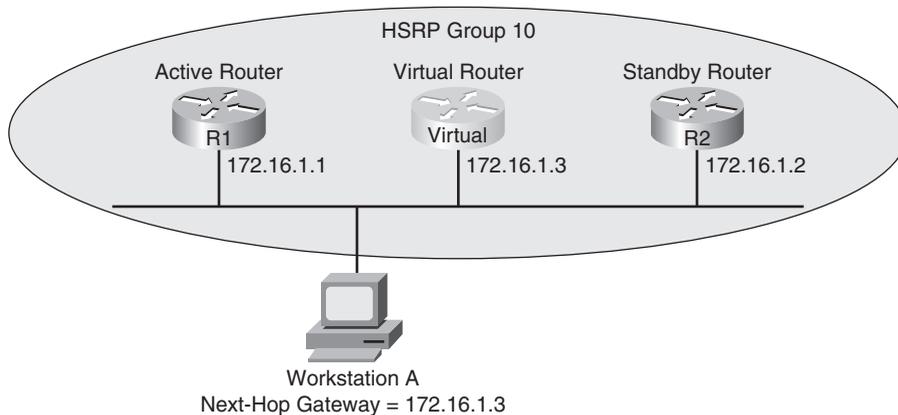


Figure 2-1 Basic HSRP Operation

VRRP

VRRP, similar to HSRP, allows a collection of routers to service traffic destined for a single IP address. Unlike HSRP, the IP address serviced by a VRRP group does not have to be a virtual IP address. The IP address can be the address of a physical interface on the virtual router master, which is the router responsible for forwarding traffic destined for the VRRP group's IP address. A VRRP group can have multiple routers acting as virtual router backups, as shown in Figure 2-2, any of which could take over in the event of the virtual router master becoming unavailable.

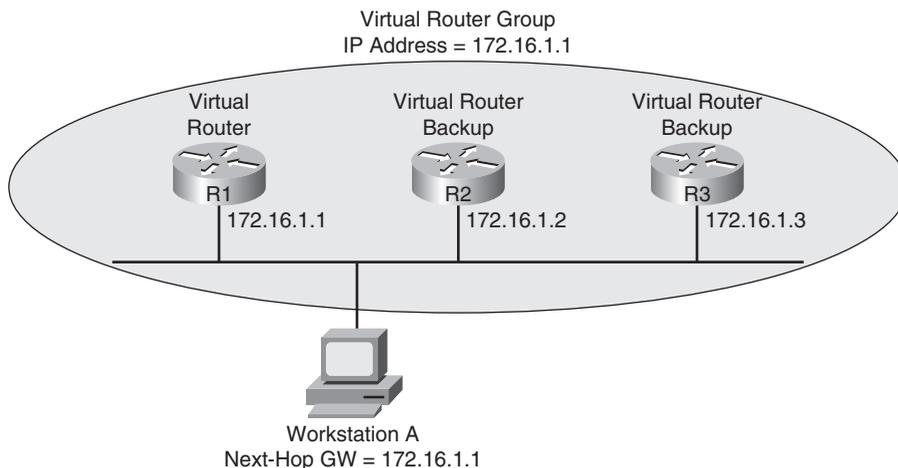


Figure 2-2 Basic VRRP Operation

GLBP

GLBP can load balance traffic destined for a next-hop gateway across a collection of routers, known as a *GLBP group*. Specifically, when a client sends an Address Resolution Protocol (ARP) request, in an attempt to determine the MAC address corresponding to a known IP address, GLBP

can respond with the MAC address of one member of the GLBP group. The next such request would receive a response containing the MAC address of a different member of the GLBP group, as depicted in Figure 2-3.

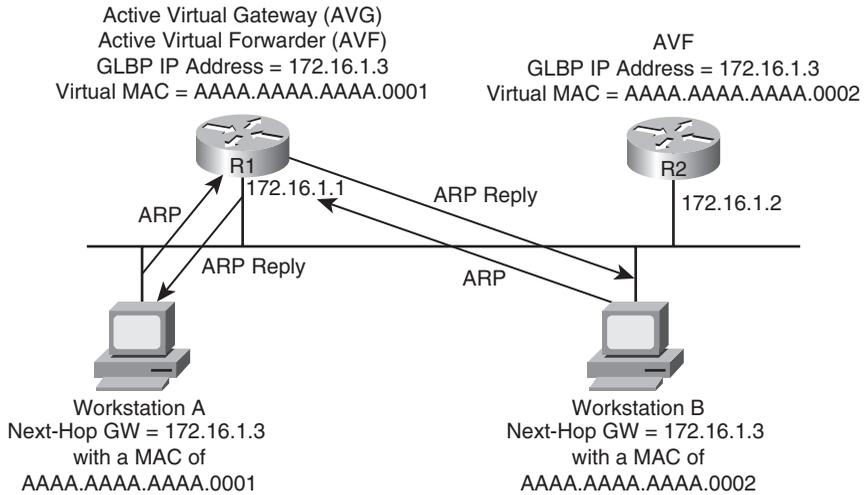


Figure 2-3 Basic GLBP Operation

Identify How HSRP Is Configured in the Lab Topology

In the lab topology, which is shown in Figure 2-4, the Fast Ethernet 0/1 interfaces on routers BB1 and BB2 are acting as HSRP interfaces for HSRP group 1. The virtual IP address being maintained by this HSRP group is 172.16.1.4.

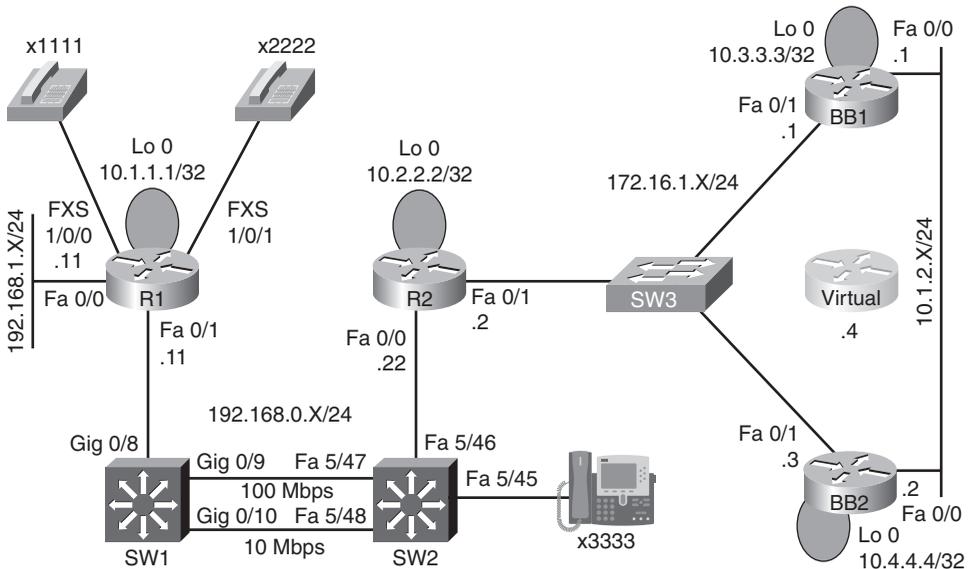


Figure 2-4 HSRP Lab Topology

Review HSRP Verification and Troubleshooting Commands

This lab also illustrates the use of **show**, **debug**, and HSRP configuration commands to verify and troubleshoot an HSRP network. The syntax for these HSRP verification and troubleshooting commands is presented in Table 2-1.

Table 2-1 HSRP Verification and Troubleshooting Syntax

Command	Description
Router# show standby <i>interface-id group</i>	Displays the HSRP configuration applied to a specified interface in a specified HSRP group.
Router# show standby brief	Provides a summary view of a router's HSRP configuration.
Router# debug standby	Shows HSRP state changes and information about sent and received HSRP packets.
Router(config-if)# standby group ip <i>virtual-ip-address</i>	Specifies the virtual IP address to be serviced by an HSRP group.
Router(config-if)# standby group priority <i>priority</i>	Configures an interface's HSRP priority (which defaults to 100).
Router(config-if)# standby group preempt	Causes a previously active HSRP router to regain its active status if it reboots.

Verify Operation of HSRP in Lab Topology

This lab issued various **show** and **debug** commands to confirm the normal operation of HSRP between routers BB1 and BB2. Examples 2-1 and 2-2 provide output from these commands.

Example 2-1 Baseline Output for Router BB1

```
BB1#show standby brief
                P indicates configured to preempt.
                |
Interface      Grp Prio P State   Active      Standby      Virtual IP
Fa0/1          1   150  Active local      172.16.1.3   172.16.1.4

BB1#debug standby
HSRP debugging is on
*Mar  1 01:14:21.487: HSRP: Fa0/1 Grp 1 Hello in 172.16.1.3 Standby pri 100 vI
P 172.16.1.4
*Mar  1 01:14:23.371: HSRP: Fa0/1 Grp 1 Hello out 172.16.1.1 Active pri 150 vI
P 172.16.1.4

BB1#u all
All possible debugging has been turned off
```

```
BB1#show standby fa 0/1 1
FastEthernet0/1 - Group 1
  State is Active
    10 state changes, last state change 00:12:40
  Virtual IP address is 172.16.1.4
  Active virtual MAC address is 0000.0c07.ac01
    Local virtual MAC address is 0000.0c07.ac01 (v1 default)
  Hello time 3 sec, hold time 10 sec
    Next hello sent in 1.536 secs
  Preemption disabled
  Active router is local
  Standby router is 172.16.1.3, priority 100 (expires in 9.684 sec)
  Priority 150 (configured 150)
  IP redundancy name is "hsrp-Fa0/1-1" (default)

BB1#show run
...output omitted...
hostname BB1
!
interface Loopback0
 ip address 10.3.3.3 255.255.255.255
!
interface FastEthernet0/0
 ip address 10.1.2.1 255.255.255.0
!
interface FastEthernet0/1
 ip address 172.16.1.1 255.255.255.0
 standby 1 ip 172.16.1.4
 standby 1 priority 150
!
router ospf 1
 network 0.0.0.0 255.255.255.255 area 0
...output omitted...
```

Example 2-2 Baseline Output for Router BB2

```

BB2#show standby brief
                P indicates configured to preempt.
                |
Interface  Grp Prio P State      Active      Standby      Virtual IP
Fa0/1      1   100 Standby    172.16.1.1  local        172.16.1.4

BB2#show run
...output omitted...
hostname BB2
!
interface Loopback0
 ip address 10.4.4.4 255.255.255.255
!
interface FastEthernet0/0
 ip address 10.1.2.2 255.255.255.0
!
interface FastEthernet0/1
 ip address 172.16.1.3 255.255.255.0
 standby 1 ip 172.16.1.4
!
router ospf 1
 network 0.0.0.0 255.255.255.255 area 0

```

A ping to the virtual IP address of 172.16.1.4 is issued from router R2 to confirm that HSRP is servicing requests for that IP address. Example 2-3 shows the output from the **ping** command.

Example 2-3 PINGING the Virtual IP Address from Router R2

```

R2#ping 172.16.1.4
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.1.4, timeout is 2 seconds:
!!!!

```

To verify that router BB2 will take over for BB1, in the event BB1's Fast Ethernet 0/1 fails, this lab shuts down the HSRP interface (that is Fast Ethernet 0/1) on BB1. In response to BB1's Fast Ethernet 0/1 interface becoming unavailable, BB2's Fast Ethernet 0/1 interface begins servicing the virtual IP address of 172.16.1.4.

Discuss Common Symptoms of an HSRP Issue

This lab identifies the following as common symptoms of an HSRP issue:

- An error message is displayed, indicating that the HSRP standby IP address is a duplicate IP address.
- The HSRP state of a router repeatedly changes.
- One router configured for HSRP does not recognize its HSRP peer.
- An HSRP router causes a port security violation on the switch to which it is attached.
- An HSRP router that was intended to be the active router is acting as the standby router, whereas the HSRP router that was intended to be the standby router is acting as the active router.

Interpret a Trouble Ticket

The following trouble ticket was presented in this lab:

A new network technician configured HSRP on routers BB1 and BB2, where BB1 was the active router. The configuration was initially working. However, now BB2 is acting as the active router, even though BB1 seems to be operational.

Troubleshoot and Resolve the Identified HSRP Issue

To begin the troubleshooting process, this lab reissues the **show standby brief** command on routers BB1 and BB2. As evidenced in Examples 2-4 and 2-5, router BB1 is administratively up with an HSRP priority of 150, whereas router BB2 is administratively up with a priority of 100.

Example 2-4 Examining the HSRP State of Router BB1's FastEthernet 0/1 Interface

```
BB1#show standby brief
                P indicates configured to preempt.
                |
Interface  Grp Prio P State      Active      Standby      Virtual IP
Fa0/1     1   150  Standby    172.16.1.3  local        172.16.1.4
```

Example 2-5 Examining the HSRP State of Router BB2's FastEthernet 0/1 Interface

```
BB2#show standby brief
                P indicates configured to preempt.
                |
Interface  Grp Prio P State      Active      Standby      Virtual IP
Fa0/1     1   100  Active     local        172.16.1.1  172.16.1.4
```

Upon examination of BB1's output, it becomes clear that the *preempt* feature is not enabled for the Fast Ethernet 0/1 interface on BB1. The absence of the preempt feature explains the reported symptom. Specifically, if BB1 had at one point been the active HSRP router for HSRP group 1, and either router BB1 or its Fast Ethernet 0/1 interface became unavailable, BB2 would have become the active router. Then, if BB1 or its Fast Ethernet 0/1 interface once again became available, BB1 would assume a standby HSRP role if BB1's FastEthernet 0/1 interface were not configured for the preempt feature.

To resolve this configuration issue, the preempt feature is added to BB1's Fast Ethernet 0/1 interface, as shown in Example 2-6. After enabling the preempt feature, notice that router BB1 regains its active HSRP role.

Example 2-6 Enabling the Preempt Feature on Router BB1's FastEthernet 0/1 Interface

```
BB1#conf term
Enter configuration commands, one per line.  End with CNTL/Z.
BB1(config)#int fa 0/1
BB1(config-if)#standby 1 preempt
BB1(config-if)#end
BB1#
*Mar  1 01:17:39.607: %HSRP-5-STATECHANGE: FastEthernet0/1 Grp 1 state Standby -
> Active

BB1#show standby brief
                P indicates configured to preempt.
                |
Interface      Grp Prio P State      Active          Standby          Virtual IP
Fa0/1          1  150 P Active     local           172.16.1.3       172.16.1.4
```

Summarize Key Elements of the HSRP Troubleshooting Process

This lab concludes by recapping common HSRP troubleshooting targets. These targets include a Layer 2 loop, network connectivity, port security, and HSRP configuration parameters.

EIGRP Troubleshooting

This *Network Troubleshooting Video Mentor* lab demonstrates approaches to troubleshooting Cisco's *Enhanced Interior Gateway Routing Protocol* (EIGRP). EIGRP exhibits characteristics of both link-state and distance-vector routing protocols. As a result, EIGRP is often classified as either an advanced distance-vector routing protocol or a balanced hybrid routing protocol.

Scenario

This lab includes the following steps:

- Step 1** Review EIGRP theory.
- Step 2** Examine lab topology.
- Step 3** Identify EIGRP verification and troubleshooting commands.
- Step 4** Verify operation of EIGRP in lab topology.
- Step 5** Discuss common EIGRP troubleshooting symptoms and resolutions.
- Step 6** Interpret a trouble ticket.
- Step 7** Troubleshoot and resolve the identified EIGRP issue.
- Step 8** Summarize key elements of the EIGRP troubleshooting process.

Review EIGRP Theory

As previously mentioned, EIGRP is considered to be a balanced hybrid routing protocol (or an advanced distance-vector routing protocol). Specifically, EIGRP advertises routes to directly attached neighbors, like a distance-vector routing protocol, while using a series of tables, similar to link-state databases.

EIGRP also offers the benefit of fast convergence after a link failure. Load balancing is supported over both equal-cost paths (a default behavior) and unequal-cost paths (through the *variance* feature).

Like most high-end routing protocols, EIGRP supports variable-length subnet masking (VLSM), and advertisements are sent via multicast (that is, an address of 224.0.0.10). In addition to IP, EIGRP can also act as the routing protocol for other routed protocols, including IPX and AppleTalk.

Parameters used to determine the best route include the following:

- **Advertised Distance (AD):** The distance from a neighbor to the destination network.
- **Feasible Distance (FD):** The AD plus the metric to reach the neighbor advertising the AD.

Routing information learned from EIGRP neighbors is inserted into the EIGRP topology table. The best route for a specific network in the IP EIGRP topology table becomes a candidate to be injected into the router's IP routing table.

EIGRP's metric is calculated with the following formula:

$$\text{EIGRP metric} = [\text{K1} * \text{bandwidth} + ((\text{K2} * \text{bandwidth}) / (256 - \text{load})) + \text{K3} * \text{delay}] * [\text{K5} / (\text{reliability} + \text{K4})]$$

By default, the K values are as follows:

- K1 = 1
- K2 = 0
- K3 = 1
- K4 = 0
- K5 = 0

As a result of these default K values, EIGRP's default metric can be calculated as follows:

$$\text{default EIGRP metric} = \text{bandwidth} + \text{delay}$$

- **bandwidth** = 10,000,000 / minimum bandwidth in kbps * 256
- **delay** = sum of delays of all interfaces in path in tens of milliseconds * 256

Examine Lab Topology

In the lab topology, which is shown in Figure 3-1, all routers have been configured with an EIGRP autonomous system of 1. Also, in router configuration mode on each router, a series of **network** statements has been issued to instruct all interfaces on each router to participate in the EIGRP routing process.

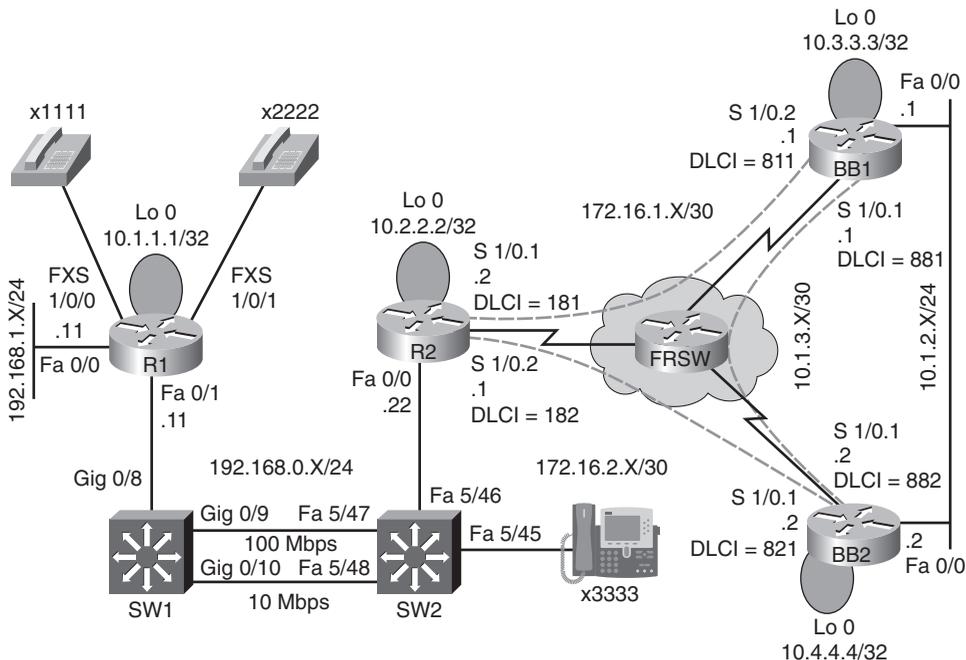


Figure 3-1 EIGRP Lab Topology

Identify EIGRP Verification and Troubleshooting Commands

This lab also illustrates the use of **show**, **debug**, and EIGRP configuration commands to verify and troubleshoot an EIGRP network. Table 3-1 presents the syntax for these EIGRP verification and troubleshooting commands.

Table 3-1 EIGRP Verification and Troubleshooting Syntax

Command	Description
Router# show ip route	Displays all routes in a router's routing table.
Router# show ip protocols	Shows current EIGRP settings.
Router# show ip eigrp topology	Lists networks learned via EIGRP.
Router# show ip eigrp neighbors	Verifies neighborships formed with EIGRP-speaking neighbors.
Router# debug eigrp packets	Provides real-time information about EIGRP packets being sent and received by a router.
Router(config)# router eigrp <i>autonomous-system-number</i>	Starts the EIGRP routing process. (NOTE: All routers that exchange EIGRP routing information must use the same autonomous system number.)
Router(config-router)# network <i>network [wildcard-mask]</i>	Specifies a connected network that will participate in the EIGRP routing process.
Router(config-router)# no auto-summary	Disables automatic network summarization.
Router(config-router)# variance multiplier	Determines the metric values over which EIGRP will load-balance traffic.

Verify Operation of EIGRP in Lab Topology

This lab issued various **show** and **ping** commands to confirm the operation of EIGRP on the routers in the topology. Example 3-1 shows the output of these commands on router R1.

Example 3-1 Baseline Output for Router R1

```
R1#show run | begin router
router eigrp 1
  network 10.1.1.1 0.0.0.0
  network 192.168.0.0
  network 192.168.1.0
  auto-summary

R1#show ip route
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
```

i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
ia - IS-IS inter area, * - candidate default, U - per-user static route
o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

```
D 172.16.0.0/16 [90/2562560] via 192.168.0.22, 00:14:17, FastEthernet0/1
   10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
D 10.0.0.0/8 is a summary, 00:14:10, Null0
C 10.1.1.1/32 is directly connected, Loopback0
C 192.168.0.0/24 is directly connected, FastEthernet0/1
C 192.168.1.0/24 is directly connected, FastEthernet0/0
```

R1#ping 172.16.1.1

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 172.16.1.1, timeout is 2 seconds:

!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 124/160/225 ms

R1#ping 172.16.1.2

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 172.16.1.2, timeout is 2 seconds:

!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 20/57/96 ms

R1#ping 172.16.2.1

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 172.16.2.1, timeout is 2 seconds:

!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 12/49/80 ms

R1#ping 172.16.2.2

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 172.16.2.2, timeout is 2 seconds:

!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 72/102/140 ms

R1#ping 10.3.3.3

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 10.3.3.3, timeout is 2 seconds:

```
.....
```

```
Success rate is 0 percent (0/5)
```

```
R1#ping 10.1.3.1
```

```
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 10.1.3.1, timeout is 2 seconds:
```

```
.....
```

```
Success rate is 0 percent (0/5)
```

```
R1#show ip protocols
```

```
Routing Protocol is "eigrp 1"
```

```
  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 1
```

```
  EIGRP NSF-aware route hold timer is 240s
```

```
  Automatic network summarization is in effect
```

```
  Automatic address summarization:
```

```
    192.168.1.0/24 for Loopback0, FastEthernet0/1
```

```
    192.168.0.0/24 for Loopback0, FastEthernet0/0
```

```
    10.0.0.0/8 for FastEthernet0/1, FastEthernet0/0
```

```
      Summarizing with metric 128256
```

```
Maximum path: 4
```

```
Routing for Networks:
```

```
  10.1.1.1/32
```

```
  192.168.0.0
```

```
  192.168.1.0
```

```
Routing Information Sources:
```

```
  Gateway          Distance      Last Update
```

```
  (this router)          90          00:15:26
```

```
  192.168.0.22          90          00:15:12
```

```
Distance: internal 90 external 170
```

```
R1#show ip eigrp neighbors
```

```
IP-EIGRP neighbors for process 1
```

H	Address	Interface	Hold Uptime	SRTT	RT0	Q	Seq
Num			(sec)	(ms)			Cnt
0	192.168.0.22	Fa0/1	14 03:11:37	190	1140	0	124

```
R1#debug eigrp packets
EIGRP Packets debugging is on
      (UPDATE, REQUEST, QUERY, REPLY, HELLO, IPXSAP, PROBE, ACK, STUB, SIAQUERY,
S
IAREPLY)
*Mar  3 03:34:07.336: EIGRP: Sending HELLO on Loopback0
*Mar  3 03:34:07.336:   AS 1, Flags 0x0, Seq 0/0 idbQ 0/0 iidbQ un/rely 0/0
*Mar  3 03:34:07.336: EIGRP: Received HELLO on Loopback0 nbr 10.1.1.1

R1#u all
Port Statistics for unclassified packets is not turned on.

All possible debugging has been turned off
```

Routers R2, BB1, and BB2 have similar EIGRP configurations. Examples 3-2, 3-3, and 3-4 show output from the **show run | begin router** command.

Example 3-2 Baseline Output for Router R2

```
R2#show run | begin router
router eigrp 1
  network 10.2.2.2 0.0.0.0
  network 172.16.1.0 0.0.0.3
  network 172.16.2.0 0.0.0.3
  network 192.168.0.0
  auto-summary
```

Example 3-3 Baseline Output for Router BB1

```
BB1#show run | begin router
router eigrp 1
  network 10.1.2.0 0.0.0.255
  network 10.1.3.0 0.0.0.3
  network 10.3.3.3 0.0.0.0
  network 172.16.1.0 0.0.0.3
  auto-summary
```

Example 3-4 Baseline Output for Router BB2

```
BB2#show run | begin router
BB2#show run | begin eigrp
router eigrp 1
  network 10.1.2.0 0.0.0.255
  network 10.1.3.0 0.0.0.3
  network 10.4.4.4 0.0.0.0
  network 172.16.2.0 0.0.0.3
  auto-summary
```

Based on the ping failures seen on router R1 and based on the route summaries appearing in router R1's IP routing table, this lab disables EIGRP's default auto-summarization behavior. Specifically, the **no auto-summary** command is issued in router configuration mode on each router.

Discuss Common EIGRP Troubleshooting Symptoms and Resolutions

This lab identifies the following as common symptoms of an EIGRP issue. Additionally, this lab discusses common resolutions and troubleshooting steps for these issues:

- EIGRP-speaking routers are not forming neighborships with neighboring EIGRP routers.
- Static or dynamic routes are not being redistributing into EIGRP.
- Some routes that should be learned via EIGRP are not appearing in the router's routing table.
- The *Not on Common Subnet* error message appears on the router console.
- Load balancing across unequal-cost paths is not functioning.
- A *Stuck in Active* error message appears on the router console.

Interpret a Trouble Ticket

The following trouble ticket is presented in this lab:

EIGRP has just been configured as the routing protocol for the network. After configuring EIGRP on all routers, instructing all router interfaces to participate in EIGRP, and disabling auto summarization, R2 does not appear to be load balancing across its links to BB1 and BB2 when sending traffic to network 10.1.2.0/24.

Troubleshoot and Resolve the Identified EIGRP Issue

To begin the troubleshooting process, this lab issues the **show ip route** command on router R2. The output of this command, as provided in Example 3-5, confirms that router R2's IP routing table only contains a single path to get to the backbone network of 10.1.2.0/24.

Example 3-5 Checking Load Balancing on Router R2

```

R2#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

    172.16.0.0/30 is subnetted, 2 subnets
C       172.16.1.0 is directly connected, Serial1/0.1
C       172.16.2.0 is directly connected, Serial1/0.2
    10.0.0.0/8 is variably subnetted, 6 subnets, 3 masks
C       10.2.2.2/32 is directly connected, Loopback0
D       10.1.3.0/30 [90/3072000] via 172.16.2.2, 00:00:34, Serial1/0.2
D       10.3.3.3/32 [90/2713600] via 172.16.2.2, 00:00:34, Serial1/0.2
D       10.1.2.0/24 [90/2585600] via 172.16.2.2, 00:00:34, Serial1/0.2
D       10.1.1.1/32 [90/409600] via 192.168.0.11, 00:00:46, FastEthernet0/0
D       10.4.4.4/32 [90/2688000] via 172.16.2.2, 00:00:34, Serial1/0.2
C       192.168.0.0/24 is directly connected, FastEthernet0/0
D       192.168.1.0/24 [90/284160] via 192.168.0.11, 00:18:33, FastEthernet0/0

```

Next, this lab checks the EIGRP Topology table to see if EIGRP has learned more than one route to reach the 10.1.2.0/24 network. The output, shown in Example 3-6, indicates that the EIGRP Topology table knows two routes that could be used to reach the 10.1.2.0/24 network.

Example 3-6 Checking the EIGRP Topology Table on Router R2

```

R2#show ip eigrp topology
IP-EIGRP Topology Table for AS(1)/ID(10.2.2.2)

Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

P 10.1.3.0/30, 1 successors, FD is 3072000
     via 172.16.2.2 (3072000/2169856), Serial1/0.2
     via 172.16.1.1 (4437248/2169856), Serial1/0.1
P 10.2.2.2/32, 1 successors, FD is 128256
     via Connected, Loopback0
P 10.1.2.0/24, 1 successors, FD is 2585600

```

```

        via 172.16.2.2 (2585600/281600), Serial1/0.2
        via 172.16.1.1 (3950848/281600), Serial1/0.1
P 10.3.3.3/32, 1 successors, FD is 2713600
        via 172.16.2.2 (2713600/409600), Serial1/0.2
        via 172.16.1.1 (4053248/128256), Serial1/0.1
P 10.1.1.1/32, 1 successors, FD is 409600
        via 192.168.0.11 (409600/128256), FastEthernet0/0
P 10.4.4.4/32, 1 successors, FD is 2688000
        via 172.16.2.2 (2688000/128256), Serial1/0.2
        via 172.16.1.1 (4078848/409600), Serial1/0.1
P 192.168.0.0/24, 1 successors, FD is 281600
        via Connected, FastEthernet0/0
P 192.168.1.0/24, 1 successors, FD is 284160
        via 192.168.0.11 (284160/28160), FastEthernet0/0
P 172.16.1.0/30, 1 successors, FD is 3925248
        via Connected, Serial1/0.1
P 172.16.2.0/30, 1 successors, FD is 2560000
        via Connected, Serial1/0.2

```

Router R2 is injecting only one of these 10.1.2.0/24 routes into the IP routing table because the feasible distances of the two routes are different. By default, EIGRP load balances over routes with equal-cost metrics (that is, equal feasible distances). However, at this point in the lab, the two routes have different metrics.

By examining the two metrics (that is, 2585600 and 3950848), this lab observes the values differ by less than a factor of two. Specifically, if you took the smallest metric of 2585600 and multiplied it by two, the result would be 5171200, which is greater than the largest metric of 3950848.

Because the metrics for the two routes vary by less than a factor of two, this lab configures EIGRP's *variance* feature on router R2, as shown in Example 3-7. Specifically, this configuration tells EIGRP on router R2 to not only inject the best EIGRP route into the IP routing table, but rather to inject the route with the best metric in addition to any route whose metric is within a factor of two of the best metric (that is, in the range 2585600–5171200). This allows the route with a metric of 3950848 to also be injected into the IP routing table.

Example 3-7 Configuring the Variance Feature on Router R2

```

R2#conf term
Enter configuration commands, one per line. End with CNTL/Z.
R2(config)#router eigrp 1
R2(config-router)#variance 2

```

To confirm that router R2 can now load balance across routers BB1 and BB2 to reach the 10.1.2.0/24 network, this lab reissues the **show ip route** command, the output of which is given in Example 3-8. This output confirms that router R2 can now load balance over two unequal-cost paths to reach the 10.1.2.0/24 network.

Example 3-8 Confirming Load Balancing on Router R2

```
R2#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

    172.16.0.0/30 is subnetted, 2 subnets
C      172.16.1.0 is directly connected, Serial1/0.1
C      172.16.2.0 is directly connected, Serial1/0.2
    10.0.0.0/8 is variably subnetted, 6 subnets, 3 masks
C      10.2.2.2/32 is directly connected, Loopback0
D      10.1.3.0/30 [90/3072000] via 172.16.2.2, 00:00:03, Serial1/0.2
          [90/4437248] via 172.16.1.1, 00:00:03, Serial1/0.1
D      10.3.3.3/32 [90/2713600] via 172.16.2.2, 00:00:03, Serial1/0.2
          [90/4053248] via 172.16.1.1, 00:00:03, Serial1/0.1
D      10.1.2.0/24 [90/2585600] via 172.16.2.2, 00:00:03, Serial1/0.2
          [90/3950848] via 172.16.1.1, 00:00:03, Serial1/0.1
D      10.1.1.1/32 [90/409600] via 192.168.0.11, 00:00:03, FastEthernet0/0
D      10.4.4.4/32 [90/2688000] via 172.16.2.2, 00:00:03, Serial1/0.2
          [90/4078848] via 172.16.1.1, 00:00:03, Serial1/0.1
C      192.168.0.0/24 is directly connected, FastEthernet0/0
D      192.168.1.0/24 [90/284160] via 192.168.0.11, 00:00:04, FastEthernet0/0
```

Summarize Key Elements of the EIGRP Troubleshooting Process

This lab concludes by recapping common EIGRP troubleshooting targets and reviewing how this lab's issues were resolved. Specifically, this lab resolved a reachability issue by disabling EIGRP's default auto-summarization behavior, and load balancing over unequal-cost paths was enabled by configuring the variance feature.

OSPF Troubleshooting

This *Network Troubleshooting Video Mentor* lab demonstrates approaches to troubleshooting the *Open Shortest Path First* (OSPF) routing protocol. OSPF falls under the *link-state* routing protocol classification. Like EIGRP, which was covered in Lab 3, “EIGRP Troubleshooting,” OSPF offers fast convergence and scalability, and is appropriate for many large enterprise networks.

Scenario

This lab includes the following steps:

- Step 1** Review OSPF theory.
- Step 2** Examine lab topology.
- Step 3** Identify OSPF verification and troubleshooting commands.
- Step 4** Verify operation of OSPF in lab topology.
- Step 5** Discuss common OSPF troubleshooting issues.
- Step 6** Interpret a trouble ticket.
- Step 7** Troubleshoot and resolve the identified OSPF issue.
- Step 8** Summarize key elements of the OSPF troubleshooting process.

Review OSPF Theory

OSPF is a link-state protocol that receives link-state advertisements (LSAs) from adjacent OSPF routers. The Dijkstra Algorithm takes the information contained in the LSAs to determine the shortest path to any destination within an area of the network.

Larger OSPF networks are often divided into areas. In a multiarea OSPF network, a backbone area (numbered *Area 0*) must exist, and all other areas must connect to Area 0. If an area is not physically adjacent to Area 0, a *virtual link* can be configured to logically connect the nonadjacent area with Area 0.

OSPF uses a metric of *cost*, which is a function of bandwidth. Cost can be calculated as follows:

$$\text{cost} = 100,000,000 / \text{bandwidth (in kbps)}$$

A multi-access network might have multiple routers residing on a common network segment. Rather than having all routers form a full-mesh of adjacencies with one another, a *designated router* (DR) can be elected, and all other routers on the segment can form an adjacency with the DR, as illustrated in Figure 4-1.

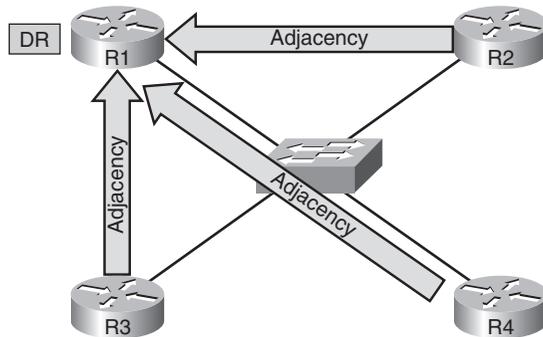


Figure 4-1 Designated Router Adjacencies

A DR is elected based on router priority, with higher-priority values being more preferable. If routers have equal priorities, the DR is elected based on the highest router ID. An OSPF router ID is the IP address of a router's loopback interface, or by the highest IP address on an active interface if the router is not configured with a loopback interface.

Examine Lab Topology

In the lab topology, which is shown in Figure 4-2, all routers have been configured to use the OSPF routing protocol. Also, in router configuration mode on each router, one or more **network** statements have been issued to instruct all interfaces on each router to participate in the OSPF.

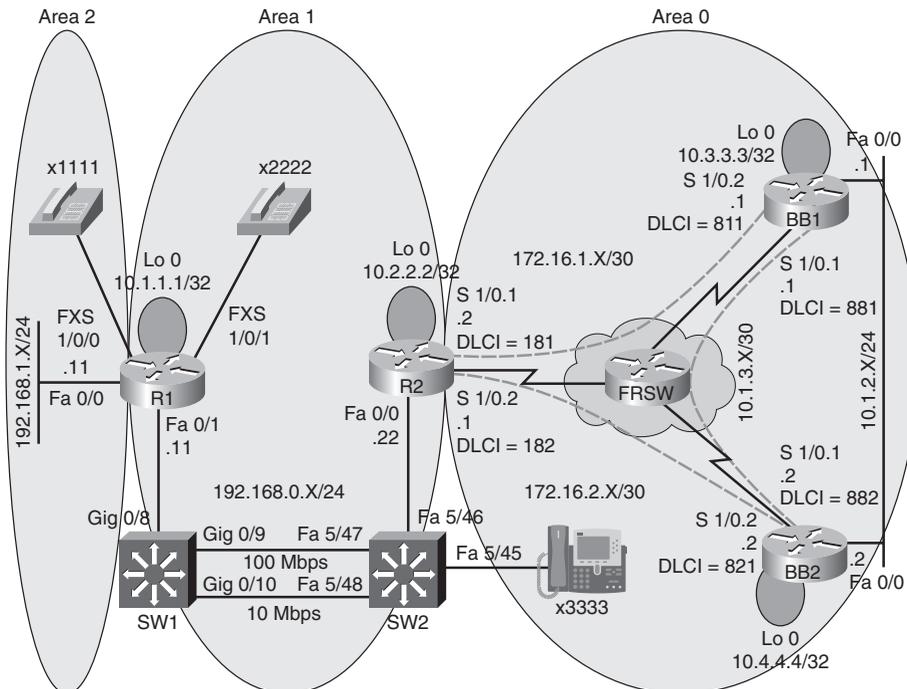


Figure 4-2 OSPF Lab Topology

Identify OSPF Verification and Troubleshooting Commands

This lab also illustrates the use of **show**, **debug**, and OSPF configuration commands to verify and troubleshoot an OSPF network. Table 4-1 presents the syntax for these OSPF verification and troubleshooting commands.

Table 4-1 OSPF Verification and Troubleshooting Syntax

Command	Description
Router# show ip route	Lists routes that have been injected into a router's IP routing table.
Router# show ip ospf neighbor	Displays OSPF neighbors learned off specified router interfaces.
Router# show ip ospf [interface <i>interface-id</i>]	Shows settings for each OSPF process running on a router or interface.
Router# show ip ospf virtual-links	Provides status information about virtual links.
Router# debug ip ospf packet	Verifies that OSPF packets are successfully flowing between routers.
Router(config)# router ospf <i>process-id</i>	Enables an OSPF process on a router.
Router(config-router)# network <i>network</i> [<i>wildcard-mask</i>] area <i>number</i>	Identifies a network participating in the OSPF process and the OSPF areas to which the network belongs.
Router(config-router)# ip ospf priority 0	Prevents a router interface from participating in a DR election.
Router(config-router)# area <i>number</i> virtual-link <i>router-id</i>	Creates a virtual link between a router and Area 0, where <i>number</i> = the area that is to be crossed to reach Area 0, and <i>router-id</i> = the OSPF router ID of the router on the opposite end of the virtual link.
Router(config-router)# neighbor <i>ip-address</i>	Statically configures a neighboring OSPF router in a nonbroadcast network.

Verify Operation of OSPF in Lab Topology

This lab issues various **show** commands to confirm the operation of OSPF on the routers in the topology. Examples 4-1, 4-2, 4-3, and 4-4 show sample output from these commands on routers R1, R2, R3, and R4.

Example 4-1 Baseline Output for Router R1

```
R1#show run | begin router
```

```
router ospf 1
```

```
area 1 virtual-link 10.2.2.2
network 10.1.1.1 0.0.0.0 area 1
network 192.168.0.0 0.0.0.255 area 1
network 192.168.1.0 0.0.0.255 area 2
```

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

Neighbor ID	Pri	State	Dead Time	Address	Interface
10.2.2.2	0	FULL/ -	-	192.168.0.22	OSPF_VL2
10.2.2.2	1	FULL/DR	00:00:38	192.168.0.22	FastEthernet0/1

```
R1#show ip route
```

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

```
Gateway of last resort is not set
```

```
172.16.0.0/30 is subnetted, 2 subnets
O    172.16.1.0 [110/134] via 192.168.0.22, 01:34:44, FastEthernet0/1
O    172.16.2.0 [110/81] via 192.168.0.22, 01:34:44, FastEthernet0/1
10.0.0.0/8 is variably subnetted, 6 subnets, 3 masks
O    10.2.2.2/32 [110/2] via 192.168.0.22, 02:24:31, FastEthernet0/1
O    10.1.3.0/30 [110/145] via 192.168.0.22, 01:34:44, FastEthernet0/1
O    10.3.3.3/32 [110/92] via 192.168.0.22, 01:34:44, FastEthernet0/1
O    10.1.2.0/24 [110/91] via 192.168.0.22, 01:34:45, FastEthernet0/1
C    10.1.1.1/32 is directly connected, Loopback0
O    10.4.4.4/32 [110/82] via 192.168.0.22, 01:34:45, FastEthernet0/1
C    192.168.0.0/24 is directly connected, FastEthernet0/1
C    192.168.1.0/24 is directly connected, FastEthernet0/0
```

```
R1#show ip ospf
```

```
Routing Process "ospf 1" with ID 10.1.1.1
```

```
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 SPFs 10000 msec
Maximum wait time between two consecutive SPFs 10000 msec
Incremental-SPF disabled
Minimum LSA interval 5 secs
Minimum LSA arrival 1000 msec
LSA group pacing timer 240 secs
Interface flood pacing timer 33 msec
Retransmission pacing timer 66 msec
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 3. 3 normal 0 stub 0 nssa
Number of areas transit capable is 1
External flood list length 0
  Area BACKBONE(0)
    Number of interfaces in this area is 1
      Area has no authentication
      SPF algorithm last executed 01:35:17.308 ago
      SPF algorithm executed 9 times
      Area ranges are
      Number of LSA 12. Checksum Sum 0x063B08
      Number of opaque link LSA 0. Checksum Sum 0x000000
      Number of DCbitless LSA 0
      Number of indication LSA 0
      Number of DoNotAge LSA 7
      Flood list length 0
    Area 1
      Number of interfaces in this area is 2 (1 loopback)
      This area has transit capability: Virtual Link Endpoint
      Area has no authentication
      SPF algorithm last executed 02:25:04.377 ago
      SPF algorithm executed 22 times
      Area ranges are
      Number of LSA 10. Checksum Sum 0x059726
      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 2
      Number of interfaces in this area is 1
      Area has no authentication
```

```

SPF algorithm last executed 02:25:15.880 ago
SPF algorithm executed 9 times
Area ranges are
Number of LSA 10. Checksum Sum 0x05F97B
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

```

R1#show ip ospf interface fa0/1

```

FastEthernet0/1 is up, line protocol is up
  Internet Address 192.168.0.11/24, Area 1
  Process ID 1, Router ID 10.1.1.1, Network Type BROADCAST, Cost: 1
  Transmit Delay is 1 sec, State BDR, Priority 1
  Designated Router (ID) 10.2.2.2, Interface address 192.168.0.22
  Backup Designated router (ID) 10.1.1.1, Interface address 192.168.0.11
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    oob-resync timeout 40
    Hello due in 00:00:00
  Supports Link-local Signaling (LLS)
  Index 2/2, flood queue length 0
  Next 0x0(0)/0x0(0)
  Last flood scan length is 1, maximum is 1
  Last flood scan time is 0 msec, maximum is 4 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 10.2.2.2 (Designated Router)
  Suppress hello for 0 neighbor(s)

```

Example 4-2 Baseline Output for Router R2

R2#show run | begin router

```

router ospf 1
  area 1 virtual-link 10.1.1.1
  network 10.2.2.2 0.0.0.0 area 1
  network 172.16.1.0 0.0.0.3 area 0
  network 172.16.2.0 0.0.0.3 area 0
  network 192.168.0.0 0.0.0.255 area 1

```

R2#show ip ospf neighbor

Neighbor ID	Pri	State	Dead Time	Address	Interface
10.4.4.4	0	FULL/ -	00:00:34	172.16.2.2	Serial1/0.2
10.3.3.3	0	FULL/ -	00:00:37	172.16.1.1	Serial1/0.1

```
10.1.1.1          0    FULL/ -           -           192.168.0.11    OSPF_VL0
10.1.1.1          1    FULL/BDR          00:00:39     192.168.0.11
FastEthernet0/
0
```

R2#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

```
       172.16.0.0/30 is subnetted, 2 subnets
C       172.16.1.0 is directly connected, Serial1/0.1
C       172.16.2.0 is directly connected, Serial1/0.2
       10.0.0.0/8 is variably subnetted, 6 subnets, 3 masks
C       10.2.2.2/32 is directly connected, Loopback0
O       10.1.3.0/30 [110/144] via 172.16.2.2, 01:34:50, Serial1/0.2
O       10.3.3.3/32 [110/91] via 172.16.2.2, 01:34:50, Serial1/0.2
O       10.1.2.0/24 [110/90] via 172.16.2.2, 01:34:50, Serial1/0.2
O       10.1.1.1/32 [110/11] via 192.168.0.11, 02:24:36, FastEthernet0/0
O       10.4.4.4/32 [110/81] via 172.16.2.2, 01:34:50, Serial1/0.2
C       192.168.0.0/24 is directly connected, FastEthernet0/0
O IA 192.168.1.0/24 [110/11] via 192.168.0.11, 01:34:50, FastEthernet0/0
```

R2#show run | begin router

```
router ospf 1
  area 1 virtual-link 10.1.1.1
  network 10.2.2.2 0.0.0.0 area 1
  network 172.16.1.0 0.0.0.3 area 0
  network 172.16.2.0 0.0.0.3 area 0
  network 192.168.0.0 0.0.0.255 area 1
```

Example 4-3 Baseline Output for Router BB1

```

BB1#show run | begin router
router ospf 1
  network 0.0.0.0 255.255.255.255 area 0

BB1#show ip ospf neighbor

Neighbor ID      Pri   State           Dead Time   Address        Interface
10.4.4.4         1    FULL/DR         00:00:38   10.1.2.2       FastEthernet0/
0
10.2.2.2         0    FULL/ -         00:00:39   172.16.1.2     Serial1/0.2
10.4.4.4         0    FULL/ -         00:00:38   10.1.3.2       Serial1/0.1

BB1#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

    172.16.0.0/30 is subnetted, 2 subnets
C      172.16.1.0 is directly connected, Serial1/0.2
O      172.16.2.0 [110/90] via 10.1.2.2, 01:35:01, FastEthernet0/0
    10.0.0.0/8 is variably subnetted, 6 subnets, 3 masks
O IA   10.2.2.2/32 [110/91] via 10.1.2.2, 01:35:01, FastEthernet0/0
C      10.1.3.0/30 is directly connected, Serial1/0.1
C      10.3.3.3/32 is directly connected, Loopback0
C      10.1.2.0/24 is directly connected, FastEthernet0/0
O IA   10.1.1.1/32 [110/101] via 10.1.2.2, 01:35:01, FastEthernet0/0
O      10.4.4.4/32 [110/11] via 10.1.2.2, 01:35:01, FastEthernet0/0
O IA   192.168.0.0/24 [110/100] via 10.1.2.2, 01:35:01, FastEthernet0/0
O IA   192.168.1.0/24 [110/101] via 10.1.2.2, 01:35:01, FastEthernet0/0

```

Example 4-4 Baseline Output for Router BB2

```

BB2#show run | begin router
router ospf 1
  network 0.0.0.0 255.255.255.255 area 0

BB2#show ip ospf neighbor

Neighbor ID      Pri   State           Dead Time   Address       Interface
10.2.2.2         0    FULL/ -        00:00:32   172.16.2.1   Serial1/0.2
10.3.3.3         0    FULL/ -        00:00:39   10.1.3.1     Serial1/0.1
10.3.3.3         1    FULL/BDR       00:00:35   10.1.2.1     FastEthernet0/
0

BB2#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

      172.16.0.0/30 is subnetted, 2 subnets
O       172.16.1.0 [110/143] via 10.1.2.1, 01:35:06, FastEthernet0/0
C       172.16.2.0 is directly connected, Serial1/0.2
      10.0.0.0/8 is variably subnetted, 6 subnets, 3 masks
O IA   10.2.2.2/32 [110/81] via 172.16.2.1, 01:35:06, Serial1/0.2
C     10.1.3.0/30 is directly connected, Serial1/0.1
O     10.3.3.3/32 [110/11] via 10.1.2.1, 01:35:06, FastEthernet0/0
C     10.1.2.0/24 is directly connected, FastEthernet0/0
O IA   10.1.1.1/32 [110/91] via 172.16.2.1, 01:35:06, Serial1/0.2
C     10.4.4.4/32 is directly connected, Loopback0
O IA  192.168.0.0/24 [110/90] via 172.16.2.1, 01:35:06, Serial1/0.2
O IA  192.168.1.0/24 [110/91] via 172.16.2.1, 01:35:06, Serial1/0.2

```

Discuss Common OSPF Troubleshooting Issues

This lab identifies the following as common symptoms of an OSPF issue. Additionally, this lab discusses common resolutions and troubleshooting steps for these issues:

- The *can't allocate router-id* message appears on the router console.
- Static or dynamic routes are not being redistributed into OSPF.
- Some routes that should be learned via OSPF are not appearing in the router's routing table.
- OSPF neighbors are not reaching the FULL state.

Interpret a Trouble Ticket

The following trouble ticket is presented in this lab:

For vendor interoperability reasons, a company changed its routing protocol from EIGRP to OSPF. The network was divided into areas to help constrain the flooding of routes, and all interfaces were instructed to participate in OSPF. However, none of the routers have full reachability to all of the subnets.

Troubleshoot and Resolve the Identified OSPF Issue

To begin the troubleshooting process, this lab issued the **show ip route** command on each router to determine which routes were missing from each router. The 192.168.1.0/24 network, connected to router R1's Fast Ethernet 0/0 interface, was not present in the IP routing tables for routers R2, BB1, and BB2. Also, the 10.4.4.4/32 network, the IP address of router BB2's Loopback 0 interface, was not present in the IP routing tables for routers R1, R2, and BB1.

After examining each router, the following configuration issues were identified and resolved:

- The virtual link configuration on router R1 was incorrect. Specifically, the transit area in the **area number virtual-link router-id** command was configured as Area 2. However, the transit area should have been Area 1. Example 4-5 shows the commands used to correct this misconfiguration.

Example 4-5 Correcting Router R1's Virtual Link Configuration

```
R1#conf term
Enter configuration commands, one per line. End with CNTL/Z.
R1(config)#router ospf 1
R1(config-router)#no area 2 virtual-link 10.2.2.2
R1(config-router)#area 1 virtual-link 10.2.2.2
```

- Subinterface Serial 1/0.1 on router BB1 had non-default Hello and Dead timers, which did not match the timers at the far end of the Frame Relay link. Example 4-6 illustrates how these non-default values were reset.

Example 4-6 Correcting Router BB1's Non-Default Timer Configuration

```
BB1#conf term
Enter configuration commands, one per line. End with CNTL/Z.
BB1(config)#int s1/0.1
BB1(config-subif)#no ip ospf hello-interval 60
BB1(config-subif)#no ip ospf dead-interval 200
```

- Interface FastEthernet 0/0 on router BB1 was configured with an incorrect OSPF network type of nonbroadcast. Example 4-7 demonstrates how this OSPF interface was reset to its default OSPF network type (that is, the broadcast OSPF network type).

Example 4-7 Correcting Router BB1's Incorrect OSPF Network Type Configuration

```
BB1#conf term
Enter configuration commands, one per line. End with CNTL/Z.
BB1(config)#int fa 0/0
BB1(config-if)#no ip ospf network non-broadcast
```

- Similar to the incorrect OSPF network type on router BB1's FastEthernet 0/0 interface, the Serial 1/0.2 subinterface on router BB2 was configured incorrectly. A point-to-point Frame Relay subinterface defaults to an OSPF network type of point-to-point. However, Serial 1/0.2 had been configured as an OSPF network type of nonbroadcast. Example 4-8 reviews how this non-default OSPF network type configuration was removed.

Example 4-8 Correcting Router BB2's Incorrect OSPF Network Type Configuration

```
BB2#conf term
Enter configuration commands, one per line. End with CNTL/Z.
BB2(config)#int s1/0.2
BB2(config-subif)#no ip ospf network non-broadcast
```

Summarize Key Elements of the OSPF Troubleshooting Process

This lab concludes by recapping common OSPF troubleshooting targets and reviewing how this lab's issues were resolved. Specifically, this lab reestablished full reachability throughout the topology by correcting a virtual link misconfiguration, two timer misconfigurations, and two OSPF network type misconfigurations.

Route Redistribution Troubleshooting

This *Network Troubleshooting Video Mentor* lab demonstrates approaches to troubleshoot route redistribution. Route redistribution allows one routing protocol (or a statically configured or a directly connected route) to be injected into another routing process. For example, routes learned via RIP could be injected into EIGRP.

Scenario

This lab includes the following steps:

- Step 1** Review route redistribution theory.
- Step 2** Examine lab topology.
- Step 3** Identify route redistribution verification and troubleshooting commands.
- Step 4** Verify operation of route redistribution in lab topology.
- Step 5** Discuss common route redistribution troubleshooting issues.
- Step 6** Interpret a trouble ticket.
- Step 7** Troubleshoot and resolve the identified route redistribution issue.
- Step 8** Summarize key elements of the route redistribution troubleshooting process.

Review Route Redistribution Theory

As previously mentioned, route redistribution allows routes learned via one method (for example, statically configured, locally connected, or routing protocol-learned routes) to be injected into a different routing process. If two routing protocols are mutually redistributed, the routes learned via each routing protocol are injected into the other routing protocol.

A network might benefit from route redistribution in the following scenarios:

- Transitioning to a more advanced routing protocol
- Merger of companies
- Different areas of administrative control

The router that sits at the boundary of the routing domains to be redistributed is known as a *boundary router*, as illustrated in Figure 5-1. A boundary router can redistribute static routes, connected routes, or routes learned via one routing protocol into another routing protocol.

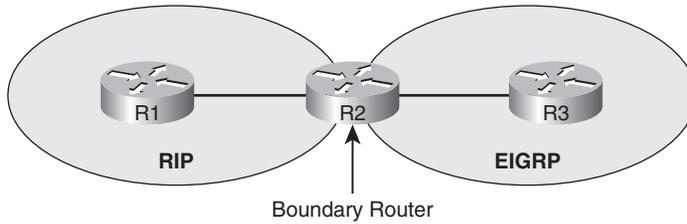


Figure 5-1 Boundary Router

Different routing protocols use different types of metrics, as illustrated in Figure 5-2. Therefore, when a route is injected into a routing protocol, a metric used by the destination routing protocol needs to be associated with the route being injected.

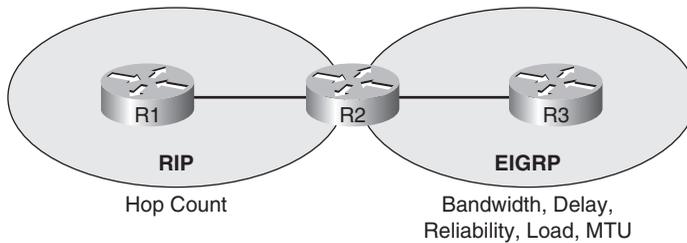


Figure 5-2 Differing Metric Parameters Between Routing Protocols

The metric assigned to a route being injected into another routing process is called a *seed metric*. The seed metric is needed to communicate relative levels of reachability between dissimilar routing protocols. A seed metric can be defined in one of three ways, as follows:

- The **default-metric** command
- The **metric** parameter in the **redistribute** command
- A route map configuration

Examine Lab Topology

In the lab topology, which is shown in Figure 5-3, route redistribution has been configured on router R2, to redistribute routes between EIGRP (running on routers R1 and R2) and OSPF running on routers R2, BB1, and BB2. Because router R2 is redistributing between OSPF and a non-OSPF protocol, router R2 is an *autonomous system boundary router* (ASBR).

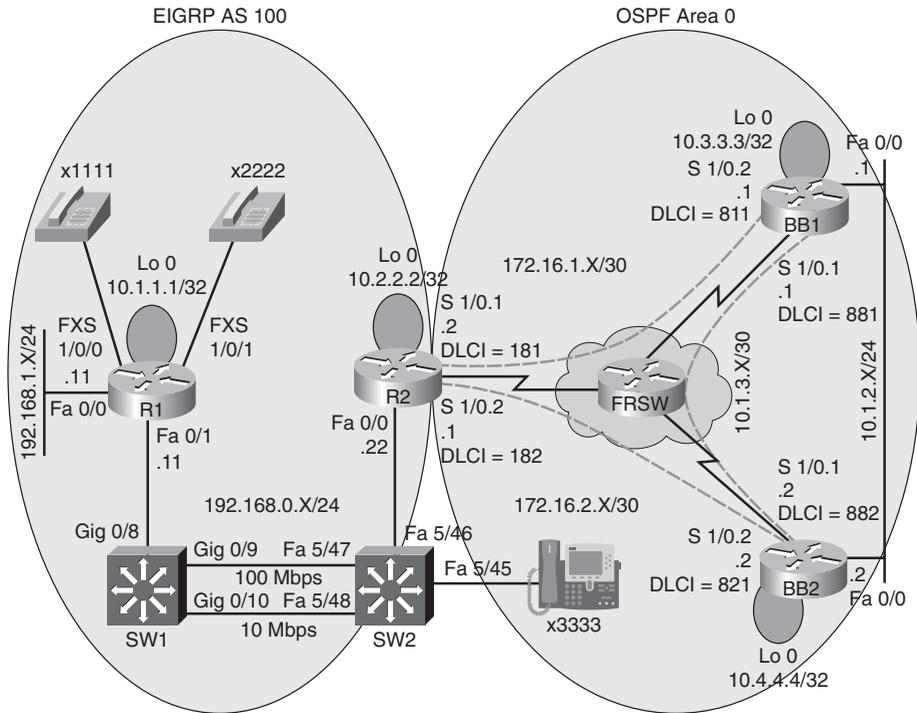


Figure 5-3 Route Redistribution Lab Topology

Identify Route Redistribution Verification and Troubleshooting Commands

This lab also illustrates the use of various **show** commands and route redistribution commands to verify and troubleshoot a reported route redistribution issue in the lab topology. Table 5-1 presents the syntax for these route redistribution verification and troubleshooting commands.

Table 5-1 Route Redistribution Verification and Troubleshooting Syntax

Command	Description
Router# show ip route	Displays routes in a router's IP routing table.
Router# debug ip eigrp notifications	Shows EIGRP routing updates in real-time.
Router(config)# router ospf process-id	Enables an OSPF process on a router.
Router(config-router)# redistribute eigrp autonomous-system-number [subnets]	Redistributes routes, including subnets, from a specified EIGRP autonomous system into OSPF.
Router(config-router)# default-metric metric	Specifies the metric used for routes redistributed into OSPF.
Router(config)# router eigrp autonomous-system-number	Enables an EIGRP routing process on a router.
Router(config-router)# redistribute ospf process-id	Redistributes routes from a specified OSPF

process ID into EIGRP.

```
Router(config-router)# default-metric
bandwidth delay reliability load mtu
```

Specifies the parameters used to calculate the seed metric for routes being redistributed into EIGRP, using the following EIGRP metric parameters:

- *bandwidth* (in kbps)
- *delay* (in tens of microseconds)
- *reliability* (maximum of 255)
- *load* (minimum of 1)
- *mtu* (in bytes)

```
Router(config-router)# distribute-list acl [in | out]
routing_protocol id
```

References an access control list (ACL) that determines which routes are advertised by a routing protocol or which routes are learned by a routing protocol, where *id* is an AS number for a routing protocol of EIGRP or a process ID for a routing protocol of OSPF.

Verify Operation of Route Redistribution in Lab Topology

This lab issued the **show ip route** command on each router to confirm the operation of route redistribution. Examples 5-1, 5-2, 5-3, and 5-4 provide the outputs from this command on each router.

Example 5-1 Baseline Output for Router R1

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

Gateway of last resort is not set

    172.16.0.0/30 is subnetted, 2 subnets
D EX   172.16.1.0 [170/1734656] via 192.168.0.22, 00:04:39, FastEthernet0/1
D EX   172.16.2.0 [170/1734656] via 192.168.0.22, 00:04:39, FastEthernet0/1
    10.0.0.0/8 is variably subnetted, 6 subnets, 3 masks
D      10.2.2.2/32 [90/156160] via 192.168.0.22, 00:04:39, FastEthernet0/1
D EX   10.1.3.0/30 [170/1734656] via 192.168.0.22, 00:04:39, FastEthernet0/1
D EX   10.3.3.3/32 [170/1734656] via 192.168.0.22, 00:04:39, FastEthernet0/1
D EX   10.1.2.0/24 [170/1734656] via 192.168.0.22, 00:04:40, FastEthernet0/1
C      10.1.1.1/32 is directly connected, Loopback0
D EX   10.4.4.4/32 [170/1734656] via 192.168.0.22, 00:04:40, FastEthernet0/1
C      192.168.0.0/24 is directly connected, FastEthernet0/1
C      192.168.1.0/24 is directly connected, FastEthernet0/0
```

Example 5-2 Baseline Output for Router R2

```

R2#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

      172.16.0.0/30 is subnetted, 2 subnets
C       172.16.1.0 is directly connected, Serial1/0.1
C       172.16.2.0 is directly connected, Serial1/0.2
      10.0.0.0/8 is variably subnetted, 6 subnets, 3 masks
C       10.2.2.2/32 is directly connected, Loopback0
O       10.1.3.0/30 [110/144] via 172.16.2.2, 00:07:12, Serial1/0.2
O       10.3.3.3/32 [110/91] via 172.16.2.2, 00:07:12, Serial1/0.2
O       10.1.2.0/24 [110/90] via 172.16.2.2, 00:07:12, Serial1/0.2
D       10.1.1.1/32 [90/409600] via 192.168.0.11, 00:04:46, FastEthernet0/0
O       10.4.4.4/32 [110/81] via 172.16.2.2, 00:07:12, Serial1/0.2
C       192.168.0.0/24 is directly connected, FastEthernet0/0
D       192.168.1.0/24 [90/284160] via 192.168.0.11, 00:04:46, FastEthernet0/0

```

Example 5-3 Baseline Output for Router BB1

```

BB1#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

      172.16.0.0/30 is subnetted, 2 subnets
C       172.16.1.0 is directly connected, Serial1/0.2
O       172.16.2.0 [110/90] via 10.1.2.2, 00:07:08, FastEthernet0/0
      10.0.0.0/8 is variably subnetted, 6 subnets, 3 masks
O E2   10.2.2.2/32 [110/64] via 10.1.2.2, 00:07:08, FastEthernet0/0
C       10.1.3.0/30 is directly connected, Serial1/0.1

```

```

C      10.3.3.3/32 is directly connected, Loopback0
C      10.1.2.0/24 is directly connected, FastEthernet0/0
O E2   10.1.1.1/32 [110/64] via 10.1.2.2, 00:04:49, FastEthernet0/0
O      10.4.4.4/32 [110/11] via 10.1.2.2, 00:07:08, FastEthernet0/0
O E2   192.168.0.0/24 [110/64] via 10.1.2.2, 00:07:08, FastEthernet0/0
O E2   192.168.1.0/24 [110/64] via 10.1.2.2, 00:04:49, FastEthernet0/0

```

Example 5-4 Baseline Output for Router BB2

```

BB2#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

      172.16.0.0/30 is subnetted, 2 subnets
O      172.16.1.0 [110/143] via 10.1.2.1, 00:08:48, FastEthernet0/0
C      172.16.2.0 is directly connected, Serial1/0.2
      10.0.0.0/8 is variably subnetted, 6 subnets, 3 masks
O E2   10.2.2.2/32 [110/64] via 172.16.2.1, 00:08:48, Serial1/0.2
C      10.1.3.0/30 is directly connected, Serial1/0.1
O      10.3.3.3/32 [110/11] via 10.1.2.1, 00:08:48, FastEthernet0/0
C      10.1.2.0/24 is directly connected, FastEthernet0/0
O E2   10.1.1.1/32 [110/64] via 172.16.2.1, 00:06:30, Serial1/0.2
C      10.4.4.4/32 is directly connected, Loopback0
O E2   192.168.0.0/24 [110/64] via 172.16.2.1, 00:08:48, Serial1/0.2
O E2   192.168.1.0/24 [110/64] via 172.16.2.1, 00:06:30, Serial1/0.2

```

Router R2, acting as a boundary router, had previously been configured for mutual route redistribution. The route redistribution configuration is illustrated in Example 5-5.

Example 5-5 Mutual Route Redistribution on Router R2

```
R2#show run | begin router
router eigrp 100
  redistribute ospf 1 metric 1500 100 255 1 1500
  network 10.2.2.2 0.0.0.0
  network 192.168.0.0
  no auto-summary
!
router ospf 1
  redistribute eigrp 100 metric 64 subnets
  network 172.16.1.0 0.0.0.3 area 0
  network 172.16.2.0 0.0.0.3 area 0
```

Discuss Common Route Redistribution Troubleshooting Issues

This lab identifies the following as common symptoms of a route redistribution issue. Additionally, this lab discusses common resolutions and troubleshooting tips for these issues:

- No external routes are being injected into a routing protocol.
- Subnets are not successfully redistributed into OSPF.
- A routing loop occurs.

Interpret a Trouble Ticket

The following trouble ticket is presented in this lab:

Company A has acquired Company B. Company A's network (that is, routers R1 and R2) uses EIGRP, whereas Company B's network (that is, routers BB1 and BB2) uses OSPF. Router R2 was configured as a boundary router, and router R2's configuration specifies that EIGRP and OSPF are mutually redistributed. However, routers R1, BB1, and BB2 do not see all of the subnets present in the network. The administrator of the newly merged network is also concerned that the mutual redistribution might lead to a routing loop.

Troubleshoot and Resolve the Identified Route Redistribution Issue

To begin the troubleshooting process, this lab issued the **show ip route** command on the routers reported to be missing IP routes, to determine exactly which routes were missing. Router R1's IP routing table lacked all OSPF-learned routes, as shown in Example 5-6.

Example 5-6 Router R1's IP Routing Table

```

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

Gateway of last resort is not set

    10.0.0.0/32 is subnetted, 2 subnets
D       10.2.2.2 [90/156160] via 192.168.0.22, 00:09:44, FastEthernet0/1
C       10.1.1.1 is directly connected, Loopback0
C       192.168.0.0/24 is directly connected, FastEthernet0/1
C       192.168.1.0/24 is directly connected, FastEthernet0/0

```

Router BB1, which was running OSPF, had some routes that originated in EIGRP. However, the 10.1.1.1/32 network, which was the IP address of router R1's Loopback 0 interface, was missing from router BB1's IP routing table, as illustrated in Example 5-7.

Example 5-7 Router BB1's IP Routing Table

```

BB1#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

    172.16.0.0/30 is subnetted, 2 subnets
C       172.16.1.0 is directly connected, Serial1/0.2
O       172.16.2.0 [110/90] via 10.1.2.2, 00:13:00, FastEthernet0/0
    10.0.0.0/8 is variably subnetted, 4 subnets, 3 masks
C       10.1.3.0/30 is directly connected, Serial1/0.1
C       10.3.3.3/32 is directly connected, Loopback0
C       10.1.2.0/24 is directly connected, FastEthernet0/0
O       10.4.4.4/32 [110/11] via 10.1.2.2, 00:13:00, FastEthernet0/0
O E2 192.168.0.0/24 [110/64] via 10.1.2.2, 00:01:14, FastEthernet0/0
O E2 192.168.1.0/24 [110/64] via 10.1.2.2, 00:01:14, FastEthernet0/0

```

Router BB2's IP routing table, as depicted in Example 5-8, was very similar to router BB1's IP routing table.

Example 5-8 Router BB2's IP Routing Table

```
BB2#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

    172.16.0.0/30 is subnetted, 2 subnets
O       172.16.1.0 [110/143] via 10.1.2.1, 00:13:39, FastEthernet0/0
C       172.16.2.0 is directly connected, Serial1/0.2
    10.0.0.0/8 is variably subnetted, 4 subnets, 3 masks
C       10.1.3.0/30 is directly connected, Serial1/0.1
O       10.3.3.3/32 [110/11] via 10.1.2.1, 00:13:39, FastEthernet0/0
C       10.1.2.0/24 is directly connected, FastEthernet0/0
C       10.4.4.4/32 is directly connected, Loopback0
O E2 192.168.0.0/24 [110/64] via 172.16.2.1, 00:01:53, Serial1/0.2
O E2 192.168.1.0/24 [110/64] via 172.16.2.1, 00:01:53, Serial1/0.2
```

After examining each router, the following configuration issues were identified and resolved:

- The EIGRP routing process lacked a default metric, which would be assigned to routes being redistributed into the EIGRP routing process. Example 5-9 shows the commands used to correct this misconfiguration.

Example 5-9 Adding a Default Metric for Router R2's EIGRP Routing Process

```
R2#conf term
Enter configuration commands, one per line. End with CNTL/Z.
R2(config)#router eigrp 100
R2(config-router)#default-metric 1500 100 255 1 1500
R2(config-router)#end
```

- The OSPF routing process lacked the **subnets** parameter at the end of the **redistribute** command, which would allow non-classful networks to be redistributed into OSPF. Example 5-10 illustrates how this configuration was corrected.

Example 5-10 Redistributing Subnets into Router R2's OSPF Routing Process

```
R2#conf term
Enter configuration commands, one per line. End with CNTL/Z.
R2(config)#router ospf 1
R2(config-router)#no redistribute eigrp 100 metric 64
R2(config-router)#redistribute eigrp 100 metric 64 subnets
R2(config-router)#end
```

- To prevent a possible routing loop from forming, due to the mutual route redistribution configuration, this lab also showed how distribute lists could be used on router R2 to prevent an EIGRP-learned route from being injected back into EIGRP and to prevent an OSPF-learned route from being injected back into OSPF. Example 5-11 demonstrates the configuration of the distribute lists.

Example 5-11 Using Distribute Lists to Prevent a Routing Loop

```
R2#conf term
Enter configuration commands, one per line. End with CNTL/Z.
R2(config)#access-list 1 deny 10.1.1.1 0.0.0.0
R2(config)#access-list 1 deny 192.168.1.0 0.0.0.255
R2(config)#access-list 1 permit any
R2(config)#access-list 2 deny 10.3.3.3 0.0.0.0
R2(config)#access-list 2 deny 10.4.4.4 0.0.0.0
R2(config)#access-list 2 deny 10.1.3.0 0.0.0.3
R2(config)#access-list 2 deny 10.1.2.0 0.0.0.255
R2(config)#access-list 2 permit any
R2(config)#router eigrp 100
R2(config-router)#distribute-list 1 out
R2(config-router)#exit
R2(config)#router ospf 1
R2(config-router)#distribute-list 2 out
R2(config-router)#end
```

Summarize Key Elements of the Route Redistribution Troubleshooting Process

This lab concludes by recapping common route redistribution troubleshooting targets and reviewing how this lab's issues were resolved. Specifically, this lab reestablished full reachability throughout the topology by adding a default metric for routes being redistributed into EIGRP and by specifying that individual subnets should be redistributed into OSPF. Finally, this lab demonstrated the use of distribute lists to prevent a routing loop, which could potentially result from a mutual route redistribution configuration.

BGP Troubleshooting

This *Network Troubleshooting Video Mentor* lab demonstrates troubleshooting a multihomed Border Gateway Protocol (BGP) environment, where an enterprise network can reach the Internet via one of two Internet Service Providers (ISPs). Unlike OSPF and EIGRP, which are considered to be interior gateway protocols (IGPs), BGP is an exterior gateway protocol (EGP). An EGP routes traffic between autonomous systems, where an autonomous system is a network under a single administrative control.

Scenario

This lab includes the following steps:

- Step 1** Review BGP theory.
- Step 2** Examine lab topology.
- Step 3** Identify BGP verification and troubleshooting commands.
- Step 4** Verify operation of BGP in lab topology.
- Step 5** Discuss common BGP troubleshooting issues.
- Step 6** Interpret a trouble ticket.
- Step 7** Troubleshoot and resolve the identified BGP issue.
- Step 8** Summarize key elements of the BGP troubleshooting process.

Review BGP Theory

A BGP routing process is associated with an autonomous system (AS). An AS is a network under a single administrative control (for example, a company). Although IGPs, such as OSPF or EIGRP, are used within an autonomous system, EGPs, such as BGP, are used between autonomous systems.

BGP is a path-vector routing protocol, similar to a distance-vector routing protocol. Specifically, BGP keeps track of the autonomous systems that must be transited to reach a destination network.

Much like EIGRP's topology table, BGP maintains a forwarding table, which lists all paths to a destination network. The best route for a network in a router's BGP forwarding table can be injected into the router's IP routing table, unless a more attractive route is already known by the IP routing table.

BGP Path Selection

Unlike OSPF and EIGRP, BGP does not consider a link's bandwidth when making a routing decision. Rather, BGP can use the following criteria when deciding how a packet should be forwarded:

BGP prefers the path with the highest weight. (**NOTE:** The BGP **weight** parameter is a Cisco-specific parameter.)

BGP prefers the path with the highest local preference value.

BGP prefers the path originated by BGP on the local router.

BGP prefers the path having the shortest autonomous system path.

BGP prefers the path with the lowest origin type. (**NOTE:** IGP < EGP < INCOMPLETE.)

BGP prefers the path with the lowest multi-exit discriminator.

BGP prefers eBGP paths over iBGP paths.

BGP prefers the path with the lowest IGP metric to the BGP next hop.

BGP prefers the path that points to a BGP router with lowest BGP router ID.

Influencing Outbound Path Selection with the Local Preference Attribute

This lab discusses how assigning local preference attributes to routes could cause a router to prefer one outbound path over another. A route map could be created to specify a local preference value, and that route map could be associated with a BGP neighbor, in either the inbound or outbound direction.

This lab also illustrates how assigning local preference values to routes coming into a router could cause that router to make its outbound routing decisions based on those local preference values. When using route maps to set local preference values, recall that higher values are preferred over lower values.

Influencing Inbound Path Selection with the ASPATH Attribute

This lab discusses how a router could influence inbound path selection by manipulating the AS path (that is, a listing of the ASes that must be transited to reach a specific destination network) advertised to a neighbor. Specifically, BGP can make routing decisions based on the smallest number of autonomous systems that must be crossed to reach a destination network. Using a route map, you can prepend one or more additional instances of your local AS to the ASPATH advertised to a router's neighbor, thereby making that path appear less attractive to its neighbor.

Examine Lab Topology

In the lab topology, which is shown in Figure 6-1, routers R1 and R2 are considered to be enterprise routers, whereas routers BB1 and BB2 are each considered to be routers for different Internet service providers (ISPs). Router R2 is the Internet boundary router for the enterprise network and has a multihomed connection to the Internet (that is, one connection via router BB1 and another connection via router BB2).

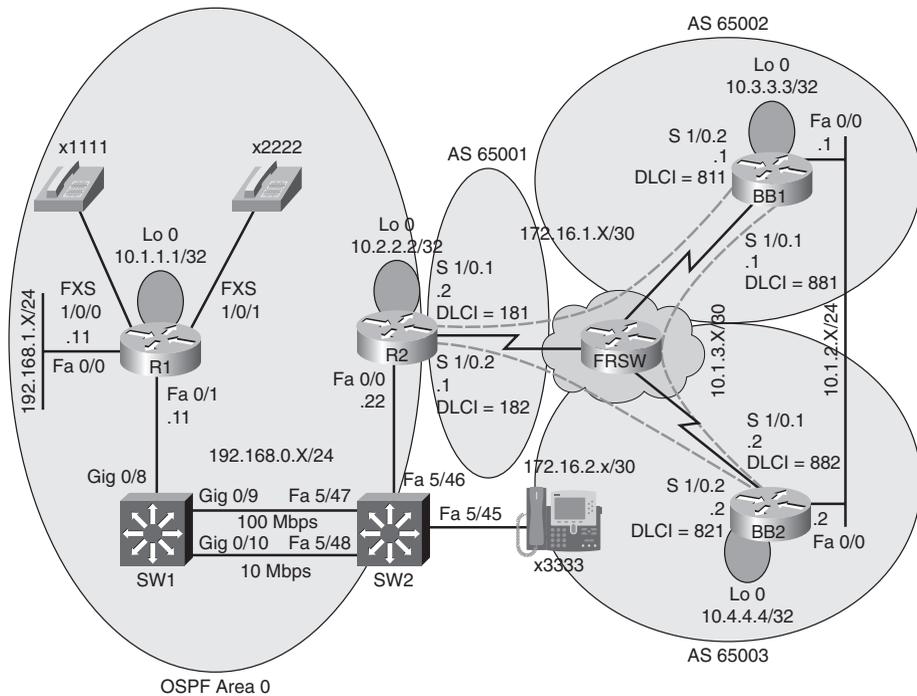


Figure 6-1 Lab 6 Topology

Identify BGP Verification and Troubleshooting Commands

This lab also illustrates the use of various **show** commands and BGP configuration commands to verify and troubleshoot a reported BGP issue in the lab topology. The syntax for these BGP commands is presented in Table 6-1.

Table 6-1 BGP Verification and Troubleshooting Syntax

Command	Description
Router# show ip route	Displays routes in a router's IP routing table.
Router# show ip bgp summary	Provides information about the status of BGP neighborships.
Router# show ip bgp	Shows the BGP forwarding table.
Router(config)# route-map tag [permit deny] [seq-num]	Creates a route map.
Router(config-route-map)# set local-preference <i>local-preference</i>	Sets the local preference BGP attribute for routes matched by a route map.
Router(config-route-map)# set as-path prepend <i>autonomous-system-number-1</i> [... <i>autonomous-system-number-n</i>]	Defines an autonomous system path to prepend to an autonomous system path known by the BGP forwarding table.
Router(config)# router bgp as-number	Enables the BGP process for a specific autonomous system.
Router(config-router)# neighbor ip-address route-map route-map-name [in out]	Applies a specified route map to routes received from or advertised to a specified BGP neighbor.

Verify Operation of BGP in Lab Topology

This lab issues the **show ip route** command on router R1 to confirm that this router has full reachability throughout the topology. Example 6-1 shows the output from this command. Notice, router R1 has a default route in its IP routing table. This default route is learned via OSPF from router R2.

Example 6-1 Baseline Output for Router R1

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

Gateway of last resort is 192.168.0.22 to network 0.0.0.0

    10.0.0.0/32 is subnetted, 2 subnets
O       10.2.2.2 [110/2] via 192.168.0.22, 00:05:33, FastEthernet0/1
C       10.1.1.1 is directly connected, Loopback0
C       192.168.0.0/24 is directly connected, FastEthernet0/1
C       192.168.1.0/24 is directly connected, FastEthernet0/0
O*E2 0.0.0.0/0 [110/1] via 192.168.0.22, 00:05:33, FastEthernet0/1
```

Router R2 is configured for both OSPF and BGP, with the BGP-learned default route being injected into OSPF, and with OSPF-learned routes being redistributed into BGP. Example 6-2 shows the initial IP routing table for router R2. Notice that the next-hop router for the default route is 172.16.1.1 (that is, router BB1).

Example 6-2 Baseline IP Routing Table on Router R2

```
R2#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 172.16.1.1 to network 0.0.0.0

    172.16.0.0/30 is subnetted, 2 subnets
C       172.16.1.0 is directly connected, Serial1/0.1
```

```

C      172.16.2.0 is directly connected, Serial1/0.2
      10.0.0.0/8 is variably subnetted, 6 subnets, 3 masks
C      10.2.2.2/32 is directly connected, Loopback0
B      10.1.3.0/30 [20/0] via 172.16.1.1, 00:01:40
B      10.3.3.3/32 [20/0] via 172.16.1.1, 00:01:40
B      10.1.2.0/24 [20/0] via 172.16.1.1, 00:01:40
O      10.1.1.1/32 [110/11] via 192.168.0.11, 00:08:17, FastEthernet0/0
B      10.4.4.4/32 [20/0] via 172.16.2.2, 00:01:40
C      192.168.0.0/24 is directly connected, FastEthernet0/0
O      192.168.1.0/24 [110/11] via 192.168.0.11, 00:08:17, FastEthernet0/0
B*    0.0.0.0/0 [20/0] via 172.16.1.1, 00:01:40

```

Example 6-3 illustrates the initial OSPF and BGP configuration on router R2.

Example 6-3 Initial Router Configuration on Router R2

```

R2#show run | begin router
router ospf 1
  log-adjacency-changes
  network 10.2.2.2 0.0.0.0 area 0
  network 192.168.0.0 0.0.0.255 area 0
  default-information originate
!
router bgp 65001
  no synchronization
  bgp log-neighbor-changes
  network 172.16.1.0 mask 255.255.255.252
  network 172.16.2.0 mask 255.255.255.252
  redistribute ospf 1
  neighbor 172.16.1.1 remote-as 65002
  neighbor 172.16.2.2 remote-as 65003
  no auto-summary

```

Example 6-4 shows the output of the **show ip bgp summary** command on router R2, which confirms that router R2 resides in BGP AS 65001. The output also confirms that BGP adjacencies have been formed with routers BB1 and BB2.

Example 6-4 BGP Configuration Summary on Router R2

```
R2#show ip bgp summary
BGP router identifier 10.2.2.2, local AS number 65001
BGP table version is 18, main routing table version 18
11 network entries using 1287 bytes of memory
20 path entries using 1040 bytes of memory
8/5 BGP path/bestpath attribute entries using 992 bytes of memory
4 BGP AS-PATH entries using 96 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 3415 total bytes of memory
BGP activity 38/27 prefixes, 75/55 paths, scan interval 60 secs
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
172.16.1.1	4	65002	102	97	18	0	0	00:02:47	7
172.16.2.2	4	65003	100	97	18	0	0	00:02:47	7

Router BB1 is configured for BGP and is sourcing a default route advertisement. Example 6-5 shows router BB1's IP routing table.

Example 6-5 Initial IP Routing Table on Router BB1

```
BB1#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 0.0.0.0 to network 0.0.0.0

172.16.0.0/30 is subnetted, 2 subnets
C    172.16.1.0 is directly connected, Serial1/0.2
B    172.16.2.0 [20/0] via 10.1.3.2, 00:03:01
10.0.0.0/8 is variably subnetted, 6 subnets, 3 masks
B    10.2.2.2/32 [20/0] via 172.16.1.2, 00:01:59
C    10.1.3.0/30 is directly connected, Serial1/0.1
C    10.3.3.3/32 is directly connected, Loopback0
C    10.1.2.0/24 is directly connected, FastEthernet0/0
B    10.1.1.1/32 [20/11] via 172.16.1.2, 00:01:59
B    10.4.4.4/32 [20/0] via 10.1.3.2, 00:40:10
B    192.168.0.0/24 [20/0] via 172.16.1.2, 00:01:59
B    192.168.1.0/24 [20/11] via 172.16.1.2, 00:01:59
S*   0.0.0.0/0 is directly connected, Null0
```

Router BB2's IP routing table, as shown in Example 6-6, is very similar to router BB1's IP routing table. Notice that router BB2 is also sourcing a default route and is advertising it via BGP to router R2. Therefore, router R2 will have two paths to reach a default route in its BGP forwarding table.

Example 6-6 Initial IP Routing Table on Router BB2

```
BB2#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 0.0.0.0 to network 0.0.0.0

    172.16.0.0/30 is subnetted, 2 subnets
B       172.16.1.0 [20/0] via 10.1.3.1, 00:03:11
C       172.16.2.0 is directly connected, Serial1/0.2
    10.0.0.0/8 is variably subnetted, 6 subnets, 3 masks
B       10.2.2.2/32 [20/0] via 172.16.2.1, 00:02:09
C       10.1.3.0/30 is directly connected, Serial1/0.1
B       10.3.3.3/32 [20/0] via 10.1.3.1, 00:40:10
C       10.1.2.0/24 is directly connected, FastEthernet0/0
B       10.1.1.1/32 [20/11] via 172.16.2.1, 00:02:09
C       10.4.4.4/32 is directly connected, Loopback0
B       192.168.0.0/24 [20/0] via 172.16.2.1, 00:02:09
B       192.168.1.0/24 [20/11] via 172.16.2.1, 00:02:09
S*     0.0.0.0/0 is directly connected, Null0
```

Discuss Common BGP Troubleshooting Issues

This lab identifies the following as common symptoms of a route redistribution issue. Additionally, this lab discusses common resolutions and troubleshooting tips for these issues:

- A BGP neighborhood is not being established or is flapping.
- BGP routes are not appearing as expected in a router's IP routing table.
- Routes are not being advertised to BGP neighbors.
- An inappropriate path is being selected in a multihomed BGP environment.

Interpret a Trouble Ticket

The following trouble ticket is presented in this lab:

Company A (that is, routers R1 and R2) has connections to two service providers (that is, BB1 and BB2). Router R2 is running BGP and is peering with both routers BB1 and BB2. The bandwidth between routers R2 and BB2 is greater than the bandwidth between routers R2 and BB1. Therefore, Company A wants to use the R2 to BB2 link as the primary link to the backbone network (that is, a default route). However, Company A noticed that the R2 to BB1 link is being used.

Troubleshoot and Resolve the Identified BGP Issue

The issue reported in the trouble ticket is a suboptimal routing issue. As shown earlier, in Example 6-2, router R2 preferred the 64-kbps link to router BB1 to reach a default route, as opposed to the 128-kbps link to router BB2. Therefore, the outbound routing from router R2 is suboptimal.

Also, the inbound routing, coming into the enterprise via router R2, is also suboptimal. To illustrate this point, consider Example 6-7, which shows the BGP forwarding table on router BB1. Notice that router BB1 prefers a next-hop router of router R2 to reach the 10.1.1.1/32 network, which resides inside the enterprise network (that is, the network comprised of routers R1 and R2). Using a next-hop router of R2 would force traffic over the 64-kbps link, rather than sending traffic from router BB1 over the 256-kbps link to router BB2, then over the 128-kbps link to router R2, and finally across the FastEthernet connection to router R1.

Example 6-7 BGP Forwarding Table on Router BB1

```
BB1#show ip bgp
BGP table version is 130, local router ID is 10.3.3.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
*  0.0.0.0          10.1.3.2          0                0 65003 i
*>                 0.0.0.0          0                32768 i
* 10.1.1.1/32      10.1.3.2          0                0 65003 65001 ?
*>                 172.16.1.2       11               0 65001 ?
* 10.1.2.0/24      10.1.3.2          0                0 65003 i
*>                 0.0.0.0          0                32768 i
* 10.1.3.0/30      10.1.3.2          0                0 65003 i
*>                 0.0.0.0          0                32768 i
* 10.2.2.2/32      10.1.3.2          0                0 65003 65001 ?
*>                 172.16.1.2       0                0 65001 ?
*> 10.3.3.3/32      0.0.0.0          0                32768 i
* 10.4.4.4/32      172.16.1.2          0                0 65001 65003 i
*>                 10.1.3.2          0                0 65003 i
* 172.16.1.0/30    172.16.1.2          0                0 65001 i
```

```

*>          0.0.0.0          0          32768 i
* 172.16.2.0/30 172.16.1.2    0          0 65001 i
*>          10.1.3.2        0          0 65003 i
* 192.168.0.0  10.1.3.2          0 65003 65001 ?
*>          172.16.1.2      0          0 65001 ?
* 192.168.1.0  10.1.3.2          0 65003 65001 ?
*>          172.16.1.2      11         0 65001 ?

```

To correct the inbound and outbound path selection issues, this lab only configured router R2. Although BGP attributes could have been manipulated on routers BB1 and BB2 to influence path selection, in a real-world environment, the administrator of the enterprise would not have privileges to configure the ISP routers.

Therefore, this lab configures local preference values for routes advertised into router R2 from routers BB1 and BB2 to prefer routes being advertised via router BB2. This configuration, which influences outbound path selection, is presented in Example 6-8.

Example 6-8 Local Preference Configuration on Router R2

```

R2(config)#route-map LOCALPREF-BB1
R2(config-route-map)#set local-preference 100
R2(config-route-map)#exit
R2(config)#route-map LOCALPREF-BB2
R2(config-route-map)#set local-preference 200
R2(config-route-map)#exit
R2(config)#router bgp 65001
R2(config-router)#neighbor 172.16.1.1 route-map LOCALPREF-BB1 in
R2(config-router)#neighbor 172.16.2.2 route-map LOCALPREF-BB2 in
R2(config-router)#exit

```

To influence inbound path selection, this lab configures a route map to prepend two additional instances of AS 65001 to routes being advertised via BGP from router R2 to router BB1. Example 6-9 shows this configuration, which causes router BB1 to use router BB2 as a next-hop router when sending traffic into the enterprise network, since the path via router BB2 appears to be fewer AS hops away from the enterprise networks.

Example 6-9 ASPATH Configuration on Router R2

```
R2(config)#route-map ASPATH 10
R2(config-route-map)#set as-path prepend 65001 65001
R2(config-route-map)#exit
R2(config)#router bgp 65001
R2(config-router)#neighbor 172.16.1.1 route-map ASPATH out
R2(config-router)#end
```

Summarize Key Elements of the BGP Troubleshooting Process

This lab concludes by reviewing common BGP troubleshooting targets and reviewing how the identified issues were resolved. Specifically, this lab overrides BGP's default path selection by statically configuring local preference and ASPATH attributes on router R2.

IPv6 and OSPFv3 Troubleshooting

This *Network Troubleshooting Video Mentor* lab demonstrates troubleshooting an OSPFv3 adjacency issue in an IPv6-based network. Many of the same troubleshooting targets exist in both the OSPFv3 and OSPFv2 environments. Therefore, many OSPF troubleshooting issues discussed in Lab 4, “OSPF Troubleshooting,” are applicable to IPv6 OSPFv3 troubleshooting.

Scenario

This lab includes the following steps:

- Step 1** Review IPv6 and OSPFv3 theory.
- Step 2** Examine lab topology.
- Step 3** Identify IPv6 and OSPFv3 verification and troubleshooting commands.
- Step 4** Verify operation of IPv6 and OSPFv3 in lab topology.
- Step 5** Discuss possible resolutions to an OSPFv3 adjacency issue.
- Step 6** Interpret a trouble ticket.
- Step 7** Troubleshoot and resolve the identified OSPFv3 adjacency issue.
- Step 8** Summarize key elements of the OSPFv3 adjacency troubleshooting process.

Review IPv6 and OSPFv3 Theory

This section reviews IPv6 address types, the IPv6 address structure, routing options for IPv6, and characteristics of OSPFv3. Unlike OSPFv2, OSPFv3 is capable of populating an IPv6 routing table.

IPv6 Address Types

IPv6 has three types of addresses, as follows:

- Unicast
- Multicast
- Anycast

Unicast

With unicast, a single IPv6 address is applied to a single interface, as illustrated in Figure 7-1. The communication flow can be one-to-one.

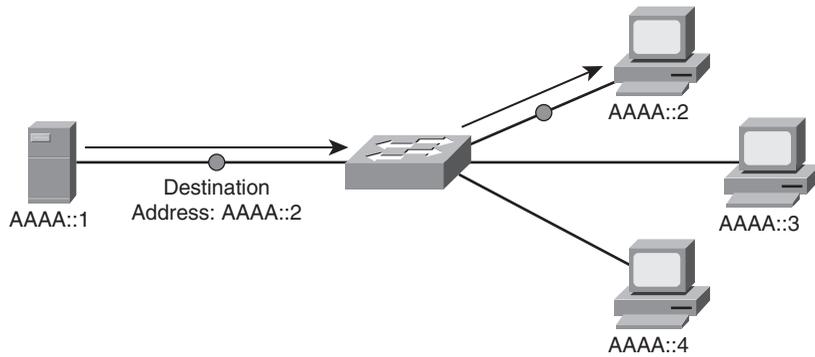


Figure 7-1 IPv6 Unicast Example

Multicast

With multicast, a single IPv6 address (that is, a multicast group) represents multiple devices on a network, as seen in Figure 7-2. The communication flow is one-to-many.

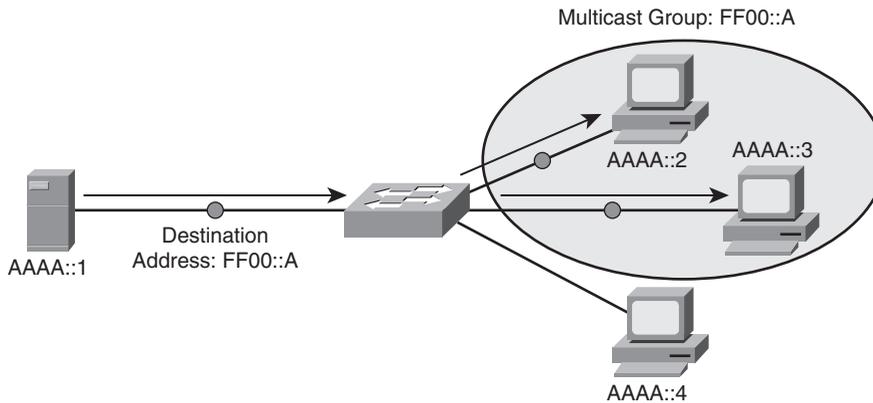


Figure 7-2 IPv6 Multicast Example

Anycast

With anycast, a single IPv6 address is assigned to multiple devices, as depicted in Figure 7-3. The communication flow is one-to-nearest (from the perspective of a router's routing table).

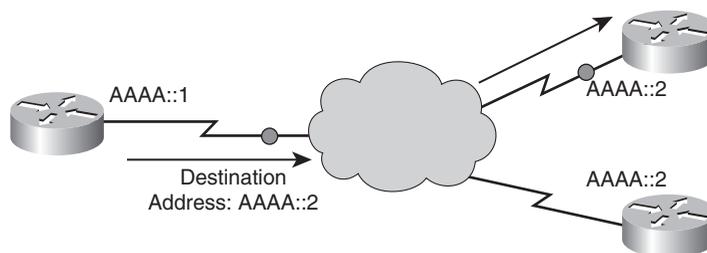


Figure 7-3 IPv6 Anycast Example

IPv6 Address Format

An IPv6 address has the following address format, where X is a hexadecimal digit in the range of 0–F:

```
XXXX:XXXX:XXXX:XXXX:XXXX:XXXX:XXXX:XXXX
```

A hexadecimal digit is four bits in size (that is, four binary bits can represent sixteen digits).

Notice that an IPv6 address has eight fields, and each field contains four hexadecimal digits. The following formula reveals why an IPv6 address is a 128-bit address:

$$4 \text{ bits per digit} * 4 \text{ digits per field} * 8 \text{ fields} = 128 \text{ bits in an IPv6 address}$$

Because IPv6 addresses can be difficult to work with due to their size, the following rules exist for abbreviating these addresses:

- Leading zeros in a field can be omitted.
- Contiguous fields containing all zeros can be represented with a double colon. (NOTE: This can only be done once for a single IPv6 address.)

As an example, consider the following IPv6 address:

```
ABCD:0123:4040:0000:0000:0000:000A:000B
```

Using the rules for abbreviation, the IPv6 address can be rewritten as follows:

```
ABCD:123:4040::A:B
```

Also, the Extended Unique Identifier (EUI-64) format can be used to cause the router to automatically populate the low-order 64 bits of an IPv6 address based on an interface's MAC address.

IPv6 Routing Options

IPv6 maintains a separated routing table from IPv4. Following are methods of populating this IPv6 routing table:

- **Static routes:** Configured similar to IPv4 static routes.
- **RIP next generation (RIPng):** Has many of the same characteristics as RIPv2 (for example, a distance vector routing protocol with a 15 hop-count maximum).
- **OSPFv3:** Builds on OSPFv2 to add support for IPv6 network characteristics (for example, 128-bit network addresses and link-local addresses).

- **IS-IS for IPv6:** Very similar to IS-IS for IPv4, with a few IPv6 extensions added (for example, new Type, Length, Value (TLV) attributes, and a new protocol ID).
- **Multiprotocol BGP:** Allows BGP to route protocols other than IPv4 (for example, IPv6).
- **EIGRP:** Configured on the interfaces with IPv6 addressing, similar to OSPFv3.

Characteristics of OSPFv3

Following are primary characteristics of the OSPFv3 routing protocol:

- Maintains several similarities with OSPFv2:
 - Uses a hierarchical structure divided into areas.
 - Requires direct connectivity from the backbone area to all other areas.
 - Uses many of the same packet types as OSPFv2 (for example, Hello packets).
 - Leverages OSPFv2's approach to forming adjacencies with neighbors.
- Contains enhancements that support IPv6:
 - Routes over links rather than over networks.
 - Uses IPv6 link-local addresses to identify neighbors.
 - Can support multiple IPv6 subnets on a single link.
 - Allows communication between two nodes connected to a common link, even though the two nodes might not share a common subnet.
 - Supports multiple instances of OSPFv3 running over a common link.

Examine Lab Topology

In the lab topology, which is shown in Figure 7-4, all routers have been configured with IPv6 addressing on their physical interfaces. Additionally, these interfaces have been configured to participate in one of two OSPF areas. Routers BB1 and BB2 participate in the OSPF backbone area of Area 0, in addition to the Serial 1/0 subinterfaces on router R2. Router R2's FastEthernet 0/0 interfaces and router R1 participate in OSPF Area 1.

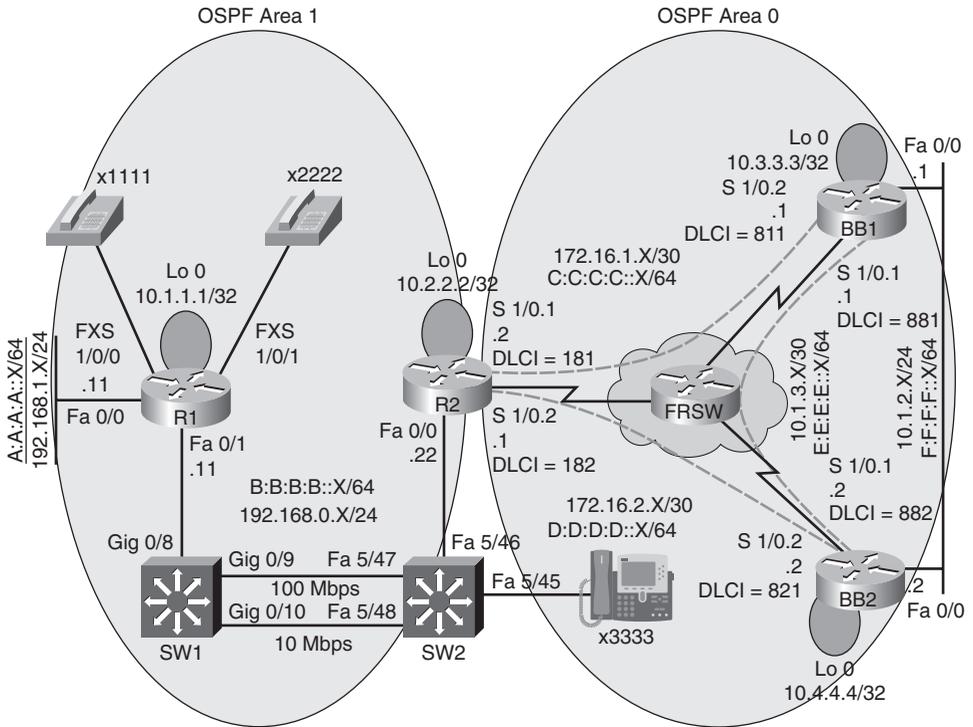


Figure 7-4 Lab 7 Topology

Identify IPv6 and OSPFv3 Verification and Troubleshooting Commands

This lab also illustrates the use of various **show**, **debug**, and IPv6 and OSPFv3 configuration commands to verify and troubleshoot a reported OSPFv3 adjacency issue in the lab topology. Table 7-1 presents the syntax for these commands.

Table 7-1 IPv6 and OSPFv3 Verification and Troubleshooting Syntax

Command	Description
Router# show ipv6 ospf	Displays OSPFv3 routing process, router ID, various timers, and information about each area on a router.
Router# show ipv6 ospf interface	Shows IPv6 link local address, area ID, process ID, router ID, and cost.
Router# show ipv6 ospf neighbor	Lists the state of a router's adjacency with all configured OSPFv3 neighbors.
Router# debug ipv6 ospf adj	Displays information about OSPFv3 adjacencies.
Router# debug ip ipv6 ospf hello	Shows OSPFv3 HELLO packet information.
Router(config)# ipv6 cef	Configures Cisco Express Forwarding for IPv6.

Table 7-1 IPv6 and OSPFv3 Verification and Troubleshooting Syntax

Command	Description
Router(config)# ipv6 unicast-routing	Globally instructs a router to forward IPv6 traffic.
Router(config-if)# ipv6 address <i>ipv6-address/prefix-length [eui-64]</i>	Assigns an IPv6 address to an interface. (NOTE: The eui-64 option allows a router to complete the low-order 64 bits of an address, based on the interface's MAC address.)
Router(config-if)# ipv6 ospf <i>process-id area area-id</i>	Allows the IPv6 address configured on an interface to participate in an OSPFv3 routing process.
Router(config)# ipv6 router ospf <i>process-id</i>	Globally enables an OSPFv3 routing process on a router.
Router(config-rtr)# router-id <i>ipv4-address</i>	Specifies an IPv4 address to be used by OSPFv3 as a router's router ID.

Verify Operation of IPv6 and OSPFv3 in Lab Topology

The routers in the topology are similarly configured for IPv6 and OSPFv3. This lab issues the **show run** command on router R2, the output of which is provided in Example 7-1, to illustrate the configuration required to globally enable IPv6 routing, to assign IPv6 addresses to interfaces, and to instruct an interface to participate in an OSPF routing process for a particular area.

Example 7-1 Running Configuration on Router R2

```
R2#show run
Building configuration...

hostname R2
-- OUTPUT OMITTED --
!
ipv6 unicast-routing
ipv6 cef
!
interface Loopback0
 ip address 10.2.2.2 255.255.255.255
!
interface FastEthernet0/0
 ip address 192.168.0.22 255.255.255.0
 duplex auto
 speed auto
 ipv6 address B:B:B:B::22/64
 ipv6 ospf 1 area 1
!
interface Serial1/0
 no ip address
```

```

encapsulation frame-relay
!
interface Serial1/0.1 point-to-point
ip address 172.16.1.2 255.255.255.252
ipv6 address C:C:C:C::2/64
ipv6 ospf 1 area 0
frame-relay interface-dlci 181
!
interface Serial1/0.2 point-to-point
ip address 172.16.2.1 255.255.255.252
ipv6 address D:D:D:D::1/64
ipv6 ospf 1 area 0
frame-relay interface-dlci 182
!
ipv6 router ospf 1
log-adjacency-changes
!
-- OUTPUT OMITTED --

```

To confirm that all routers had full reachability throughout the topology, the **show ipv6 route** command is issued on each router. The outputs confirmed that each router contained routes to all IPv6 networks configured in the topology. Example 7-2 provides a sample example of the **show ipv6 route** command, as displayed on Router R2.

Example 7-2 Baseline IP Routing Table on Router R2

```

R2#show ipv6 route
IPv6 Routing Table - 11 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B - BGP
       U - Per-user Static route
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
O   A:A:A:A::/64 [110/11]
    via FE80::209:B7FF:FEFA:D1E1, FastEthernet0/0
C   B:B:B:B::/64 [0/0]
    via ::, FastEthernet0/0
L   B:B:B:B::22/128 [0/0]
    via ::, FastEthernet0/0
C   C:C:C:C::/64 [0/0]
    via ::, Serial1/0.1
L   C:C:C:C::2/128 [0/0]
    via ::, Serial1/0.1
C   D:D:D:D::/64 [0/0]

```

```

    via ::, Serial1/0.2
L   D:D:D:D::1/128 [0/0]
    via ::, Serial1/0.2
O   E:E:E:E::/64 [110/128]
    via FE80::C202:8FF:FE98:0, Serial1/0.1
    via FE80::C200:8FF:FE2C:0, Serial1/0.2
O   F:F:F:F::/64 [110/74]
    via FE80::C202:8FF:FE98:0, Serial1/0.1
    via FE80::C200:8FF:FE2C:0, Serial1/0.2
L   FE80::/10 [0/0]
    via ::, Null0
L   FF00::/8 [0/0]
    via ::, Null0

```

As depicted in Example 7-3, a few additional OSPFv3 verification commands are issued to demonstrate the output available in these commands.

Example 7-3 IPv6 and OSPFv6 Verification Commands on Router R2

```

R2#show ipv6 ospf interface serial 1/0.2
Serial1/0.2 is up, line protocol is up
  Link Local Address FE80::C201:8FF:FE2C:0, Interface ID 14
  Area 0, Process ID 1, Instance ID 0, Router ID 10.2.2.2
  Network Type POINT_TO_POINT, Cost: 64
  Transmit Delay is 1 sec, State POINT_TO_POINT,
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:04
  Index 1/2/3, flood queue length 0
  Next 0x0(0)/0x0(0)/0x0(0)
  Last flood scan length is 1, maximum is 2
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 10.4.4.4
  Suppress hello for 0 neighbor(s)
R2#show ipv6 ospf
  Routing Process "ospfv3 1" with ID 10.2.2.2
  It is an area border router
  SPF schedule delay 5 secs, Hold time between two SPFs 10 secs
  Minimum LSA interval 5 secs. Minimum LSA arrival 1 secs
  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 areas in this router is 2. 2 normal 0 stub 0 nssa

```

```
Reference bandwidth unit is 100 mbps
```

```
Area BACKBONE(0)
```

```
Number of interfaces in this area is 2
SPF algorithm executed 8 times
Number of LSA 14. Checksum Sum 0x063F6C
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
SPF algorithm executed 5 times
Number of LSA 11. Checksum Sum 0x0481E4
Number of DCbitless LSA 0
Number of indication LSA 0
Number of DoNotAge LSA 0
Flood list length 0
```

```
R2#show ipv6 ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Interface ID	Interface
10.4.4.4 Serial1/0.2	1	FULL/ -	00:00:36	14	
10.3.3.3 Serial1/0.1	1	FULL/ -	00:00:36	14	
10.1.1.1 FastEthernet0/0	1	FULL/BDR	00:00:39	4	

Finally, in the verification portion of the lab, a series of pings were sent to each router to confirm reachability throughout the topology. All the pings issued in the lab were successful.

Discuss Possible Resolutions to an OSPFv3 Adjacency Issue

Many OSPFv3 troubleshooting issues are similar to OSPFv2 troubleshooting issues, as discussed in Lab 4. Therefore, this *Network Troubleshooting Video Mentor* lab focuses on resolving a specific OSPFv3 issue. The issue addressed is a failure to form or maintain an OSPFv3 adjacency between neighbors. Following is a listing of potential reasons an OSPFv3 adjacency might not be formed:

- Mismatched HELLO parameters
- Mismatched IP MTU setting
- Interface configured as passive
- Mismatched area type

Interpret a Trouble Ticket

The following trouble ticket is presented in this lab:

Company A recently added IPv6 addressing to its existing IPv4 addressing. OSPFv3 is the protocol being used to route the IPv6 traffic.

However, several OSPFv3 adjacencies are not coming up. Full IPv6 reachability throughout the topology needs to be established.

Troubleshoot and Resolve the Identified OSPFv3 Adjacency Issue

To determine the extent of the reported issue, this lab issues the **show ipv6 ospf neighbor** command on each router. Routers R1 and R2 report that they have no OSPF neighbors, whereas routers BB1 and BB2 form an adjacency over their FastEthernet link but not over their Frame Relay link.

Upon close inspection, the following configuration errors are identified:

- Router R2's HELLO timer on the FastEthernet 0/0 is set to a non-default value, whereas the other end of the link is still set to the default.
- Router R2's OSPF network type on subinterface Serial 1/0.2 is set to point-to-multipoint, whereas the other end of the link is set to the default of point-to-point.
- Router R2 has its OSPFv3 process configured with the passive-interface default command, which prevents any of router R2's interfaces from forming OSPFv3 adjacencies.
- Router BB1 has its Serial 1/0.1 subinterface configured for an IPv6 MTU of 1400 bytes, whereas the other end of the link is configured with a default MTU of 1500 bytes.

Example 7-4 shows how router R2's configuration was corrected.

Example 7-4 Correcting Router R2's Configuration Errors

```
R2#conf term
Enter configuration commands, one per line. End with CNTL/Z.
R2(config)#int fa 0/0
R2(config-if)#no ipv6 ospf hello-interval 60
R2(config-if)#exit
R2(config)#int s1/0.2
R2(config-subif)#no ipv6 ospf network point-to-multipoint
R2(config-subif)#exit
R2(config)#ipv6 router no passive-interface default ospf 1
R2(config-rtr)#
```

Example 7-5 shows how router BB1's configuration was corrected.

Example 7-5 Correcting Router BB1's Configuration Errors

```
BB1#conf term
Enter configuration commands, one per line. End with CNTL/Z.
BB1(config)#int s1/0.1
BB1(config-subif)#ipv6 mtu 1500
```

Summarize Key Elements of the OSPFv3 Adjacency Troubleshooting Process

This lab concludes by reviewing common approaches to resolving an OSPFv3 adjacency issue and how the identified issue in this lab was resolved. Specifically, this lab corrects the following misconfigurations on router R2: HELLO interval on interface FastEthernet 0/0, OSPF network type on subinterface Serial 1/0.2, and a **passive-interface default** command configured for the OSPFv3 routing process. Additionally, this lab corrects an incorrect MTU setting for router BB1's Serial 1/0.1 subinterface.

IPv6 and RIPng Troubleshooting

This *Network Troubleshooting Video Mentor* lab demonstrates troubleshooting a RIPng issue in an IPv6-based network. Lab 7, “IPv6 and OSPFv3 Troubleshooting,” reviewed the fundamentals of IPv6, in addition to the fundamentals of OSPFv3. Therefore, this lab does not review IPv6 and instead focuses on RIPng.

Scenario

This lab includes the following steps:

- Step 1** Review RIPng theory.
- Step 2** Examine lab topology.
- Step 3** Identify RIPng verification and troubleshooting commands.
- Step 4** Verify operation of IPv6 and RIPng in lab topology.
- Step 5** Discuss possible symptoms and resolutions of a RIPng issue.
- Step 6** Interpret a trouble ticket.
- Step 7** Troubleshoot and resolve the identified RIPng issue.
- Step 8** Summarize key elements of the RIPng troubleshooting process.

Review RIPng Theory

RIP next generation, better known as *RIPng*, is an enhancement to RIPv2. However, RIPng has several characteristics similar to RIPv2, as follows:

- Distance-vector routing protocol
- Hop count metric
- Maximum hop count of fifteen
- Sends routing updates via multicast
 - RIPv2: 224.0.0.9
 - RIPng: FF02::9

The following characteristics are enhancements of RIPng over RIPv2:

- Supports the routing of 128-bit IPv6 network addresses.
- Link-local addresses used for next-hop addresses.
- Next-hop addresses stored in Routing Information Base (RIB).
- Networks added to a RIP routing process in interface-configuration mode.

Examine Lab Topology

In the lab topology, which is shown in Figure 8-1, all routers have been configured with IPv6 addressing on their physical interfaces and subinterfaces. Routers R1 and R2 are considered to be enterprise routers, whereas routers BB1 and BB2 are considered to be Internet service provider (ISP) routers, with whom the enterprise is dual homed.

The backbone ISP routers are configured to send IPv6 default route advertisements (that is, advertisements for route `::/0`) to the enterprise routers only. Therefore, routers R1 and R2 do not have entries for the `E::E::/64` and `F::F::/64` networks.

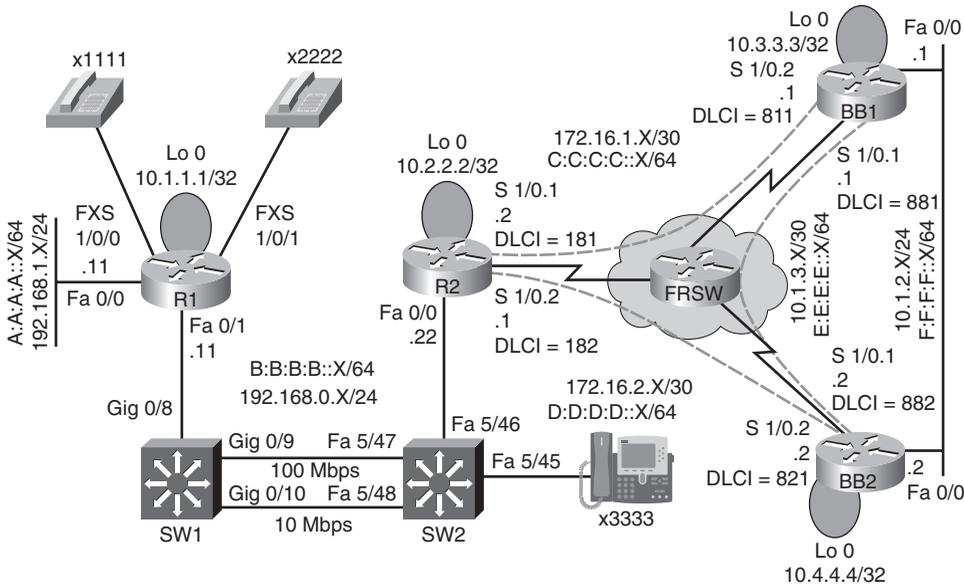


Figure 8-1 Lab 8 Topology

Identify RIPng Verification and Troubleshooting Commands

This lab also illustrates the use of various **show**, **debug**, and RIPng configuration commands to verify and troubleshoot a reported RIPng issue in the lab topology. Table 8-1 presents the syntax for these commands.

Table 8-1 RIPng Verification and Troubleshooting Syntax

Command	Description
Router# show ipv6 rip [<i>process-name</i>] [<i>database</i> <i>next-hops</i>]	Displays information about the specified RIPng routing process, and optionally the contents of the RIPng database and a listing of next-hop addresses.
Router# show ipv6 route	Shows the contents of the IPv6 routing table.
Router# debug ipv6 rip	Provides real-time information about RIPng messages.
Router(config-if)# ipv6 rip process-name enable	Instructs an interface to participate in the specified RIPng routing process.

Table 8-1 RIPng Verification and Troubleshooting Syntax

Command	Description
Router(config-if)# ipv6 rip <i>process-name</i> default-information { only originate }	Causes an interface to originate a default route advertisement (that is, an advertisement for network ::/0) and optionally suppress the advertisement of all other routes.
Router(config)# ipv6 router rip <i>process-name</i>	Enters router configuration mode for the specified RIPng routing process.
Router(config-rtr)# maximum-paths <i>number</i>	Specifies the number of equal-cost paths across which RIPng can load balance (defaults to 16 with a valid range of 1–64).

Verify Operation of IPv6 and RIPng in Lab Topology

All routers in the topology are configured with IPv6 addresses for each of their physical interfaces and subinterfaces. Router R1's configuration, which is representative of the basic IPv6 and RIPng configuration present on all the routers in the topology, is presented in Example 8-1.

Example 8-1 Running Configuration on Router R1

```
R1#show run
...OUTPUT OMITTED...
hostname R1
!
ipv6 unicast-routing
ipv6 cef
!
interface Loopback0
 ip address 10.1.1.1 255.255.255.255
!
interface FastEthernet0/0
 ip address 192.168.1.11 255.255.255.0
 duplex auto
 speed auto
 ipv6 address A:A:A:A::11/64
 ipv6 rip PROCESS1 enable
!
interface FastEthernet0/1
 ip address 192.168.0.11 255.255.255.0
 speed auto
 half-duplex
 ipv6 address B:B:B:B::11/64
 ipv6 rip PROCESS1 enable
!
ipv6 router rip PROCESS1
...OUTPUT OMITTED...
```

Router R1's IPv6 routing table is shown in Example 8-2. Notice the absence of specific routes to the backbone IPv6 networks of E:E:E:E::/64 and F:F:F:F::/64. Instead, router R1 could reach these networks via the default route (::/0) in its routing table.

Example 8-2 IPv6 Routing Table on Router R1

```
R1#show ipv6 route
IPv6 Routing Table - 9 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B - BGP
       U - Per-user Static route
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
R  ::/0 [120/3]
   via FE80::C201:EFF:FE64:0, FastEthernet0/1
C  A:A:A:A::/64 [0/0]
   via ::, FastEthernet0/0
L  A:A:A:A::11/128 [0/0]
   via ::, FastEthernet0/0
C  B:B:B:B::/64 [0/0]
   via ::, FastEthernet0/1
L  B:B:B:B::11/128 [0/0]
   via ::, FastEthernet0/1
R  C:C:C:C::/64 [120/2]
   via FE80::C201:EFF:FE64:0, FastEthernet0/1
R  D:D:D:D::/64 [120/2]
   via FE80::C201:EFF:FE64:0, FastEthernet0/1
L  FE80::/10 [0/0]
   via ::, Null0
L  FF00::/8 [0/0]
   via ::, Null0
```

Router R1's reachability to all IPv6 networks in the topology is confirmed by issuing a series of pings from R1, one to each IPv6 network in the topology. The successful ping results are provided in Example 8-3.

Example 8-3 Ping Results for Router R1

```
R1#ping a:a:a::11

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to A:A:A:A::11, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
R1#ping b:b:b::22

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to B:B:B:B::22, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 20/54/88 ms
R1#ping c:c:c::1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to C:C:C:C::1, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 48/89/128 ms
R1#ping d:d:d::2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to D:D:D:D::2, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 52/97/173 ms
R1#ping e:e:e::1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to E:E:E:E::1, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 60/85/132 ms
R1#ping f:f:f::2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to F:F:F:F::2, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 64/145/337 ms
```

Router R2's routing table, provided in Example 8-4, shows that router R2 has two paths to reach the default network of ::0: one path via router BB1 and one path via router BB2. These dual paths are made possible because RIPng (in addition to RIPv1 and RIPv2) can load balance across equal-cost paths. All versions of RIP use *hop count* as their metrics.

Example 8-4 IPv6 Routing Table on Router R2

```
R2#show ipv6 route
IPv6 Routing Table - 10 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B - BGP
       U - Per-user Static route
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
R   ::/0 [120/2]
    via FE80::C202:EFF:FEBC:0, Serial1/0.1
    via FE80::C200:EFF:FE64:0, Serial1/0.2
R   A:A:A:A::/64 [120/2]
    via FE80::209:B7FF:FEFA:D1E1, FastEthernet0/0
C   B:B:B:B::/64 [0/0]
    via ::, FastEthernet0/0
L   B:B:B:B::22/128 [0/0]
    via ::, FastEthernet0/0
C   C:C:C:C::/64 [0/0]
    via ::, Serial1/0.1
L   C:C:C:C::2/128 [0/0]
    via ::, Serial1/0.1
C   D:D:D:D::/64 [0/0]
    via ::, Serial1/0.2
L   D:D:D:D::1/128 [0/0]
    via ::, Serial1/0.2
L   FE80::/10 [0/0]
    via ::, Null0
L   FF00::/8 [0/0]
    via ::, Null0
```

Routers BB1 and BB2 both have listings for all of the topology's IPv6 networks in their IPv6 routing tables. These routing tables are presented in Examples 8-5 and 8-6.

Example 8-5 IPv6 Routing Table on Router BB1

```
BB1#show ipv6 route
IPv6 Routing Table - 11 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B - BGP
       U - Per-user Static route
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
R   A:A:A:A::/64 [120/3]
    via FE80::C201:EFF:FE64:0, Serial1/0.2
R   B:B:B:B::/64 [120/2]
    via FE80::C201:EFF:FE64:0, Serial1/0.2
C   C:C:C:C::/64 [0/0]
    via ::, Serial1/0.2
L   C:C:C:C::1/128 [0/0]
    via ::, Serial1/0.2
R   D:D:D:D::/64 [120/2]
    via FE80::C200:EFF:FE64:0, FastEthernet0/0
    via FE80::C200:EFF:FE64:0, Serial1/0.1
    via FE80::C201:EFF:FE64:0, Serial1/0.2
C   E:E:E:E::/64 [0/0]
    via ::, Serial1/0.1
L   E:E:E:E::1/128 [0/0]
    via ::, Serial1/0.1
C   F:F:F:F::/64 [0/0]
    via ::, FastEthernet0/0
L   F:F:F:F::1/128 [0/0]
    via ::, FastEthernet0/0
L   FE80::/10 [0/0]
    via ::, Null0
L   FF00::/8 [0/0]
    via ::, Null0
```

Example 8-6 IPv6 Routing Table on Router BB2

```

BB2#show ipv6 route
IPv6 Routing Table - 11 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B - BGP
       U - Per-user Static route
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
R  A:A:A:A::/64 [120/3]
   via FE80::C201:EFF:FE64:0, Serial1/0.2
R  B:B:B:B::/64 [120/2]
   via FE80::C201:EFF:FE64:0, Serial1/0.2
R  C:C:C:C::/64 [120/2]
   via FE80::C201:EFF:FE64:0, Serial1/0.2
   via FE80::C202:EFF:FEBC:0, FastEthernet0/0
   via FE80::C202:EFF:FEBC:0, Serial1/0.1
C  D:D:D:D::/64 [0/0]
   via ::, Serial1/0.2
L  D:D:D:D::2/128 [0/0]
   via ::, Serial1/0.2
C  E:E:E:E::/64 [0/0]
   via ::, Serial1/0.1
L  E:E:E:E::2/128 [0/0]
   via ::, Serial1/0.1
C  F:F:F:F::/64 [0/0]
   via ::, FastEthernet0/0
L  F:F:F:F::2/128 [0/0]
   via ::, FastEthernet0/0
L  FE80::/10 [0/0]
   via ::, Null0
L  FF00::/8 [0/0]
   via ::, Null0

```

Notice that the IPv6 routing tables for routers BB1 and BB2 do not have a default route entry. The absence of the default route is because routers BB1 and BB2 are sourcing default route information to router R2, while suppressing more specific route information to router R2. Example 8-7 shows the **ipv6 rip process-name default-information originate** command that both routers BB1 and BB2 use to accomplish this default route advertisement.

Example 8-7 Advertising Default Route Information on Router BB1

```
BB1#show run | begin Serial1/0.2
interface Serial1/0.2 point-to-point
 ip address 172.16.1.1 255.255.255.252
 ipv6 address C:C:C:C::1/64
 ipv6 rip PROCESS1 enable
 ipv6 rip PROCESS1 default-information only
 frame-relay interface-dlci 811
```

Discuss Possible Symptoms and Resolutions of a RIPng Issue

This lab identifies the following as common symptoms of a RIPng issue. Additionally, this lab discusses common resolutions and troubleshooting tips for these symptoms:

- RIPng routes not appearing in the IPv6 routing table.
- RIPng not performing appropriate load balancing.
- Interface not sending RIPng updates.
- Individual network advertisements not suppressed when sourcing a default route.

Interpret a Trouble Ticket

The following trouble ticket is presented in this lab:

Company A (that is, routers R1 and R2) is dual-homed to two Internet Service Providers (ISPs). The ISP routers are BB1 and BB2. However, router R2 only sees a single path to reach a default route (rather than one path from each ISP) in its IPv6 routing table. Also, router R2 is seeing other ISP-advertised routes (specifically, E:E:E:E::/64 and F:F:F:F::/64) rather than just a default route in its IPv6 routing table. All routes router R2 receives from the ISP routers, except a default route, should be suppressed.

Troubleshoot and Resolve the Identified RIPng Adjacency Issue

The **show ipv6 route** command is issued on router R2 to confirm that the IPv6 routing table includes only a single path to reach the default network of ::/0. The output from this command, which also confirms the presence of the backbone routes E:E:E:E::/64 and F:F:F:F::/64 in the IPv6 routing table, is provided in Example 8-8.

Example 8-8 Confirmation of Troubleshooting Issues on Router R2

```
R2#show ipv6 route
IPv6 Routing Table - 12 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B - BGP
       U - Per-user Static route
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
R   ::/0 [120/2]
    via FE80::C200:EFF:FE64:0, Serial1/0.2
R   A:A:A:A::/64 [120/2]
    via FE80::209:B7FF:FEFA:D1E1, FastEthernet0/0
C   B:B:B:B::/64 [0/0]
    via ::, FastEthernet0/0
L   B:B:B:B::22/128 [0/0]
    via ::, FastEthernet0/0
C   C:C:C:C::/64 [0/0]
    via ::, Serial1/0.1
L   C:C:C:C::2/128 [0/0]
    via ::, Serial1/0.1
C   D:D:D:D::/64 [0/0]
    via ::, Serial1/0.2
L   D:D:D:D::1/128 [0/0]
    via ::, Serial1/0.2
R   E:E:E:E::/64 [120/2]
    via FE80::C200:EFF:FE64:0, Serial1/0.2
R   F:F:F:F::/64 [120/2]
    via FE80::C200:EFF:FE64:0, Serial1/0.2
L   FE80::/10 [0/0]
    via ::, Null0
L   FF00::/8 [0/0]
    via ::, Null0
```

The **show ipv6 rip database** command, as shown in Example 8-9, proves that router R2 received two default route advertisements. However, only one of those route advertisements is injected into the IPv6 routing table.

Example 8-9 RIP Database on Router R2

```

R2#show ipv6 rip database
RIP process "PROCESS1", local RIB
  A:A:A:A::/64, metric 2, installed
    FastEthernet0/0/FE80::209:B7FF:FEFA:D1E1, expires in 174 secs
  B:B:B:B::/64, metric 2
    FastEthernet0/0/FE80::209:B7FF:FEFA:D1E1, expires in 174 secs
  D:D:D:D::/64, metric 2
    Serial1/0.2/FE80::C200:EFF:FE64:0, expires in 160 secs
  E:E:E:E::/64, metric 2, installed
    Serial1/0.2/FE80::C200:EFF:FE64:0, expires in 160 secs
  F:F:F:F::/64, metric 2, installed
    Serial1/0.2/FE80::C200:EFF:FE64:0, expires in 160 secs
  ::/0, metric 2, installed
    Serial1/0.2/FE80::C200:EFF:FE64:0, expires in 160 secs
    Serial1/0.1/FE80::C202:EFF:FEBC:0, expires in 170 secs

```

A review of router R2's running configuration reveals the **maximum-paths 1** command in router configuration mode for the RIPng routing process. This command prevented two default route paths from appearing in the IPv6 routing table. Example 8-10 shows how this command was removed from router R2's configuration in order to restore load balancing to the default route.

Example 8-10 Restoring Load Balancing to a Default Route on Router R2

```

R2#conf term
Enter configuration commands, one per line. End with CNTL/Z.
R2(config)#ipv6 router rip PROCESS1
R2(config-rtr)#no maximum-paths 1

```

The **debug ipv6 rip** command is issued on router R2 to see if router BB2 is sending both default route information and specific route information. The output from this command, as presented in Example 8-11, confirms that router BB2 was not suppressing specific route information.

Example 8-11 Debugging RIPng Traffic on Router R2

```

R2#debug ipv6 rip
...OUTPUT OMITTED...
*Mar 1 00:33:30.747: RIPng: response received from FE80::C200:EFF:FE64:0 on
Ser
ial1/0.2 for PROCESS1
*Mar 1 00:33:30.751:          src=FE80::C200:EFF:FE64:0 (Serial1/0.2)
*Mar 1 00:33:30.751:          dst=FF02::9
*Mar 1 00:33:30.755:          sport=521, dport=521, length=92
*Mar 1 00:33:30.755:          command=2, version=1, mbz=0, #rte=4
*Mar 1 00:33:30.755:          tag=0, metric=1, prefix=F:F:F:F::/64
*Mar 1 00:33:30.755:          tag=0, metric=1, prefix=E:E:E:E::/64
*Mar 1 00:33:30.755:          tag=0, metric=1, prefix=D:D:D:D::/64
*Mar 1 00:33:30.755:          tag=0, metric=1, prefix=::/0
...OUTPUT OMITTED...

```

The focus of this lab's troubleshooting efforts then shifts to router BB2. An inspection of router BB2's running configuration reveals the **ipv6 rip PROCESS1 default-information originate** command under subinterface configuration mode for Serial 1/0.2. The **originate** option at the end of this command sources a default router advertisement, but it does not suppress the sending of more specific routes. Example 8-12 shows how this configuration was changed to use the **only** parameter. The **only** parameter causes the interface to originate only default route information, while suppressing the more specific routes.

Example 8-12 Suppressing Specific Route Information on Router BB2's Serial 1/0.2 Subinterface

```

BB2#conf term
Enter configuration commands, one per line. End with CNTL/Z.
BB2(config)#int s1/0.2
BB2(config-subif)#ipv6 rip PROCESS1 default-information originate only

```

After giving the E:E:E:E::/64 and F:F:F:F::/64 routes sufficient time to timeout of router R2's IPv6 routing table, the **show ipv6 route** was issued once again. The output, as shown in Example 8-13, confirms that the issues reported in the trouble ticket are resolved.

Example 8-13 Router R2's IPv6 Routing Table After Troubleshooting

```

R2#show ipv6 route
IPv6 Routing Table - 10 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B - BGP
       U - Per-user Static route
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
R   ::/0 [120/2]
    via FE80::C200:EFF:FE64:0, Serial1/0.2
    via FE80::C202:EFF:FEBC:0, Serial1/0.1
R   A:A:A:A::/64 [120/2]
    via FE80::209:B7FF:FEFA:D1E1, FastEthernet0/0
C   B:B:B:B::/64 [0/0]
    via ::, FastEthernet0/0
L   B:B:B:B::22/128 [0/0]
    via ::, FastEthernet0/0
C   C:C:C:C::/64 [0/0]
    via ::, Serial1/0.1
L   C:C:C:C::2/128 [0/0]
    via ::, Serial1/0.1
C   D:D:D:D::/64 [0/0]
    via ::, Serial1/0.2
L   D:D:D:D::1/128 [0/0]
    via ::, Serial1/0.2
L   FE80::/10 [0/0]
    via ::, Null0
L   FF00::/8 [0/0]
    via ::, Null0

```

Summarize Key Elements of the RIPng Troubleshooting Process

This lab concludes by discussing RIPng troubleshooting considerations and reviewing how this lab resolves the reported issues. Specifically, this lab corrects a RIPng load-balancing misconfiguration and a misconfiguration that failed to suppress specific route information while sourcing default route information.

Cisco IOS Security Troubleshooting

Cisco IOS offers a wide array of security solutions. These solutions can often complement additional Cisco security solutions (for example, Cisco's Adaptive Security Appliance (ASA), which can provide firewall, virtual private network (VPN), and intrusion prevention system (IPS) services).

Because troubleshooting Cisco IOS security is a study in itself, this *Network Troubleshooting Video Mentor* lab reviews, at a high level, the functions of major Cisco IOS security solutions. Common Cisco IOS troubleshooting targets are discussed. Then, in the live interface, this lab demonstrates tips for recovering from a forgotten password and having a connection time out too quickly. Because access control list (ACL) configuration is an extremely common area of security misconfiguration, this lab concludes by troubleshooting an ACL issue.

Scenario

This lab includes the following steps:

- Step 1** Review Cisco IOS security features.
- Step 2** Examine lab topology.
- Step 3** Discuss examples of potential Cisco IOS security issues.
- Step 4** Interpret a trouble ticket.
- Step 5** Identify a sampling of Cisco IOS security verification and troubleshooting commands.
- Step 6** Troubleshoot and resolve the identified Cisco IOS security issues.
- Step 7** Summarize potential Cisco IOS security troubleshooting targets.

Review Cisco IOS Security Features

Cisco IOS includes the following security features:

- **Router password protection:** As illustrated in Figure 9-1, access to a router's prompt can be secured through password protection. Passwords can be applied to a router's console port (commonly used for connectivity with a terminal), auxiliary port (commonly used for connectivity with a modem), and VTY lines (commonly used for connectivity via a Telnet or a Secure Shell (SSH) connection).

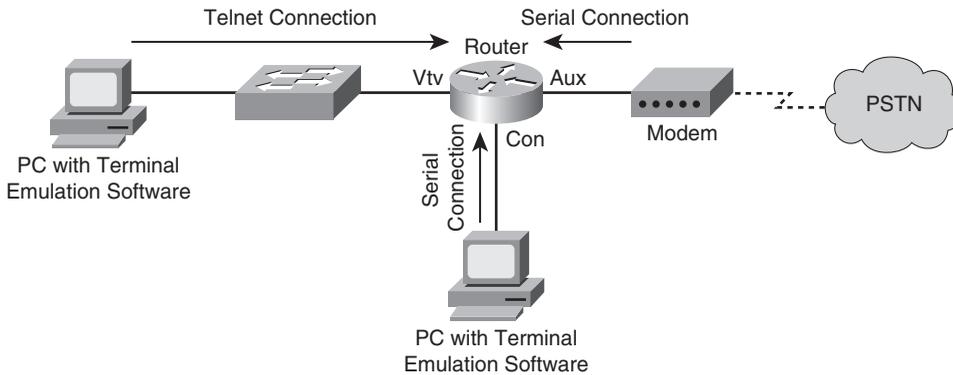


Figure 9-1 Router Password Protection

- Authentication, Authorization, and Accounting (AAA):** Although username/password information can be stored locally on a router, for scalability and for increased granularity in the assigning of privilege levels, an AAA server can be used to store username/password information. AAA can be used not just for permitting connections into a router, but also to permit connections passing through a router, as shown in Figure 9-2.

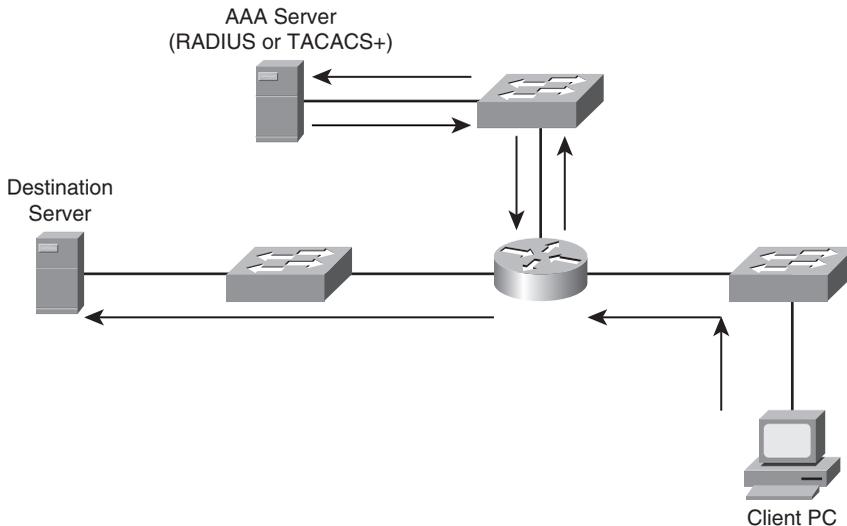


Figure 9-2 AAA Services

- Cisco IOS firewall:** A router running the Cisco IOS firewall feature set can provide stateful firewall features. A stateful firewall can, for example, inspect sessions originating on the inside of a company's network destined for a device outside the company's network, such that return traffic for that session can reenter the company's network. However, if that same outside device had initiated a session destined for the company's network, the stateful firewall could deny that traffic, as depicted in Figure 9-3.

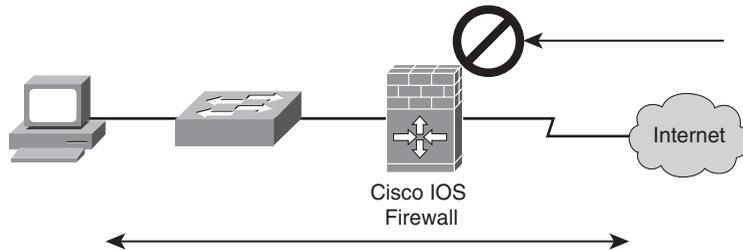


Figure 9-3 Cisco IOS Firewall

- **ACL:** An ACL contains a set of rules, which dictate permitted and denied traffic. The rules within an ACL are processed top-down, and an implicit rule resides at the bottom of an ACL that denies all traffic not explicitly permitted in the ACL.
- **IPS:** Although Cisco offers dedicated IPS hardware, Cisco IOS routers can also be configured to provide threat prevention via IPS. A router's flash can store an .SDF (Signature Definition File) file, which contains the signatures of well-known security threats. As illustrated in Figure 9-4, a router running IPS services might be used to complement one or more additional security solutions in a network.

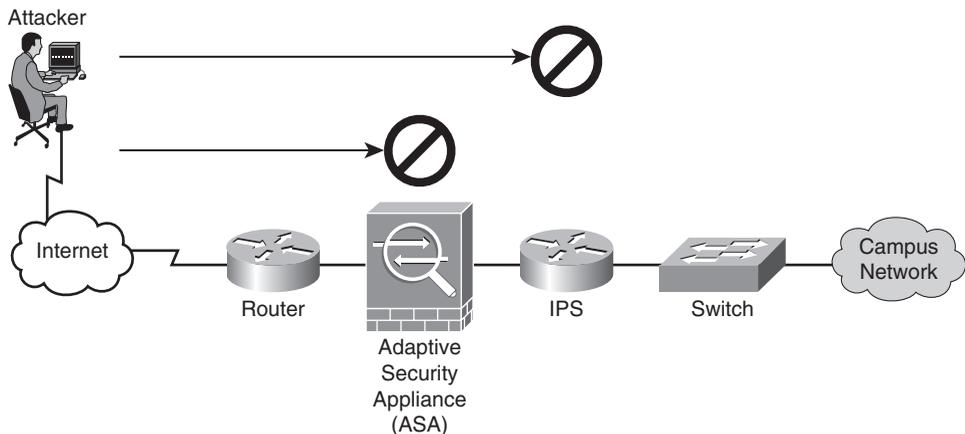


Figure 9-4 Cisco IOS IPS

- **IPsec VPN:** A VPN creates a logical tunnel through another network (for example, the Internet). Traffic flowing across that logical tunnel can be protected using IP Security (IPsec), as shown in Figure 9-5. For example, if an attacker were to use a packet-capture utility to view packets flowing across an IPsec VPN that was providing encryption services, the attacker would not be able to interpret the captured packets.

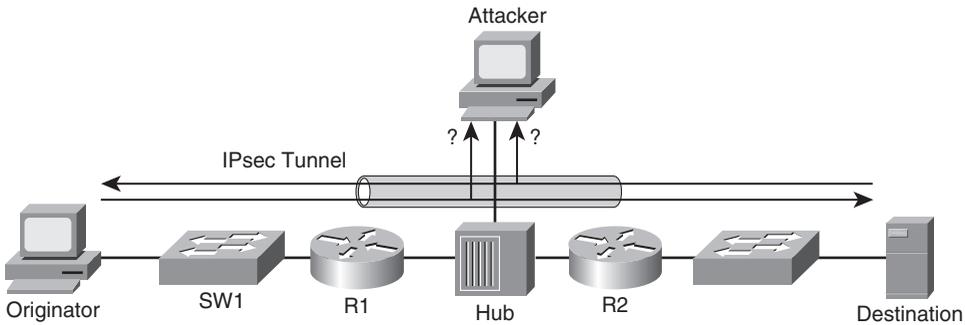


Figure 9-5 IPsec VPN

Examine Lab Topology

In the lab topology, which is shown in Figure 9-6, all routers have been configured to have all of their interfaces participate in OSPF Area 0. Therefore, all routers have full reachability throughout the network. Routers R1 and R2 are considered to be enterprise routers, whereas routers BB1 and BB2 are considered to be Internet Service Provider (ISP) routers.

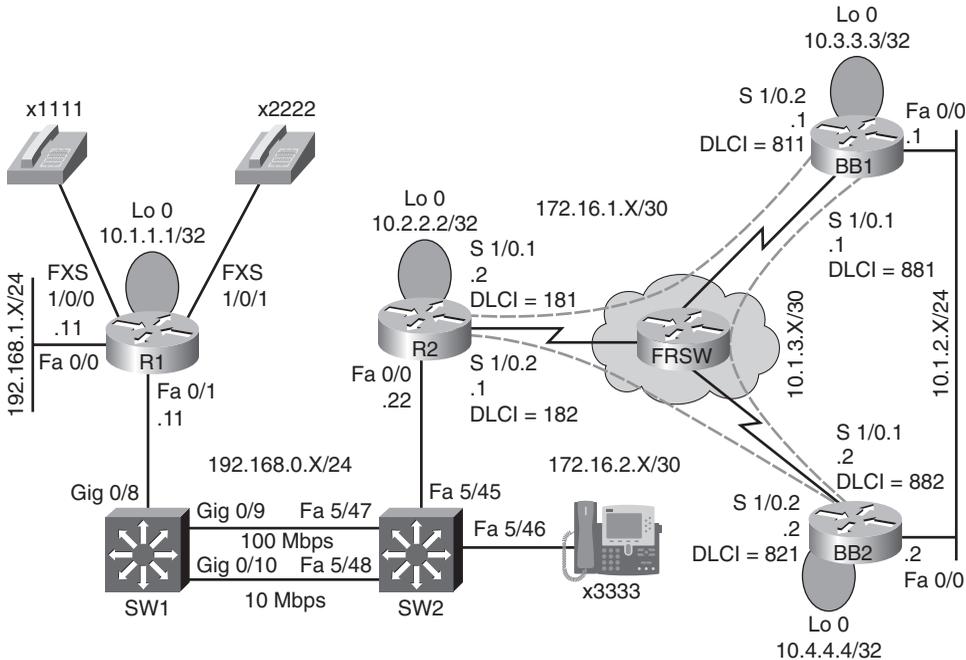


Figure 9-6 Lab 6 Topology

Discuss Examples of Potential Cisco IOS Security Issues

This lab discusses troubleshooting the following Cisco IOS issues:

- Forgetting a router's enable password
- Setting the **exec-timeout** to too short of a time period
- Traffic blocked due to incorrect access lists
- Inappropriate traffic crossing an IPsec VPN
- Malicious traffic not detected by a router's IPS feature
- Cisco IOS firewall not permitting appropriate traffic
- Unable to log in to a router because the AAA server is offline

Interpret a Trouble Ticket

The following trouble ticket is presented in this lab:

A new administrator for Company A has forgotten the enable secret password assigned to router R1 and can no longer log in. Also, when this administrator connects to router R2 via Telnet, the connection is timed out after only one second. The administrator reports that this short timeout does not give him sufficient time to correct the configuration. Also, the administrator configured an access list on router R2 to prevent anyone on the backbone (that is, connections coming into router R2 via the Frame Relay network) from connecting to the Loopback 0 interfaces on routers R1 or R2 via Telnet. However, the access list does not seem to be working.

Identify a Sampling of Cisco IOS Security Verification and Troubleshooting Commands

Because a plethora of Cisco IOS security troubleshooting and configuration commands exist, this lab illustrates verification and configuration commands relevant to the presented trouble ticket. Table 9-1 presents the syntax for these commands.

Table 9-1 Sampling of Cisco IOS Security Troubleshooting Syntax

Command	Description
Router(config-line)# exec-timeout <i>minutes</i> [<i>seconds</i>]	Specifies how long the EXEC process running on a line waits for user input before timing out the connection (defaults to 10 minutes).
Router(config)# access-list <i>number</i> { deny permit } <i>protocol source wildcard-mask destination wildcard-mask</i> [eq port-number] [log]	Creates an extended IP access list, where the access list number is in the range 100–199.
Router# show access-lists	Displays access lists configured on a router.
Router# show logging	Shows output collected from logged access list entries.

Troubleshoot and Resolve the Identified Cisco IOS Security Issues

The trouble ticket identified three issues, as follows:

- A forgotten enable secret password
- An **exec-timeout** parameter set too low
- An ACL misconfiguration

Password Recovery

The first issue addressed in this lab is password recovery. The administrator reportedly forgot the enable secret password on his Cisco 2611XM router (that is, router R1). The router was rebooted, and during the first few seconds of the router booting up, a **Break** was sent from the terminal emulator to the router. The **Break** caused the ROM Monitor prompt (that is, *rommon*) to appear on router R1's console.

The configuration register was set to 0x2142 with the command **confreg 0x2142**. Setting the configuration register to this value causes the router to ignore the startup configuration when the router boots. The router was then rebooted by issuing the **reset** command at the *rommon* prompt.

Because the router ignored the startup configuration, after the router booted, a prompt was presented, asking the administrator if he wished to go through the setup dialog. A **no** was entered at this prompt. The **enable** command was entered to go into privileged configuration mode. From privileged mode, the startup configuration, stored in the router's NVRAM, was merged with the existing running configuration using the command **copy star run**. This command does not *replace* the running configuration with the startup configuration. Rather, these two configurations are *merged*. After this merger, all the physical interfaces were administratively shut down. Therefore, the **no shutdown** command was entered for interfaces FastEthernet 0/0 and FastEthernet 0/1.

The enable secret password was reset to *cisco* using the command **enable secret cisco**. Next, the configuration register was set back to its normal value of 0x2102 with the command **config-register 0x2102**. The running configuration was copied to the startup configuration with the command **copy run star**. The router was then rebooted with the **reload** command.

Incorrect exec-timeout Parameter Recovery

The second issue addressed in this lab was recovering from a misconfiguration on router R2, which caused a Telnet session to timeout after only one second of inactivity. The challenge with such a misconfiguration is that when an administrator Telnets to the router to correct the configuration, he might be logged out if he pauses for as little as a single second.

The fix demonstrated in this lab was to continuously tap on the keyboard's down arrow with one hand, while using the other hand to enter the commands required to correct the **exec-timeout** misconfiguration. Example 9-1 shows the commands entered to set the **exec-timeout** parameter such that a Telnet session never times out.

Example 9-1 Correcting an exec-timeout Misconfiguration

```
R2#conf term
R2(config)#line vty 0 4
R2(config-line)#exec-timeout 0 0
```

ACL Configuration Correction

This lab's final troubleshooting issue was an ACL misconfiguration. The goal of the ACL on router R2 was to prevent Telnet traffic coming in from the backbone (that is, coming in over subinterfaces Serial 1/0.1 or Serial 1/0.2) destined for the loopback interface on router R1 or R2 (that is, IP addresses 10.1.1.1 or 10.2.2.2).

Upon examination, the ACL (an extended IP ACL numbered 100) on router R2 appeared to be configured correctly. The output of the **show access-list** command is provided in Example 9-2.

Example 9-2 Output of the show access-list Command on Router R2

```
R2#show access-list
Extended IP access list 100
 10 deny tcp any host 10.1.1.1 eq telnet
 20 deny tcp any host 10.2.2.2 eq telnet
 30 permit ip any any (116 matches)
```

However, ACL 100 was applied in the outbound direction on router R2's Frame Relay subinterfaces. This lab replaced the incorrect **ip access-group** commands, as shown in Example 9-3.

Example 9-3 Correcting the Application of ACL 100 on Router R2

```
R2#conf term
R2(config)#interface s1/0.1
R2(config-if)#no ip access-group 100 out
R2(config-if)#ip access-group 100 in
R2(config-if)#interface s1/0.2
R2(config-if)#no ip access-group 100 out
R2(config-if)#ip access-group 100 in
```

Summarize Potential Cisco IOS Security Troubleshooting Targets

This lab concludes by discussing common Cisco IOS security troubleshooting targets and reviewing how this lab resolved the reported issues. Specifically, this lab demonstrates how to perform password recovery, how to correct an **exec-timeout** parameter set to only one second, and how to correctly apply an ACL.

DHCP Troubleshooting

This *Network Troubleshooting Video Mentor* lab focuses on troubleshooting Dynamic Host Configuration Protocol (DHCP). DHCP serves as one of the most common methods of assigning IP address information to a network host. Specifically, DHCP allows a DHCP client to obtain an IP address, subnet mask, default gateway IP address, DNS server IP address, and other types of IP address information from a DHCP server.

If you have a cable modem or DSL connection in your home, your cable modem or DSL router might obtain its IP address from your service provider via DHCP. In many corporate networks, when a PC boots up, that PC receives its IP address configuration information from a corporate DHCP server.

Cisco IOS routers can act as a DHCP server. Also, some router interfaces can obtain their IP address information from a DHCP server. Although some troubleshooting issues might stem from misconfiguration, a major DHCP troubleshooting target results from DHCP's use of broadcast messages. Specifically, because broadcast traffic cannot, by default, pass through a router, a DHCP message intended to discover DHCP servers on a network is constrained to the DHCP client's subnet.

Scenario

This lab includes the following steps:

- Step 1** Review DHCP theory.
- Step 2** Examine lab topology.
- Step 3** Identify DHCP verification and troubleshooting commands.
- Step 4** Discuss possible symptoms and resolutions of a DHCP issue.
- Step 5** Interpret a trouble ticket.
- Step 6** Troubleshoot and resolve the identified DHCP issue.
- Step 7** Summarize key elements of the DHCP troubleshooting process.

Review DHCP Theory

Dynamic Host Configuration Protocol (DHCP) allows a client to obtain an IP address, subnet mask, default gateway IP address, DNS server IP address, and other types of information from a server. This collection of IP address information is known as a *DHCP lease*. A lease is obtained through the exchange of four messages between a DHCP client and DHCP server. The acronym of *DORA* might help you remember these four DHCP messages, as follows:

- DHCPDISCOVER
- DHCPOFFER
- DHCPREQUEST
- DHCPACK

The exchange of these messages is demonstrated in Figure 10-1.

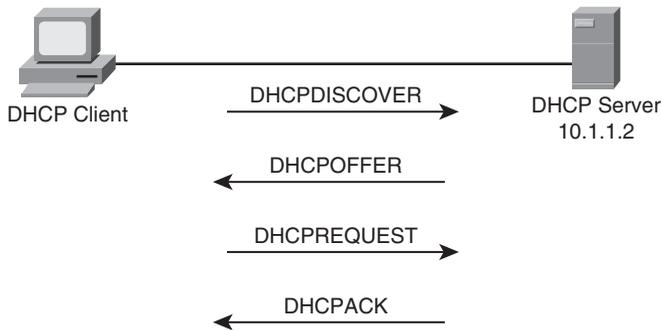


Figure 10-1 DHCP Messages Exchanged Between a DHCP Client and a DHCP Server

By default, routers do not forward broadcasts. This confines a DHCPDISCOVER message, which is sent as a broadcast, to a DHCP client's local subnet, as illustrated in Figure 10-2.

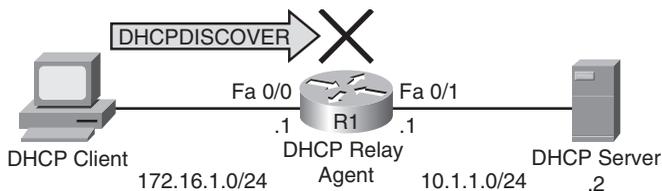


Figure 10-2 Router Blocking Broadcast Traffic

Fortunately, you can add configuration to a router interface, off of which a DHCP client is attached, to permit some types of broadcast traffic (including DHCP broadcast traffic) to be forwarded to either a specific IP address or to a specific subnet. The **ip helper-address** *ip-address* command, issued in interface configuration mode, allows this selective forwarding of broadcast traffic. Example 10-1 provides a sample configuration, which allows the DHCP client in Figure 10-2 to have its DHCPDISCOVER message forwarded to a DHCP server at an IP address of 10.1.1.2.

Example 10-1 Sample Configuration Using the ip helper-address Command

```

Router(config)#interface fa 0/0
Router(config-if)#ip address 172.16.1.1 255.255.255.0
Router(config-if)#ip helper-address 10.1.1.2
  
```

After the addition of this configuration, the DHCPDISCOVER message can be forwarded to a DHCP server, as seen in Figure 10-3.

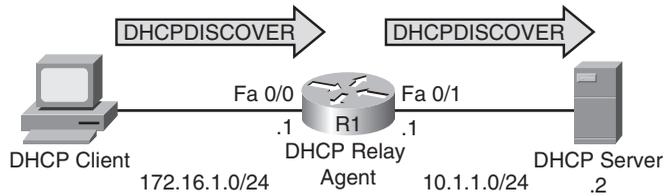


Figure 10-3 A DHCPDISCOVER Message Forwarded by a Router

Examine Lab Topology

In the lab topology, which is shown in Figure 10-4, all routers have been configured to have full reachability to one another, using OSPF as the routing protocol. An IP phone with a directory number of 3333 acts as a DHCP client, with router BB1 acting as a DHCP server.

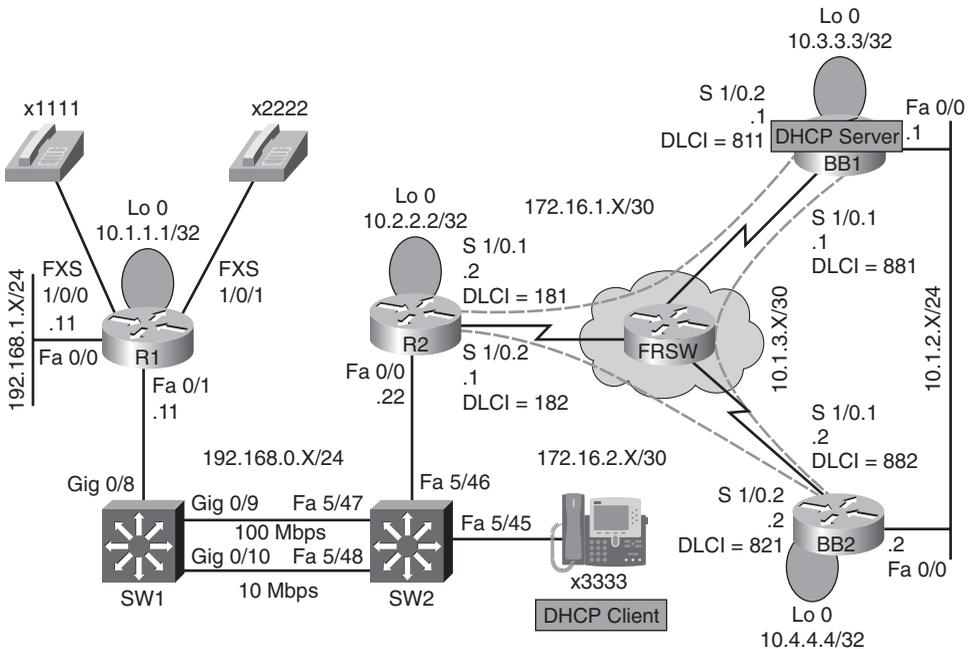


Figure 10-4 Lab 10 Topology

Identify DHCP Verification and Troubleshooting Commands

Effectively troubleshooting a DHCP issue might require knowledge of Cisco IOS DHCP configuration commands. Therefore, this lab presents a collection of commands, which can be used to configure, verify, and troubleshoot potential DHCP issues. Table 10-1 presents the syntax for these commands.

Table 10-1 DHCP Configuration, Verification, and Troubleshooting Syntax

Command	Description
Router# show ip dhcp binding	Displays IP addresses assigned by an IOS DHCP server, the corresponding MAC addresses, and the lease expiration.
Router# show ip dhcp pool	Shows information about a router's DHCP pools, including the pools' subnets and utilization levels.
Router# debug ip dhcp server events	Provides real-time information about DHCP address assignments and database updates.
Router# debug ip dhcp server packet	Displays real-time decodes of DHCP packets.
Router(config-if)# ip helper-address ip-address	Causes an interface to forward received UDP broadcasts to the destination IP address, which could be either a specific IP address or a directed broadcast.
Router(config)# ip dhcp excluded-address beginning-ip-address [ending-ip-address]	Specifies a range of IP addresses not to be assigned to DHCP clients.
Router(config)# ip dhcp pool pool-name	Creates a DHCP pool.
Router(dhcp-config)# network network-address subnet-mask	Identifies a subnet to be used by a DHCP pool.
Router(dhcp-config)# default-router ip-address	Specifies the IP address of a default gateway to be given to a DHCP client.
Router(dhcp-config)# dns-server ip-address	Configures the IP address of a DNS server to be given to a DHCP client.
Router(dhcp-config)# netbios-name-server ip-address	Defines the IP address of a WINS server to be given to a DHCP client.
Router(dhcp-config)# lease {days hours minutes infinite}	Determines the duration of a DHCP lease given to a DHCP client.
Router(dhcp-config)# option 150 ip ip-address	Specifies the IP address of a TFTP server given a DHCP client.
Router(config-if)# ip address dhcp	Tells an interface to obtain its IP address via DHCP.

Discuss Possible Symptoms and Resolutions of a DHCP Issue

This lab discusses troubleshooting the following DHCP issues:

- The DHCP client cannot obtain an IP address from a DHCP server located on the same subnet.
- The DHCP client cannot obtain an IP address from a DHCP server located on a different subnet.
- A POOL EXHAUSTED error message is received from an IOS router acting as a DHCP server.
- IP address assigned by DHCP server overlaps with an existing IP address.
- The DHCP pool configured for an interface's secondary network address is not being used to assign IP addresses to DHCP clients.

Interpret a Trouble Ticket

The following trouble ticket is presented in this lab:

Router BB1 is acting as a DHCP server. However, an IP phone (x3333) is failing to obtain a DHCP lease.

Troubleshoot and Resolve the Identified DHCP Issue

This lab begins by confirming that the IP phone has not obtained a DHCP lease from router BB1, which acts as the topology's DHCP server. Specifically, the **show ip dhcp binding** command is issued on router BB1 and reveals that no leases have been handed out. Example 10-2 shows the output from this command.

Example 10-2 Viewing Router BB1's DHCP Bindings

```
BB1#show ip dhcp binding
Bindings from all pools not associated with VRF:
IP address          Client-ID/                Lease expiration        Type
                   Hardware address/
                   User name
```

Router BB1's running configuration is next examined, focusing on the DHCP configuration commands shown in Example 10-3.

Example 10-3 Viewing Router BB1's DHCP Configuration

```
BB1#show run
...OUTPUT OMITTED...
ip dhcp excluded-address 192.168.0.1 192.168.0.100
!
ip dhcp pool TSHOOT
  network 192.168.1.0 255.255.255.0
  option 150 ip 10.1.1.1
  default-router 192.168.0.22
...OUTPUT OMITTED...
```

The **network** command, as highlighted in Example 10-3, specifies an incorrect network address. The network specified should be 192.168.0.0/24 rather than 192.168.1.0/24. This misconfiguration is corrected, as seen in Example 10-4.

Example 10-4 Correcting Router BB1's Misconfiguration

```

BB1#conf term
Enter configuration commands, one per line. End with CNTL/Z.
BB1(config)#ip dhcp pool TSH00T
BB1(dhcp-config)#no network 192.168.1.0 255.255.255.0
BB1(dhcp-config)#network 192.168.0.0 255.255.255.0
BB1(dhcp-config)#end
BB1#

```

Even after correcting router BB1's misconfiguration, the IP phone still fails to obtain a DHCP lease, as evidenced by the issuance of another **show ip dhcp binding** command, as seen in Example 10-5.

Example 10-5 Checking Router BB1's DHCP Bindings after Correcting Router BB1's Misconfiguration

```

BB1#show ip dhcp binding
Bindings from all pools not associated with VRF:
IP address           Client-ID/                Lease expiration        Type
                   Hardware address/
                   User name

```

This lab's troubleshooting focus then shifts to router R2, which should be acting as the default gateway for the IP phone attempting to obtain a DHCP lease. Example 10-6 shows the portion of router R2's running configuration containing configuration information for interface FastEthernet 0/0, the interface assigned the IP address that the IP phone should point to as its default gateway.

Example 10-6 Examining Router R2's Running Configuration

```

R2#show run | begin interface FastEthernet0/0
interface FastEthernet0/0
 ip address 192.168.0.22 255.255.255.0
 duplex auto
 speed auto

```

This lab observes that router R2's configuration for interface FastEthernet 0/0 lacks an **ip helper-address ip-address** command, which would permit a DHCP broadcast to be forwarded through router R2 and be sent to a specified IP address or a specified network address. Example 10-7 demonstrates the addition of this command on router R2.

Example 10-7 Adding the ip helper-address Command to Router R2's FastEthernet 0/0's Interface

```

R2#conf term
Enter configuration commands, one per line. End with CNTL/Z.
R2(config)#int fa0/0
R2(config-if)#ip helper-address 10.3.3.3
R2(config-if)#end
R2#

```

After adding an IP helper address to router R2, the IP phone is able to obtain an IP address from router BB1, which is acting as a DHCP server. Example 10-8 confirms the lease assignment, as seen in Example 10-8. Specifically, the output shows that an IP address of 192.168.0.101 is assigned to the IP phone.

Example 10-8 Viewing the Lease Assigned by Router BB1 to the IP Phone

```

BB1#show ip dhcp binding
Bindings from all pools not associated with VRF:
IP address          Client-ID/
                   Hardware address/
                   User name
192.168.0.101      0100.08a3.b895.c4
                   Mar 02 2002 01:16 AM
                   Automatic

```

Summarize Key Elements of the DHCP Troubleshooting Process

This lab concludes by discussing common DHCP troubleshooting targets and reviewing how this lab resolves the reported issue. This lab resolves two DHCP configuration issues, as follows:

- Correcting the **network** *network-address subnet-mask* command on router BB1.
- Adding an appropriate **ip helper-address** *ip-address* command on router R2.

NAT Troubleshooting

Some IP addresses are routable through the public Internet, whereas other IP addresses are considered private, and are intended for use within an organization. Because these private IP addresses might need to communicate outside of their local networks, Network Address Translation (NAT) allows private IP addresses (as defined in RFC 1918) to be translated into Internet-routable IP addresses (that is, public IP addresses).

Effectively troubleshooting a NAT issue requires knowledge of various NAT terminology and configuration syntax, in addition to a collection of **show** and **debug** commands. This *Network Troubleshooting Video Mentor* lab reviews basic NAT operation, in addition to configuration and verification syntax. Also, you are presented with a listing of common NAT troubleshooting targets and a specific trouble ticket to resolve.

Scenario

This lab includes the following steps:

- Step 1** Review NAT theory.
- Step 2** Examine lab topology.
- Step 3** Identify NAT verification and troubleshooting commands.
- Step 4** Discuss possible causes for a NAT translation to fail.
- Step 5** Interpret a trouble ticket.
- Step 6** Troubleshoot and resolve the identified NAT issue.
- Step 7** Summarize key elements of the NAT troubleshooting process.

Review NAT Theory

Consider Figure 11-1, which shows a basic NAT topology. Note that even though the IP addresses of 172.16.1.1 and 192.168.1.1 are actually private IP addresses, for the purpose of this discussion, assume they are publicly routable IP addresses. The reason for the use of these private IP addresses to represent public IP addresses is to avoid using an entity's registered IP addresses in the example.

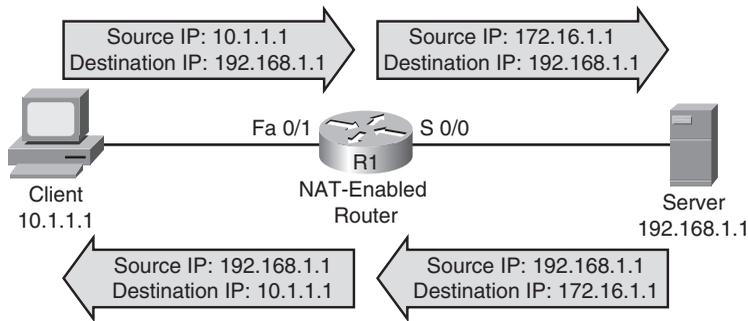


Figure 11-1 Basic NAT Topology

In the topology, a client with a private IP address of 10.1.1.1 wants to communicate with a server on the public Internet. The server's IP address is 192.168.1.1. Router R1 is configured for NAT. Router R1 takes packets coming from 10.1.1.1 destined for 192.168.1.1 and changes the source IP address in the packets' headers to 172.16.1.1 (which we are assuming is a publicly routable IP address for the purposes of this discussion). When the server at IP address 192.168.1.1 receives traffic from the client, the server's return traffic is sent to a destination address of 172.16.1.1. When router R1 receives traffic from the outside network destined for 172.16.1.1, the router translates the destination IP address to 10.1.1.1 and forwards the traffic to the inside network where the client receives the traffic.

To effectively troubleshoot a NAT configuration, you should be familiar with the terminology describing the various IP addresses involved in a NAT translation. Table 11-1 acts as a reference for these NAT IP address classifications.

Table 11-1 Names of NAT IP Addresses

NAT IP Address	Definition
Inside Local	A private IP address referencing an inside device.
Inside Global	A public IP address referencing an inside device.
Outside Local	A private IP address referencing an outside device.
Outside Global	A public IP address referencing an outside device.

As a memory aid, remember that *inside* always refers to an inside device, whereas *outside* always refers to an outside device. Also, think of the word *local* as being similar to the Spanish word *loco*, which means “crazy.” That is what a local address could be thought of—it is a crazy made-up address (that is, a private IP address not routable on the Internet). Finally, let the *g* in *global* remind you of the *g* in *good*, because a global address is a good (that is, routable on the Internet) IP address.

Based on these definitions, Table 11-2 categorizes the IP addresses previously shown in Figure 11-1.

Table 11-2 Classifying the NAT IP Addresses in Figure 11-1

NAT IP Address	NAT IP Address Type
Inside Local	10.1.1.1
Inside Global	172.16.1.1
Outside Local	None
Outside Global	192.168.1.1

Again, refer back to Figure 11-1. Example 11-1 shows how you could configure router R1 in that figure for dynamic NAT to support the translation shown.

Example 11-1 Dynamic NAT Sample Configuration

```
R1#show run
...OUTPUT OMITTED...
interface FastEthernet1/0
 ip address 10.1.1.100 255.255.255.0
 ip nat inside
!
interface Serial 0/0
 ip address 172.16.1.100 255.255.255.0
 ip nat outside
!
ip nat pool OUTSIDE_POOL 172.16.1.1 172.16.1.10 netmask 255.255.255.0
ip nat inside source list 1 pool OUTSIDE_POOL
!
access-list 1 permit 10.0.0.0 0.0.0.255
...OUTPUT OMITTED...
```

In the example, ACL 1 is used to identify the inside addresses (the 10.1.1.0/24 network in this example) to be translated. A pool of addresses named OUTSIDE_POOL is defined as IP addresses in the range 172.16.1.1–172.16.1.10. The **ip nat inside source list 1 pool OUTSIDE_POOL** command associates the internal range of addresses defined by ACL 1 with the range of outside addresses defined by the OUTSIDE_POOL pool. Finally, you need to indicate what router interface is acting as the inside interface and what interface is acting as the outside interface. Note that you could have multiple interfaces acting as inside or outside interfaces. The **ip nat inside** command is issued for interface Fast Ethernet 1/0, and the **ip nat outside** command is issued for interface Serial 0/0.

Examine Lab Topology

In the lab topology, which is shown in Figure 11-2, all routers have been configured to have full reachability to one another, using OSPF as the routing protocol. Router R2 has been configured to run NAT, where IP addresses belonging to the 192.168.0.0/24 network are translated into an IP

address of 172.16.1.2 (that is, the IP address of subinterface Serial 1/0.1 on router R2), and IP addresses belonging to the 192.168.1.0/24 network are translated into an IP address of 172.16.2.1 (that is, the IP address of subinterface Serial 1/0.2 on router R2).

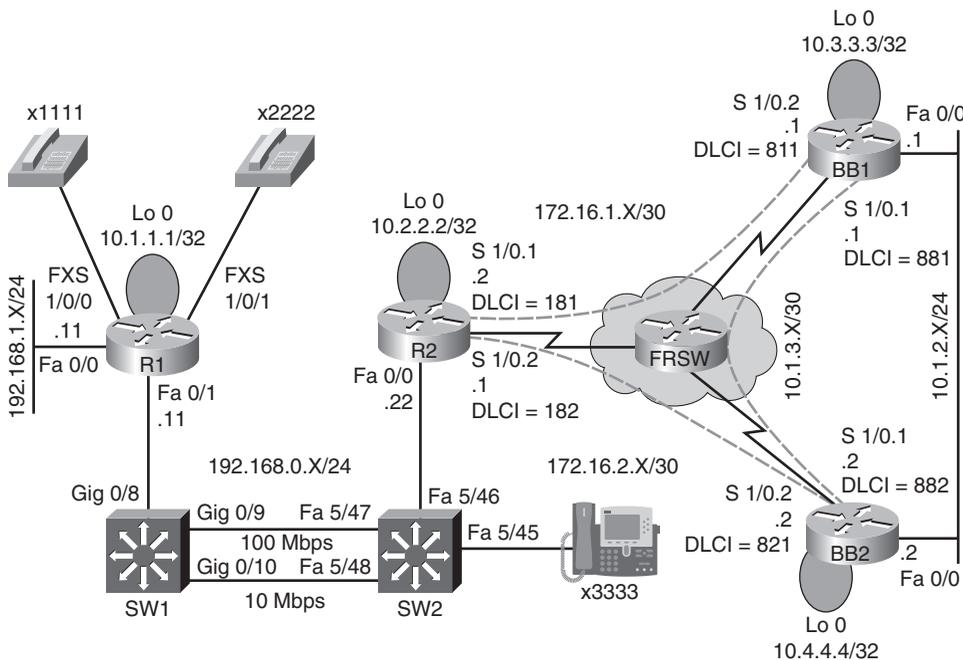


Figure 11-2 Lab 11 Topology

Identify NAT Verification and Troubleshooting Commands

Troubleshooting a NAT issue might require knowledge of Cisco IOS NAT configuration commands, in addition to **show** and **debug** commands. Therefore, this lab presents a collection of commands that you can use to configure, verify, and troubleshoot potential NAT issues. The syntax for these commands is presented in Table 11-3.

Table 11-3 NAT Configuration, Verification, and Troubleshooting Syntax

Command	Description
Router# show ip nat statistics	Displays statistics about NAT operation, including such information as the number of translations performed and identification of the inside and outside interfaces.
Router# show ip nat translations	Shows a table containing the addresses used in active NAT translations.
Router# clear ip nat translations *	Removes all dynamic entries from a router's NAT translation table.
Router(config-if)# ip nat {inside outside}	Designates an interface as an inside or outside NAT interface.

Table 11-3 NAT Configuration, Verification, and Troubleshooting Syntax

Command	Description
Router(config)# ip nat pool <i>pool-name start-ip end-ip</i> { netmask subnet-mask prefix-length prefix-length }	Defines a pool of inside global addresses into which inside local addresses can be translated.
Router(config)# ip nat inside source list <i>access-list pool pool-name</i> [overload]	Creates a pool of inside local addresses to be translated into one or more inside global addresses. (NOTE: The overload option allows multiple inside local addresses to be translated into a single inside global address.)
Router(config)# ip nat translation max-entries <i>number</i>	Specifies the maximum number of entries permitted in a router's NAT table.

Discuss Possible Causes for a NAT Translation to Fail

This lab discusses troubleshooting the following NAT issues:

- Inside and outside NAT interfaces specified incorrectly.
- Overlap of IP addresses assigned by dynamic pools of IP addresses or a static IP address.
- Incorrect configuration of access list used to specify a pool of IP addresses.
- Translated IP address remains cached in a system after timing out of a router's NAT translation table.

Interpret a Trouble Ticket

The following trouble ticket is presented in this lab:

Company A is dual-homed out to the Internet (that is, routers BB1 and BB2, where each router represents a different ISP). Inside IP addresses in the 192.168.0.0/24 subnet should be translated into the IP address of sub-interface Serial 1/0.1 on router R2, whereas inside IP addresses in the 192.168.1.0/24 subnet should be translated into the IP address of sub-interface Serial 1/0.2 on router R2. Router R2's NAT translation table shows two active translations. The configuration, therefore, seems to be partially working. However, no additional NAT translations can be set up.

Troubleshoot and Resolve the Identified NAT Issue

Because router R2 is the router configured for NAT, this lab begins by viewing active NAT translations on router R2, as shown in Example 11-2.

Example 11-2 Viewing Router R2's NAT Translations

R2#show ip nat translations				
Pro	Inside global	Inside local	Outside local	Outside global
icmp	172.16.1.2:12	192.168.0.11:12	10.4.4.4:12	10.4.4.4:12
icmp	172.16.2.1:512	192.168.1.50:512	10.1.3.2:512	10.1.3.2:512

Next, the **debug ip nat** command is issued to view NAT translations as they occurred. Example 11-3 shows sample output from this command.

Example 11-3 Viewing Real-Time NAT Translations on Router R2

```
R2#debug ip nat
IP NAT debugging is on
*Mar 1 01:18:22.515: NAT*: s=10.4.4.4, d=172.16.1.2->192.168.0.11 [2257]
*Mar 1 01:18:22.539: NAT*: s=192.168.0.11->172.16.1.2, d=10.4.4.4 [2258]
*Mar 1 01:18:22.739: NAT*: s=10.4.4.4, d=172.16.1.2->192.168.0.11 [2258]
*Mar 1 01:18:22.859: NAT*: s=192.168.0.11->172.16.1.2, d=10.4.4.4 [2259]
*Mar 1 01:18:23.115: NAT*: s=10.4.4.4, d=172.16.1.2->192.168.0.11 [2259]
*Mar 1 01:18:23.155: NAT*: s=192.168.0.11->172.16.1.2, d=10.4.4.4 [2260]
*Mar 1 01:18:23.171: NAT*: s=192.168.1.50->172.16.2.1, d=10.1.3.2 [34875]
*Mar 1 01:18:23.291: NAT*: s=10.4.4.4, d=172.16.1.2->192.168.0.11 [2260]
*Mar 1 01:18:23.311: NAT*: s=10.1.3.2, d=172.16.2.1->192.168.1.50 [34875]
*Mar 1 01:18:23.355: NAT*: s=192.168.0.11->172.16.1.2, d=10.4.4.4 [2261]
*Mar 1 01:18:23.483: NAT*: s=10.4.4.4, d=172.16.1.2->192.168.0.11 [2261]
*Mar 1 01:18:23.563: NAT*: s=192.168.0.11->172.16.1.2, d=10.4.4.4 [2262]
*Mar 1 01:18:23.731: NAT*: s=10.4.4.4, d=172.16.1.2->192.168.0.11 [2262]
*Mar 1 01:18:23.827: NAT*: s=192.168.0.11->172.16.1.2, d=10.4.4.4 [2263]
*Mar 1 01:18:23.999: NAT*: s=10.4.4.4, d=172.16.1.2->192.168.0.11 [2263]
*Mar 1 01:18:24.075: NAT*: s=192.168.0.11->172.16.1.2, d=10.4.4.4 [2264]
*Mar 1 01:18:24.099: NAT*: s=192.168.1.50->172.16.2.1, d=10.1.3.2 [35083]
*Mar 1 01:18:24.299: NAT*: s=10.4.4.4, d=172.16.1.2->192.168.0.11 [2264]
R2#u all
```

To confirm the reported issue that router R2 only supports two simultaneous NAT translations, an attempt is made to establish a third NAT translation. This failed attempt is illustrated in Example 11-4.

Example 11-4 Failed Ping Attempt from Router R1 to a Router BB2's Loopback Interface

```
R1#ping 10.4.4.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.4.4.4, timeout is 2 seconds:
U.U.U
Success rate is 0 percent (0/5)
```

To determine if the failed ping is a result of NAT or a different issue, router R2's NAT translation table is cleared with the **clear ip nat translation *** command, after which router R1 could successfully ping 10.4.4.4. These results lead to the conclusion that the issue reported in the trouble ticket is accurate. Specifically, router R2 is only capable of setting up a maximum of two simultaneous NAT translations.

A common reason why a router configured for NAT cannot establish multiple NAT translations is that the NAT configuration has more inside local addresses than available inside global addresses. Port Address Translation (PAT) uses the **overload** keyword as part of the **ip nat inside source** command to allow a single inside global address to be used for multiple inside local addresses.

To determine if the reported issue is resulting from the omission of the **overload** keyword in the **ip nat inside source** command, this lab views the running configurations of router R2, a portion of which is shown in Example 11-5.

Example 11-5 Router R2's Running Configuration

```
R2#show run
...OUTPUT OMITTED...
interface Loopback0
 ip address 10.2.2.2 255.255.255.255
!
interface FastEthernet0/0
 ip address 192.168.0.22 255.255.255.0
 ip nat inside
!
interface Serial1/0
 no ip address
 encapsulation frame-relay
!
interface Serial1/0.1 point-to-point
 ip address 172.16.1.2 255.255.255.252
 ip nat outside
 frame-relay interface-dlci 181
!
interface Serial1/0.2 point-to-point
 ip address 172.16.2.1 255.255.255.252
 ip nat outside
 frame-relay interface-dlci 182
!
router ospf 1
 network 0.0.0.0 255.255.255.255 area 0
!
 ip nat translation max-entries 2
 ip nat inside source list 1 interface Serial1/0.2 overload
 ip nat inside source list 2 interface Serial1/0.1 overload
!
 access-list 1 permit 192.168.1.0 0.0.0.255
 access-list 2 permit 192.168.0.0 0.0.0.255
!
...OUTPUT OMITTED...
```

The **overload** option is specified as part of the **ip nat inside source** command, thus eliminating that as the cause of the reported issue. However, the **ip nat translation max-entries 2** command is observed in router R2's running configuration. This command limits the maximum number of simultaneous NAT translations on router R2 to only two. Therefore, this command is removed, as shown in Example 11-6.

Example 11-6 Removing the ip nat translation max-entries 2 Command from Router R2

```
R2#conf term
R2(config)#no ip nat translation max-entries 2
R2(config)#end
```

After removing the **ip nat translation max-entries 2** command, router R2 is able to simultaneously support multiple NAT translations, as evidenced in Example 11-7.

Example 11-7 Router R2 Supporting Multiple Simultaneous NAT Translations

```
R2#show ip nat translations
Pro Inside global      Inside local      Outside local      Outside global
icmp 172.16.1.2:14     192.168.0.11:14  10.4.4.4:14        10.4.4.4:14
icmp 172.16.1.2:15     192.168.0.11:15  10.3.3.3:15        10.3.3.3:15
icmp 172.16.1.2:16     192.168.0.11:16  10.4.4.4:16        10.4.4.4:16
icmp 172.16.2.1:512    192.168.1.50:512 10.1.3.2:512       10.1.3.2:512
```

Summarize Key Elements of the NAT Troubleshooting Process

This lab concludes by discussing common NAT troubleshooting targets and reviewing how this lab resolves the reported issue. Specifically, this lab confirms that router R2 can support a maximum of two simultaneous NAT translations. Then, after interrogating router R2's running configuration, the **ip nat translation max-entries 2** command is observed to be causing the issue. After the removal of this command, router R2 can support multiple simultaneous NAT translations.

VoIP Troubleshooting

Voice over IP (VoIP) allows the spoken voice to be sent over an IP-based network. Because VoIP traffic is latency-sensitive, the underlying data network needs to be able to recognize voice traffic and treat it with high priority. Although a full discussion of quality of service (QoS) technologies is well beyond the scope of the TSHOOT curriculum, this lab introduces fundamental configuration, verification, and troubleshooting commands that can get you started in your configuration or troubleshooting of VoIP quality of service configurations.

This lab begins, however, with an introduction to VoIP technologies. Then, a collection of QoS technologies is reviewed, followed by an examination of this lab's topology. As a reference, you are given multiple verification, troubleshooting, and configuration commands that are useful when addressing a VoIP quality issue. Next, this lab discusses possible causes for a VoIP quality issue, and challenges you with a trouble ticket.

Scenario

This lab includes the following steps:

- Step 1** Review VoIP and QoS theory.
- Step 2** Examine lab topology.
- Step 3** Identify verification and troubleshooting commands for VoIP quality issues.
- Step 4** Discuss possible causes and resolutions of a VoIP quality issue.
- Step 5** Interpret a trouble ticket.
- Step 6** Troubleshoot and resolve the identified VoIP quality issue.
- Step 7** Summarize key elements of the VoIP quality troubleshooting process.

Review VoIP and QoS Theory

Whereas *VoIP* is a technology that sends the spoken voice (for example, voice originating from analog phones connected to a voice gateway) across an IP network, *IP telephony* is an extension of VoIP, which uses voice endpoints (for example, IP phones) that natively speak IP.

An IP telephony network not only duplicates the features offered in traditional corporate Private Branch Exchange (PBX) telephony systems, but IP telephony expands on those features. You often find the following sampling of components in an IP telephony network, as shown in Figure 12-1:

- **IP phone:** An IP phone provides IP voice to the desktop.
- **Gatekeeper:** A gatekeeper provides call admission control (CAC), bandwidth control and management, and address translation features.

- **Gateway:** A gateway provides translation between VoIP and non-VoIP networks, such as the Public Switched Telephone Network (PSTN). A gateway also provides physical access for local analog and digital voice devices, such as telephones, fax machines, key sets, and PBXs.
- **Multipoint Control Unit (MCU):** An MCU mixes audio or video streams, thus allowing participants in multiple locations to attend the same conference.
- **Call agent:** A call agent provides call control for IP phones, CAC, bandwidth control and management, and address translation. Although a call agent could be server-based, Cisco also supports a router-based call agent, known as Cisco Unified Communications Manager Express (UCME).

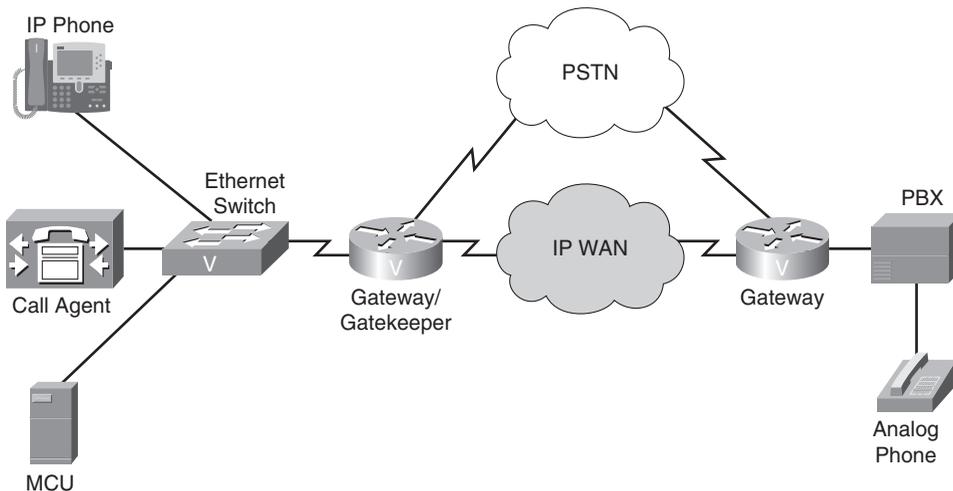


Figure 12-1 IP Telephony Topology

VoIP and IP telephony technologies rely on multiple protocols to set up, maintain, and tear down a call. Protocols discussed in this lab include the following:

- **Real-time Transport Protocol (RTP):** The protocol used to carry voice media.
- **RTP Control Protocol (RTCP):** A supervisory protocol for RTP, which can report quality information for an RTP stream.
- **Skinnny Client Control Protocol (SCCP):** A client-server protocol often used to send signaling messages between a Cisco IP Phone and a Unified Communications Manager (UCM) server.
- **Media Gateway Control Protocol (MGCP):** A client-server protocol that allows an endpoint (for example, a voice port) in a gateway (for example, a voice-enabled router) to register with a call agent (for example, a Cisco UCM server).
- **H.323:** A mature peer-to-peer protocol that allows call routing logic to be configured on H.323 terminals and gateways, rather than relying on an external server.
- **Session Initiation Protocol (SIP):** A newer peer-to-peer protocol that tends to be very vendor interoperable and can optionally rely on external servers for its call-routing logic.

Although troubleshooting voice networks encompasses many potential troubleshooting targets, this lab focuses on troubleshooting a voice quality issue. When voice traffic is added to an existing data network, voice packets start to contend with data (and perhaps voice) for bandwidth. As a result, the perceived quality of a voice call might be poor. Fortunately, Cisco offers several QoS mechanisms that can recognize voice packets as high-priority traffic and then treat that special traffic in a special way. As a few examples, this lab discussed the following QoS features:

- Classification and marking
 - Network-Based Application Recognition (NBAR)
 - IP Precedence
 - Differentiated Services Code Point (DSCP)
- Congestion management
 - Class-Based Weighted Fair Queuing (CB-WFQ)
 - Low Latency Queuing (LLQ)
- Congestion avoidance
 - Weighted Random Early Detection (WRED)
 - Explicit Congestion Notification (ECN)
- Traffic conditioning
 - Policing
 - Shaping
- Link efficiency
 - Compression
 - Link Fragmentation and Interleaving (LFI)

Examine Lab Topology

The topology used in this lab is presented in Figure 12-2. Router R1 is acting as a call agent for the Cisco IP Phone at directory number 3333. Specifically, router R1 is configured as a Cisco Unified Communications Manager Express (UCME) router. Two analog phones (with directory numbers 1111 and 2222) are also attached to router R1. Also, router R2 is configured to provide priority treatment to traffic marked with a Differentiated Services Code Point (DSCP) value of Expedited Forwarding (EF) as that traffic is sent out over the Frame Relay IP WAN.

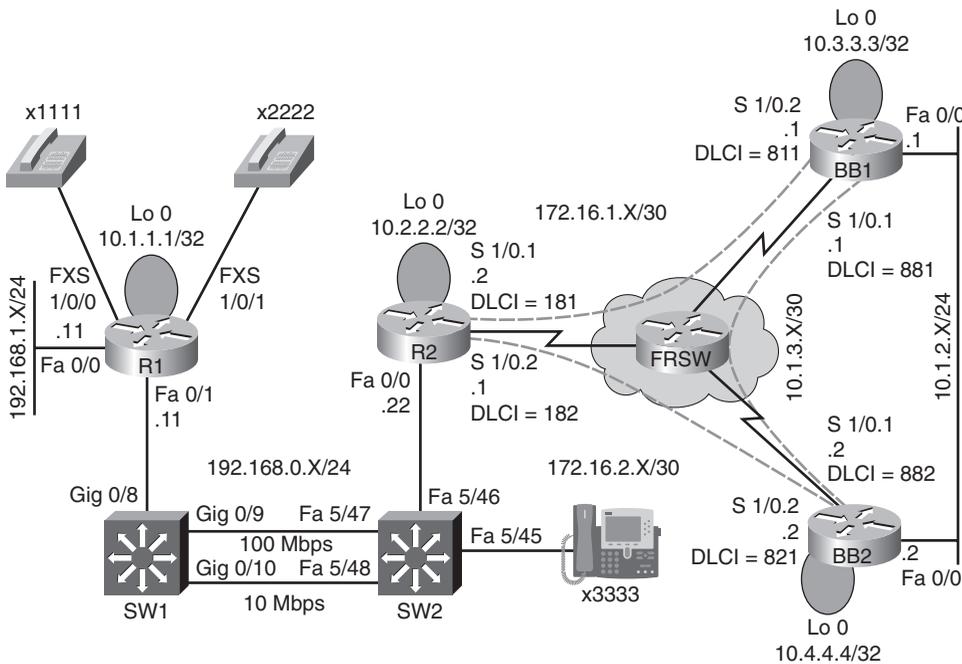


Figure 12-2 Lab 12 Topology

Identify Verification and Troubleshooting Commands for VoIP Quality Issues

In order to effectively troubleshoot a VoIP quality issue, beyond knowing how to verify existing VoIP QoS configurations, you might need to know the basics of configuring QoS mechanisms to help optimize voice quality. Therefore, this lab presents a collection of commands, which can be used to configure, verify, and troubleshoot potential VoIP quality issues. The syntax for these commands is presented in Table 12-1.

Table 12-1 Voice Quality Troubleshooting Syntax

Command	Description
Router# show class-map	Displays a router's class maps, including a description of how traffic was classified.
Router# show policy-map	Shows the policy applied to specific classes of traffic.
Router# show policy-map interface <i>interface-id</i>	Presents class map and policy map configuration information for a policy map applied to an interface, and traffic counts for packets matching the listed class maps.
Router(config)# class-map [match-any match-all] <i>class-map-name</i>	Creates a class map using either match any or match all logic (with a default of match all).
Router(config-cmap)# match <i>match-criterion</i>	Classifies traffic into a class map, using a variety of match criteria.

Table 12-1 Voice Quality Troubleshooting Syntax

Command	Description
Router(config)# policy-map <i>policy-map-name</i>	Creates a policy map and enters policy map configuration mode.
Router(config-pmap)# class <i>class-map-name</i>	Enters policy map class configuration mode.
Router(config-pmap-c)# bandwidth <i>bandwidth</i>	Specifies a minimum bandwidth guarantee (in kbps) for a class of traffic (CB-WFQ).
Router(config-pmap-c)# priority <i>bandwidth</i>	Specifies the maximum bandwidth guarantee (in kbps) for a class of latency-sensitive traffic (LLQ).
Switch(config-if)# auto qos voip { trust cisco-phone }	Optimizes a switch port's QoS configuration for VoIP traffic.
Router(config-if)# auto qos voip	Optimizes a router interface's QoS configuration for VoIP traffic.

Discuss Possible Causes and Resolutions of a VoIP Quality Issue

This lab discussed the following causes and symptoms of a VoIP quality issue:

- Lack of bandwidth
- Excessive end-to-end delay
- Jitter
- Dropped packets

Interpret a Trouble Ticket

The following trouble ticket is presented in this lab:

When placing calls across the Frame Relay WAN, users are complaining that the voice quality is poor for calls originating on an analog phone, whereas the voice quality is fine for calls originating on an IP phone.

Troubleshoot and Resolve the Identified VoIP Quality Issue

The voice quality issue being reported is focused on calls going over the Frame Relay WAN, out of router R2's Serial 1/0 subinterfaces. Therefore, this lab first examines the QoS configuration on router R2, as seen in Example 12-1.

Example 12-1 Viewing Router R2's QoS Configuration

```
R2#show run
Building configuration...
...OUTPUT OMITTED...
hostname R2
!
class-map match-all VOICE
  match dscp ef
!
policy-map TSHOOT
  class VOICE
    priority 64
!
interface Loopback0
  ip address 10.2.2.2 255.255.255.255
!
interface FastEthernet0/0
  ip address 192.168.0.22 255.255.255.0
  duplex auto
  speed auto
!
i
interface Serial1/0
  no ip address
  encapsulation frame-relay
  service-policy output TSHOOT
!
interface Serial1/0.1 point-to-point
  ip address 172.16.1.2 255.255.255.252
  frame-relay interface-dlci 181
!
interface Serial1/0.2 point-to-point
  ip address 172.16.2.1 255.255.255.252
  frame-relay interface-dlci 182
!
router ospf 1
  network 0.0.0.0 255.255.255.255 area 0
!
...OUTPUT OMITTED...
```

As seen in the example, traffic marked with a DSCP value of Expedited Forwarding, or EF, is being given priority treatment. Specifically, if voice traffic is marked with a DSCP value of EF, it will be sent ahead of other traffic as it exits router R2's Serial 1/0 subinterfaces.

Cisco IP Phones, by default, mark voice packets with a DSCP marking of EF, which is the recommendation for voice. Therefore, voice packets exiting router R2's Serial 1/0 subinterfaces will be given priority treatment.

However, two analog phones are attached to router R1. This lab viewed the running configuration of router R1, as illustrated in Example 12-2, to determine if router R1 had any QoS configuration that would give priority treatment to voice traffic originating from these analog phones.

Example 12-2 Viewing Router R1's Running Configuration

```
R1#show run
Building configuration...
...OUTPUT OMITTED...
hostname R1
!
ip dhcp pool TSHOOT
    network 192.168.0.0 255.255.255.0
    option 150 ip 192.168.0.11
    default-router 192.168.0.11
!
interface Loopback0
    ip address 10.1.1.1 255.255.255.255
!
interface FastEthernet0/0
    ip address 192.168.1.11 255.255.255.0
!
interface FastEthernet0/1
    ip address 192.168.0.11 255.255.255.0
!
router ospf 1
    network 0.0.0.0 255.255.255.255 area 0
!
voice-port 1/0/0
!
voice-port 1/0/1
!
dial-peer voice 1 pots
    destination-pattern 1111
    port 1/0/0
!
dial-peer voice 2 pots
    destination-pattern 2222
    port 1/0/1
!
telephony-service
```

```
max-ephones 5
max-dn 10
ip source-address 192.168.0.11 port 2000
create cnf-files version-stamp Jan 01 2002 00:00:00
max-conferences 4 gain -6
!
ephone-dn 1
  number 3333
!
ephone 2
  mac-address 0008.A3B8.95C4
  button 1:1
...OUTPUT OMITTED...
```

Although the output from the previous example does provide a glimpse into the configuration of Cisco Unified Communications Manager Express, notice that router R1 does not seem to have any QoS configuration that would give priority treatment originating from one of its attached analog phones. This lab then challenged you to determine an approach to marking router R1's voice traffic as high priority.

One solution of many to this issue is to leverage Cisco's AutoQoS VoIP feature, which with one command (that is, the **auto qos voip** command issued in interface configuration mode) can prioritize a router's voice traffic exiting a specific interface. Example 12-3 demonstrates how this lab enabled the AutoQoS VoIP feature on router R1's Fast Ethernet 0/1 interface.

Example 12-3 Configuring AutoQoS VoIP on Router R1's Fast Ethernet 0/1 Interface

```
R1#config term
R1(config)#int fa 0/1
R1(config-if)#auto qos voip
R1(config-if)#end
```

Example 12-4 highlights the QoS commands automatically added by the AutoQoS VoIP feature.

Example 12-4 Viewing QoS Commands Added by AutoQoS VoIP

```
R1#show run
Building configuration...
...OUTPUT OMITTED...
hostname R1
!
ip dhcp excluded-address 192.168.0.1 192.168.0.100
!
ip dhcp pool TSHOOT
    network 192.168.0.0 255.255.255.0
    option 150 ip 192.168.0.11
    default-router 192.168.0.11
!
class-map match-any AutoQoS-VoIP-Remark
    match ip dscp ef
    match ip dscp cs3
    match ip dscp af31
class-map match-any AutoQoS-VoIP-Control-UnTrust
    match access-group name AutoQoS-VoIP-Control
class-map match-any AutoQoS-VoIP-RTP-UnTrust
    match protocol rtp audio
    match access-group name AutoQoS-VoIP-RTCP
!
policy-map AutoQoS-Policy-UnTrust
    class AutoQoS-VoIP-RTP-UnTrust
        priority percent 70
        set dscp ef
    class AutoQoS-VoIP-Control-UnTrust
        bandwidth percent 5
        set dscp af31
    class AutoQoS-VoIP-Remark
        set dscp default
    class class-default
        fair-queue
!
interface Loopback0
    ip address 10.1.1.1 255.255.255.255
!
interface FastEthernet0/0
    ip address 192.168.1.11 255.255.255.0
!
interface FastEthernet0/1
    ip address 192.168.0.11 255.255.255.0
    auto qos voip
    service-policy output AutoQoS-Policy-UnTrust
```

```
!  
router ospf 1  
  network 0.0.0.0 255.255.255.255 area 0  
!  
ip access-list extended AutoQoS-VoIP-Control  
  permit tcp any any eq 1720  
  permit tcp any any range 11000 11999  
  permit udp any any eq 2427  
  permit tcp any any eq 2428  
  permit tcp any any range 2000 2002  
  permit udp any any eq 1719  
  permit udp any any eq 5060  
ip access-list extended AutoQoS-VoIP-RTCP  
  permit udp any any range 16384 32767  
!  
rmon event 33333 log trap AutoQoS description "AutoQoS SNMP traps for Voice  
Drops" owner AutoQoS  
rmon alarm 33333 cbQosCMDropBitRate.1081.1083 30 absolute rising-threshold 1  
33333 falling-threshold 0 owner AutoQoS  
!  
voice-port 1/0/0  
!  
voice-port 1/0/1  
!  
dial-peer voice 1 pots  
  destination-pattern 1111  
  port 1/0/0  
!  
dial-peer voice 2 pots  
  destination-pattern 2222  
  port 1/0/1  
!  
telephony-service  
  max-ephones 5  
  max-dn 10  
  ip source-address 192.168.0.11 port 2000  
  create cnf-files version-stamp Jan 01 2002 00:00:00  
  max-conferences 4 gain -6  
!  
ephone-dn 1  
  number 3333  
!  
ephone 2  
  mac-address 0008.A3B8.95C4  
  button 1:1  
!  
...OUTPUT OMITTED...
```

By enabling the AutoQoS VoIP feature, router R1 can now mark traffic exiting its Fa 0/1 interface as having a DSCP marking of EF. As a result, router R2 can now prioritize voice traffic, originating from router R1's analog phones, as it sends that traffic out over the Frame Relay WAN.

Summarize Key Elements of the VoIP Quality Troubleshooting Process

This lab concludes by discussing common VoIP technologies, QoS technologies, and how this lab resolved the reported issue. Specifically, this lab confirmed that router R2 was giving priority treatment to traffic marked with a DSCP value of EF, a value that Cisco IP Phones automatically use to mark voice traffic. However, even though voice traffic originating from a Cisco IP Phone was appropriately marked with a DSCP value of EF, an examination of router R1's configuration revealed that voice traffic originating on analog phones was not receiving priority markings. Although multiple QoS solutions exist for marking those voice packets appropriately, this lab leveraged Cisco's AutoQoS VoIP feature to cause router R1 to mark its voice traffic appropriately.

InformIT is a brand of Pearson and the online presence for the world's leading technology publishers. It's your source for reliable and qualified content and knowledge, providing access to the top brands, authors, and contributors from the tech community.

LearnIT at InformIT

Looking for a book, eBook, or training video on a new technology? Seeking timely and relevant information and tutorials? Looking for expert opinions, advice, and tips? **InformIT has the solution.**

- Learn about new releases and special promotions by subscribing to a wide variety of newsletters. Visit **informit.com/newsletters**.
- Access FREE podcasts from experts at **informit.com/podcasts**.
- Read the latest author articles and sample chapters at **informit.com/articles**.
- Access thousands of books and videos in the Safari Books Online digital library at **safari.informit.com**.
- Get tips from expert blogs at **informit.com/blogs**.

Visit **informit.com/learn** to discover all the ways you can access the hottest technology content.

Are You Part of the IT Crowd?

Connect with Pearson authors and editors via RSS feeds, Facebook, Twitter, YouTube, and more! Visit **informit.com/socialconnect**.



Try Safari Books Online FREE

Get online access to 5,000+ Books and Videos



Safari[®]
Books Online

FREE TRIAL—GET STARTED TODAY!
www.informit.com/safaritrial



Find trusted answers, fast

Only Safari lets you search across thousands of best-selling books from the top technology publishers, including Addison-Wesley Professional, Cisco Press, O'Reilly, Prentice Hall, Que, and Sams.



Master the latest tools and techniques

In addition to gaining access to an incredible inventory of technical books, Safari's extensive collection of video tutorials lets you learn from the leading video training experts.

WAIT, THERE'S MORE!



Keep your competitive edge

With Rough Cuts, get access to the developing manuscript and be among the first to learn the newest technologies.



Stay current with emerging technologies

Short Cuts and Quick Reference Sheets are short, concise, focused content created to get you up-to-speed quickly on new and cutting-edge technologies.



Adobe Press



Cisco Press



Microsoft Press



O'REILLY



que



SAMS

