

TSHOOT

Troubleshooting and Maintaining Cisco IP Networks

Version 1.0

Lab Guide

Text Part Number: 97-2820-01



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

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

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

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

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

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

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

Table of Contents

Lab Guide	1
Overview	1
Outline	1
Lab 1-1: Lab Access	2
Activity Objective	2
Information Packet	2
Required Resources	3
Job Aids	3
Task 1: Verify Console Connections	3
Task 2: Verify Remote Desktop Connections	3
Lab 2-1: Introduction to Troubleshooting	5
Activity Objective	5
Information Packet	5
Required Resources	6
Job Aids	6
Trouble Ticket: No Connectivity to the Server	6
Instructions	6
Lab Setup	6
Troubleshooting Log	7
Lab Debrief Notes	9
Lab 2-1: Alternate Solutions	9
Lab 2-1: Alternate Methods and Processes	10
Lab 2-1: Procedure and Communication Improvements	11
Lab 2-1: Important Commands and Tools	12
Lab 3-1: Maintenance and Troubleshooting Tools	13
Activity Objective	13
Information Packet	13
Required Resources	14
Lab Setup	14
Scenario	14
Task 1: Assign Responsibilities	14
Task 2: Review the Physical Lab Topology	16
Task 3: Review the Logical Lab Topology	23
Task 4: Review Troubleshooting and Maintenance Tools	31
Reference Configurations	33
Lab Debrief Notes	87
Lab 3-1: Alternate Solutions	87
Lab 3-1: Alternate Methods and Processes	88
Lab 3-1: Procedure and Communication Improvements	89
Lab 3-1: Important Commands and Tools	90
Lab 4-1: Layer 2 Connectivity and Spanning Tree	91
Activity Objective	91
Information Packet	91
Required Resources	91
Job Aids	92
Trouble Ticket A: Switch Replacement Gone Bad	92
Trouble Ticket B: Guest Access Problem in Branch	92
Trouble Ticket C: Internet Service Provider 1 Seems to Be Down	92
Instructions	93
Lab Setup	93
Troubleshooting Log	94
Troubleshooting Log	96
Troubleshooting Log	98
Lab 4-1: Sample Troubleshooting Flows	100

Lab Debrief Notes	114
Lab 4-1: Alternate Solutions	114
Lab 4-1: Alternate Methods and Processes	115
Lab 4-1: Procedure and Communication Improvements	116
Lab 4-1: Important Commands and Tools	117
Lab 4-1: References	118
Lab 4-2: Layer 3 Switching and First-Hop Redundancy	119
Activity Objective	119
Information Packet	119
Required Resources	120
Job Aids	120
Trouble Ticket D: Switch ASW1 Cannot Be Managed from Server SRV1	120
Trouble Ticket E: Failover not Functioning as Expected	120
Trouble Ticket F: Verify HSRP Authentication	121
Trouble Ticket G: HSRP and GLBP Comparison	121
Instructions	122
Lab Setup	122
Troubleshooting Log	123
Troubleshooting Log	125
Troubleshooting Log	127
Troubleshooting Log	129
Lab 4-2: Sample Troubleshooting Flows	131
Lab Debrief Notes	149
Lab 4-2: Alternate Solutions	149
Lab 4-2: Alternate Methods and Processes	150
Lab 4-2: Procedure and Communication Improvements	151
Lab 4-2: Important Commands and Tools	152
Lab 4-2: References	153
Lab 5-1: Layer 3 Connectivity and EIGRP	154
Activity Objective	154
Information Packet	154
Required Resources	154
Job Aids	155
Trouble Ticket H: Preparation for CCTV Pilot	155
Trouble Ticket I: Fire in the Server Room	155
Trouble Ticket J: User in Branch Cannot Access the Internet	156
Instructions	156
Lab Setup	156
Troubleshooting Log	157
Troubleshooting Log	159
Troubleshooting Log	161
Lab 5-1: Sample Troubleshooting Flows	163
Lab Debrief Notes	181
Lab 5-1: Alternate Solutions	181
Lab 5-1: Alternate Methods and Processes	182
Lab 5-1: Procedure and Communication Improvements	183
Lab 5-1: Important Commands and Tools	184
Lab 5-1: References	185
Lab 5-2: OSPF and Route Redistribution	186
Activity Objective	186
Information Packet	186
Required Resources	186
Job Aids	187
Introduction: Migration to OSPF	187
Trouble Ticket K: No Connectivity from Client PC CLT2	188
Trouble Ticket L: No Connectivity from Client PC CLT3	189
Trouble Ticket M: Internet not Reachable from Client PC CLT1	189
Trouble Ticket N: OSPF Authentication Not Working	189
Instructions	189
Lab Setup	190

Troubleshooting Log	191
Troubleshooting Log	193
Troubleshooting Log	195
Troubleshooting Log	197
Lab 5-2: Sample Troubleshooting Flows	199
Lab Debrief Notes	221
Lab 5-2: Alternate Solutions	221
Lab 5-2: Alternate Methods and Processes	222
Lab 5-2: Procedure and Communication Improvements	223
Lab 5-2: Important Commands and Tools	224
Lab 5-2: References	225
Lab 5-3: Border Gateway Protocol	226
Activity Objective	226
Information Packet	226
Required Resources	226
Job Aids	227
Introduction: Implementation of BGP	227
Trouble Ticket O: BGP Peering to Router ISP1 Not Established	228
Trouble Ticket P: Client CLT1 Cannot Reach the Internet	228
Instructions	228
Lab Setup	229
Troubleshooting Log	230
Troubleshooting Log	232
Lab 5-3: Sample Troubleshooting Flows	234
Lab Debrief Notes	245
Lab 5-3: Alternate Solutions	245
Lab 5-3: Alternate Methods and Processes	246
Lab 5-3: Procedure and Communication Improvements	247
Lab 5-3: Important Commands and Tools	248
Lab 5-3: References	249
Lab 5-4: Router Performance	250
Activity Objective	250
Information Packet	250
Required Resources	251
Job Aids	251
Trouble Ticket Q: Problems with Connectivity	251
Instructions	252
Lab Setup	254
Troubleshooting Log	255
Lab Debrief Notes	257
Lab 5-4: Alternate Solutions	257
Lab 5-4: Alternate Methods and Processes	258
Lab 5-4: Procedure and Communication Improvements	259
Lab 5-4: Important Commands and Tools	260
Lab 6-1: Introduction to Network Security	261
Activity Objective	261
Information Packet	261
Required Resources	261
Job Aids	262
Introduction: Increased Network Security	262
Trouble Ticket R: Internet Not Reachable from Client PC CLT1	263
Trouble Ticket S: Internet Not Reachable from Client PC CLT3	263
Trouble Ticket T: Client PC CLT2 Has No Network Connectivity	263
Instructions	264
Lab Setup	264
Troubleshooting Log	265
Troubleshooting Log	267
Troubleshooting Log	269
Lab 6-1: Sample Troubleshooting Flows	271

Lab Debrief Notes	288
Lab 6-1: Alternate Solutions	288
Lab 6-1: Alternate Methods and Processes	289
Lab 6-1: Procedure and Communication Improvements	290
Lab 6-1: Important Commands and Tools	291
Lab 6-1: References	292
Lab 6-2: Cisco IOS Security Features	293
Activity Objective	293
Information Packet	293
Required Resources	293
Job Aids	294
Introduction: Improving Network Security	294
Trouble Ticket U: No Connectivity from Client PCs CLT2 and CLT3	294
Trouble Ticket V: No Connectivity from Client PC CLT1	295
Trouble Ticket W: No Connectivity to server SRV1	295
Trouble Ticket X: Lost Remote Connectivity to All Routers	295
Trouble Ticket Y: Port Security Problems on Switch BSW1	296
Instructions	296
Lab Setup	296
Troubleshooting Log	297
Troubleshooting Log	299
Troubleshooting Log	301
Troubleshooting Log	303
Troubleshooting Log	305
Lab 6-2: Sample Troubleshooting Flows	307
Lab Debrief Notes	318
Lab 6-2: Alternate Solutions	318
Lab 6-2: Alternate Methods and Processes	319
Lab 6-2: Procedure and Communication Improvements	320
Lab 6-2: Important Commands and Tools	321
Lab 6-2: References	322
Lab 7-1: Troubleshooting Complex Environments	323
Activity Objective	323
Information Packet	323
Required Resources	323
Job Aids	324
Introduction: The Enterprise Network	324
Trouble Ticket A: No Connectivity from CLT1 to SRV1	327
Trouble Ticket B: No Internet Access from CLT1	327
Trouble Ticket C: No Connectivity Between Headquarters and Branch Office	327
Trouble Ticket D: No Internet Access for Guest Users	328
Network Maintenance: Verify Network Operation	328
Instructions	328
Lab Setup	328
Trouble Ticket A Troubleshooting Log	329
Trouble Ticket A Change Log	331
Trouble Ticket B Troubleshooting Log	333
Trouble Ticket B Change Log	335
Trouble Ticket C Troubleshooting Log	337
Trouble Ticket C Change Log	339
Trouble Ticket D Troubleshooting Log	341
Trouble Ticket D Change Log	343
Network Maintenance Process Log	345
Network Maintenance Change Log	347

Lab 7-1: Sample Troubleshooting Flows	350
Lab Debrief Notes	351
Lab 7-1: Alternate Solutions	351
Lab 7-1: Alternate Methods and Processes	352
Lab 7-1: Procedure and Communication Improvements	353
Lab 7-1: Important Commands and Tools	354
Lab 7-1: References	355
Answer Key	356
Lab 2-1 Answer Key: Introduction to Troubleshooting	356
Student Notes	357
Student Notes	358
Lab 3-1 Answer Key: Maintenance and Troubleshooting Tools	359
Student Notes	361
Student Notes	362
Lab 4-1 Answer Key: Layer 2 Connectivity and Spanning Tree	363
Student Notes	365
Student Notes	366
Lab 4-2 Answer Key: Layer 3 Switching and First-Hop Redundancy	367
Student Notes	370
Student Notes	371
Lab 5-1 Answer Key: Layer 3 Connectivity and EIGRP	372
Student Notes	376
Student Notes	377
Lab 5-2 Answer Key: OSPF and Route Redistribution	378
Student Notes	381
Student Notes	382
Lab 5-3 Answer Key: Border Gateway Protocol	383
Student Notes	385
Student Notes	386
Lab 5-4 Answer Key: Router Performance	387
Student Notes	389
Student Notes	390
Lab 6-1 Answer Key: Introduction to Network Security	391
Student Notes	395
Student Notes	396
Lab 6-2 Answer Key: Cisco IOS Security Features	397
Student Notes	401
Student Notes	402
Lab 7-1 Answer Key: Troubleshooting Complex Environments	403
Student Notes	406
Student Notes	407

Lab Guide

Overview

This guide presents the instructions and other information concerning the lab activities for this course. You can find the solutions in the lab activity Answer Key.

Outline

This guide includes these activities:

- Lab 1-1: Lab Access
- Lab 2-1: Introduction to Troubleshooting
- Lab 3-1: Maintenance and Troubleshooting Tools
- Lab 4-1: Layer 2 Connectivity and Spanning Tree
- Lab 4-2: Layer 3 Switching and First-Hop Redundancy
- Lab 5-1: Layer 3 Connectivity and EIGRP
- Lab 5-2: OSPF and Route Redistribution
- Lab 5-3: Border Gateway Protocol
- Lab 5-4: Router Performance
- Lab 6-1: Introduction to Network Security
- Lab 6-2: Cisco IOS Security Features
- Lab 7-1: Troubleshooting Complex Environments
- Answer Key

Lab 1-1: Lab Access

Complete this lab activity to verify connectivity to the lab equipment.

Activity Objective


In this activity, you will learn how to access the equipment that is used during the lab exercises. After completing this activity, you will be able to meet these objectives:

- Access the consoles of the routers and switches used in the lab
- Access the desktop of the server and clients used in the lab

Information Packet

The figure illustrates what you will accomplish in this activity.

Visual Objective



Required Resources

These are the resources and equipment that are required to complete this activity:

- One PC with Internet access per team member
- Four Cisco Catalyst 3560 Series Switches per team
- Six Cisco 1841 Integrated Services Routers per team
- Three client PCs per team
- One server per team

Job Aids

This job aid is available to help you complete the lab activity.

- Lab access instructions obtained from instructor

Task 1: Verify Console Connections

In this task, you will test access to the consoles of the routers and switches in your assigned pod.

Activity Procedure

Complete these steps:

- Step 1** The instructor will assign a pod of lab equipment to your team and provide you with the details that you need to connect to the consoles of the routers and switches in your assigned pod.
- Step 2** Work together with your team members to verify that you can access each of the consoles of the six routers (IRO1, IRO2, CRO1, CRO2, BRO1, and BRO2) and four switches (ASW1, BSW1, CSW1, and CSW2) in your assigned pod.

Activity Verification

You have completed this task when you attain this result:

- You have verified that you can access the consoles of the routers and switches that were assigned to your team.

Task 2: Verify Remote Desktop Connections

In this task, you will test access to the desktop of the clients and the server in your assigned pod.

Activity Procedure

Complete these steps:

- Step 1** The instructor will provide you with the details that you need to connect to the desktop of the clients and server in your assigned pod.
- Step 2** Work together with your team members to verify that you can access each of the desktops of the three clients (CLT1, CLT2, and CLT3) and the server (SRV1) in your assigned pod.

Activity Verification

You have completed this task when you attain this result:

- You have verified that you can access the desktop of the clients and server that were assigned to your team.

Lab 2-1: Introduction to Troubleshooting

Complete this lab activity to practice what you learned in the related module.

Activity Objective


In this activity, you will experience the challenges of troubleshooting in an unknown environment. After completing this activity, you will be able to meet these objectives:

- Identify the minimal documentation that is needed for you to troubleshoot effectively
- Evaluate troubleshooting methods, communication, and planning

Information Packet

The figure illustrates what you will accomplish in this activity.

Visual Objective



Required Resources

These are the resources and equipment that are required to complete this activity:

- One PC with Internet access per team member
- Four Cisco Catalyst 3560 Series Switches per team
- Six Cisco 1841 Integrated Services Routers per team
- Three client PCs per team
- One server per team

Job Aids

These job aids are available to help you complete the lab activity.

- Trouble ticket
- Troubleshooting log

Trouble Ticket: No Connectivity to the Server

You have just started your new job as a network engineer together with a few other engineers who are also newly hired. It is your first day at work, and your new team lead has just shown everybody to their desks and is busy arranging cell phones and all the other things that you need to get started. He takes a quick look at his PC and then tells you that a trouble ticket has just come in and that he would appreciate it if you and your other new teammates could do the initial troubleshooting while he is getting your things together. You are given the passwords to the routers and switches. He tells you to be careful in making changes, but fix the problem if you can. He would at least like you to give him a diagnosis as soon as he returns, which will be in 15 minutes.

The trouble ticket reads:

“A user in Branch1 (PC CLT2) reports problems accessing the shared folder “\\SRV1\Public” on server SRV1. The user had to leave for a meeting that will take all morning, but expects it to work when he returns after lunch.”

Your task is to diagnose the issue, fix it if possible, and report to your team lead in 15 minutes.

Instructions

Together with your team members, diagnose the problem.

No console password has been set for the routers and switches. The enable secret password is “cisco” and the administrator password for the PCs is also “cisco”. To connect to the routers via Telnet or SSH, use the username “admin” and password “cisco”.

Lab Setup

The instructor will provide you with directions to prepare the lab equipment for this lab. After the instructor indicates that the lab is fully prepared, you are ready to start troubleshooting.

Troubleshooting Log

Use this log to document your actions and results during the troubleshooting process.

Task Description

Your task is to diagnose the issue, fix it if possible, and report to your team lead in 15 minutes.

Note Refer to the Activity Verification items at the end of this log to verify that you have successfully completed this task.

[illegible]

Lab Debrief Notes

Use these notes sections to write down the primary learning points that are discussed during the Lab Debrief.

Lab 2-1: Alternate Solutions

[illegible]

[illegible]

[illegible]

Lab 2-1: Important Commands and Tools

[illegible]

Lab 3-1: Maintenance and Troubleshooting Tools

Complete this lab activity to practice what you learned in the related module.

Activity Objective


In this activity, you will survey the network, review and supplement the documentation of the network, and assess and assemble the tools that are available for maintenance and troubleshooting tasks. After completing this activity, you will be able to meet these objectives:

- Distribute troubleshooting tasks among team members based on assigned responsibilities
- Document the physical topology to support future troubleshooting tasks
- Document the logical topology to support future troubleshooting tasks
- Use the available tools to support future troubleshooting tasks

Information Packet

The figure illustrates what you will accomplish in this activity.

Visual Objective



Required Resources

These are the resources and equipment that are required to complete this activity:

- One PC with Internet access per team member
- Four Cisco Catalyst 3560 Series Switches per team
- Six Cisco 1841 Integrated Services Routers per team
- Three client PCs per team
- One server per team

Lab Setup

The instructor will provide you with directions to prepare the lab equipment for this lab. After the instructor indicates that the lab is fully prepared, you are ready to start the lab.

Scenario

After you reviewed the performance of your team in handling the reported routing problem, your team decided together with your supervisor that they needed to become more familiar with the company network before they can start performing network support and troubleshooting tasks. Therefore, the next task that you have been assigned by your supervisor is to update and supplement the network documentation. This task serves two purposes. It will help you to become familiar with the design and implementation of the company network and it will ensure that you have access to up-to-date and accurate network documentation to reference during future troubleshooting procedures.

Note	In this task, you will have a chance to review and document the baseline configuration of the network. No problems are introduced and you can assume that all documentation that is provided is correct.

Task 1: Assign Responsibilities

In this task, you will assign responsibilities to each team member.

Activity Procedure

Complete these steps:

- Step 1** Review the lab topology together with your team members.
- Step 2** Assign the primary responsibility for each of the devices to a team member. The team member who has primary responsibility for a device is in control of the console of that device and changes to the devices. This means that no other team member should access the console, make changes to the device or execute disruptive actions such as reloading or debugging without permission from the controlling team member. All team members can access all devices via Telnet or SSH for nondisruptive diagnostic action without permission from the controlling member. Responsibilities can be reassigned during later labs if necessary.
- Step 3** Document the responsibilities in the following table.

Device	Responsible team member
ASW1	
CSW1	
CSW2	
IRO1	
IRO2	
CRO1	
CRO2	
BRO1	
BRO2	
BSW1	

Activity Verification


You have completed this task when you attain this result:

- You have assigned responsibility for each of the devices to the team members.

Task 2: Review the Physical Lab Topology


In this task, you will review the lab topology and verify the operation of the core protocols implemented in the lab.

Your supervisor has provided you with a set of diagrams and tables that document the physical connections of the headquarters, WAN, and branch networks.




This figure shows the physical connections of the LAN at headquarters.

Lab 3-1: WAN Physical Topology



This figure shows the physical connections in the WAN.

Lab 3-1: Branch LAN Physical Topology



This figure shows the physical connections in the branch office LAN.

This table lists the VLANs that are used in the LAN at headquarters and the branch LAN.

Location	Description	VLAN	Name	VLAN members
Headquarters	Headquarters LAN			
Headquarters	Floor 1 ASW1 Office VLAN	17	F1S1-OFFICE	ASW1, CSW1, CSW2
Headquarters	Floor 1 ASW1 Voice VLAN	18	F1S1-VOICE	ASW1, CSW1, CSW2
Headquarters	Floor 1 ASW1 Guest VLAN	19	F1S1-GUEST	ASW1, CSW1, CSW2
Headquarters	Floor 1 ASW2 Office VLAN	21	F1S2-OFFICE	CSW1, CSW2
Headquarters	Floor 1 ASW2 Voice VLAN	22	F1S2-VOICE	CSW1, CSW2
Headquarters	Floor 1 ASW2 Guest VLAN	23	F1S2-GUEST	CSW1, CSW2
Headquarters	Floor 1 ASW3 Office VLAN	25	F1S3-OFFICE	CSW1, CSW2
Headquarters	Floor 1 ASW3 Voice VLAN	26	F1S3-VOICE	CSW1, CSW2
Headquarters	Floor 1 ASW3 Guest VLAN	27	F1S3-GUEST	CSW1, CSW2
Headquarters	Floor 2 ASW1 Office VLAN	33	F2S1-OFFICE	CSW1, CSW2
Headquarters	Floor 2 ASW1 Voice VLAN	34	F2S1-VOICE	CSW1, CSW2
Headquarters	Floor 2 ASW1 Guest VLAN	35	F2S1-GUEST	CSW1, CSW2
Headquarters	Floor 2 ASW2 Office VLAN	37	F2S2-OFFICE	CSW1, CSW2
Headquarters	Floor 2 ASW2 Voice VLAN	38	F2S2-VOICE	CSW1, CSW2
Headquarters	Floor 2 ASW2 Guest VLAN	39	F2S2-GUEST	CSW1, CSW2
Headquarters	Floor 2 ASW3 Office VLAN	41	F2S3-OFFICE	CSW1, CSW2
Headquarters	Floor 2 ASW3 Voice VLAN	42	F2S3-VOICE	CSW1, CSW2
Headquarters	Floor 2 ASW3 Guest VLAN	43	F2S3-GUEST	CSW1, CSW2
Headquarters	Internal Servers	112	INT-SERVER	SRV1, CSW1, CSW2
Headquarters	Management VLAN	128	MGMT	ASW1, CSW1, CSW2
Headquarters	Internet Transit LAN	129	TRANSIT	IRO1, IRO2, CSW1, CSW2
Branches	Branch LANs			
Branches	BSW1 Server VLAN	16	B1S1-SERVER	BRO1, BRO2
Branches	BSW1 Office VLAN	17	B1S1-OFFICE	CLT2, BRO1, BRO2
Branches	BSW1 Voice VLAN	18	B1S1-VOICE	BRO1, BRO2
Branches	BSW1 Guest VLAN	19	B1S1-GUEST	CLT3, BRO1, BRO2
Branches	BRO1 - BRO2	30	TRANSIT	BRO1, BRO2
Branches	Management VLAN	128	MGMT	BSW1, BRO1, BRO2
Internet	ISP Metro Links			
Internet	ISP1 FE	11	ISP1	ISP1, IRO1
Internet	ISP2 FE	12	ISP2	ISP2, IRO2

Note	Not all floors and access switches have been implemented at this time. Only access switches ASW1, which resides on floor 1 at headquarters, and BSW1, which resides at the branch office, are present in your lab. The additional VLANs have been provisioned on the core switches CSW1 and CSW2 for future use, but the corresponding access switches are not present. In addition, not all provisioned VLANs have client devices in them. Clients may be moved to different VLANs for testing purposes as required in future exercises.

Activity Procedure

Complete these steps:

- Step 1** Review the lab diagrams. For your convenience, larger versions of these diagrams have been provided in the back of this lab guide.
- Step 2** Use the Cisco Discovery Protocol to verify the physical connection diagrams of your lab pod. The diagrams for the headquarters and WAN networks have been completely filled out, but for the branch office, the interface designators are missing in the diagram. Use the Cisco Discovery Protocol to discover the interfaces that are associated with these links and fill in the correct interface designators.
- Step 3** Verify that all physical links that are shown in the diagram are operational.
- Step 4** Map the VLANs used in the labs to the diagrams. Again, for the headquarters, the VLANs have already been mapped in the diagram. Fill in the missing VLANs for the branch office.
- Step 5** Review the configurations of the devices that you control for use of Layer 1 and Layer 2 features, such as trunks, EtherChannels, and spanning tree. Document these features and discuss your findings with your teammates to ensure that everybody understands the physical design of the network. It is recommended that you review, document, and discuss at least the following aspects of the physical topology:
- The type of spanning tree that is used in the Layer 2 switched domains of the network and the configured spanning-tree priorities and other parameters
 - The resulting spanning-tree topology for all VLANs that have client devices connected
 - The Layer 2 protocols used in the WAN
- Step 6** Document anything that you deem noteworthy about the physical configuration of the devices.

Note	At this point, only physical connections should be examined and documented. Documentation of aspects of the logical topology, such as subnets, IP addresses, and routing protocols do not need to be discovered and documented at this point, but will be addressed during a later part of this lab.

Student Notes

Use this Student Notes section to write down any physical configuration details that you think are important to document for future troubleshooting.

This image shows a single page from a notebook or ledger. It features approximately 20 evenly spaced, thin grey horizontal lines running across its entire width. The background is plain white, and there are no margins, text, or other markings present.

Student Notes

Use this Student Notes section to write down any physical configuration details that you think are important to document for future troubleshooting.

[illegible]

Activity Verification


You have completed this task when you attain these results:

- You have verified that all links shown in the topology diagrams are operational.
- You have discovered and filled in all missing interface designators in the physical topology diagrams.
- You have mapped all host devices such as clients and servers to the VLAN they are a member of.
- You have discovered and documented the spanning-tree topology for all relevant VLANs.
- You have documented all other noteworthy aspects of the physical structure of your lab pod.

Task 3: Review the Logical Lab Topology


In this task, you will review the lab topology and verify the operation of the core protocols implemented in the lab.

Your supervisor has also provided you with a set of diagrams and tables that document the logical connections of the headquarters, WAN, and branch networks.



This figure shows the logical layout of the LAN at the headquarters.

Lab 3-1: WAN Logical Topology




© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0—11

This figure shows the logical layout of the WAN.

Lab 3-1: Branch Logical Topology



© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0—12

This figure shows the logical layout of the branch office LAN.

This table lists the IP subnets that are used in the lab network.

Location	Description	Subnet	Prefix	Devices
Headquarters	Headquarters LAN	10.1.128.0	/19	
Headquarters	Floor 1 ASW1 Office VLAN	10.1.128.64	/26	CLT1, CSW1, CSW2
Headquarters	Floor 1 ASW1 Voice VLAN	10.1.128.128	/26	CSW1, CSW2
Headquarters	Floor 1 ASW1 Guest VLAN	10.1.128.192	/26	CSW1, CSW2
Headquarters	Floor 1 ASW2 Office VLAN	10.1.129.64	/26	CSW1, CSW2
Headquarters	Floor 1 ASW2 Voice VLAN	10.1.129.128	/26	CSW1, CSW2
Headquarters	Floor 1 ASW2 Guest VLAN	10.1.129.192	/26	CSW1, CSW2
Headquarters	Floor 1 ASW3 Office VLAN	10.1.130.64	/26	CSW1, CSW2
Headquarters	Floor 1 ASW3 Voice VLAN	10.1.130.128	/26	CSW1, CSW2
Headquarters	Floor 1 ASW3 Guest VLAN	10.1.130.192	/26	CSW1, CSW2
Headquarters	Floor 2 ASW1 Office VLAN	10.1.132.64	/26	CSW1, CSW2
Headquarters	Floor 2 ASW1 Voice VLAN	10.1.132.128	/26	CSW1, CSW2
Headquarters	Floor 2 ASW1 Guest VLAN	10.1.132.192	/26	CSW1, CSW2
Headquarters	Floor 2 ASW2 Office VLAN	10.1.133.64	/26	CSW1, CSW2
Headquarters	Floor 2 ASW2 Voice VLAN	10.1.133.128	/26	CSW1, CSW2
Headquarters	Floor 2 ASW2 Guest VLAN	10.1.133.192	/26	CSW1, CSW2
Headquarters	Floor 2 ASW3 Office VLAN	10.1.134.64	/26	CSW1, CSW2
Headquarters	Floor 2 ASW3 Voice VLAN	10.1.134.128	/26	CSW1, CSW2
Headquarters	Floor 2 ASW3 Guest VLAN	10.1.134.192	/26	CSW1, CSW2
Headquarters	Internal Servers	10.1.152.0	/24	SRV1, CSW1, CSW2
Headquarters	Management VLAN	10.1.156.0	/22	ASW1, CSW1, CSW2
Branches	Branch LANs	10.1.160.0	/19	
Branches	BSW1 Server VLAN	10.1.160.0	/26	BRO1, BRO2
Branches	BSW1 Office VLAN	10.1.160.64	/26	CLT2, BRO1, BRO2
Branches	BSW1 Voice VLAN	10.1.160.128	/26	BRO1, BRO2
Branches	BSW1 Guest VLAN	10.1.160.192	/26	CLT3, BRO1, BRO2
Branches	BRO1 - BRO2	10.1.163.128	/30	BRO1, BRO2
Branches	Management VLAN	10.1.163.192	/26	BSW1, BRO1, BRO2
WAN	WAN links	10.1.192.0	/19	
Headquarters	CSW1 - CRO1	10.1.192.0	/30	CSW1, CRO1
Headquarters	CSW1 - CRO2	10.1.192.4	/30	CSW1, CRO2
Headquarters	CSW2 - CRO1	10.1.192.8	/30	CSW2, CRO1
Headquarters	CSW2 - CRO2	10.1.192.12	/30	CSW2, CRO2
Headquarters	Internet Transit LAN	10.1.192.16	/29	IRO1, IRO2, CSW1, CSW2

Location	Description	Subnet	Prefix	Devices
WAN	CRO1 - BRO1	10.1.193.0	/30	CRO1, BRO1
WAN	CRO2 - BRO2	10.1.193.4	/30	CRO2, BRO2
WAN	CRO1 - BRO1	10.1.194.0	/30	CRO1, BRO1
WAN	CRO1 - BRO2	10.1.194.4	/30	CRO1, BRO2
WAN	CRO2 - BRO1	10.1.194.8	/30	CRO2, BRO1
WAN	CRO2 - BRO2	10.1.194.12	/30	CRO2, BRO2
WAN	HQ Loopbacks	10.1.220.0	/24	CRO1, CRO2, IRO1, IRO2
WAN	Branch Loopbacks	10.1.221.0	/24	BRO1, BRO2
Internet	ISP1 public block	192.168.224.240	/28	IRO1, ISP1
Internet	ISP2 public block	172.24.244.80	/29	IRO2, ISP2

Note Not all floors and access switches have been implemented at this time. Only access switches ASW1, which resides on floor 1 at headquarters, and BSW1, which resides at the branch office, are present in your lab. The additional subnets have been provisioned on the core switches CSW1 and CSW2 for future use. In addition, not all provisioned subnets have client devices in them. Clients may be moved to different subnets for testing purposes as required in future exercises.

Activity Procedure

Complete these steps:

- Step 1** Review the lab diagrams provided. For your convenience, larger versions of these diagrams have been provided in the back of this lab guide.
- Step 2** Research routing tables and interface IP addresses to map the subnets scheme to the diagrams. The subnets have already been documented on the diagrams, but the host part of the addresses has not been documented. Document the host part of the IP addresses of all devices in the diagrams.

Note Typically, the host part of an IP address can be denoted by the last octet of the full IP address. For example, for IP address 10.1.128.65/26, the host part can be represented as “.65”. For addresses that are part of a subnet that is larger than a /24 prefix, it may be necessary for you to document the last two octets instead of just the last octet.

- Step 3** Review the configurations of the devices that you control and look for the use of control plane features like routing protocols, first-hop redundancy protocols, DHCP and NAT. Discuss your findings with your teammates to ensure that all team members understand the high-level design of the network. It is recommended to review, document, and discuss at least the following aspects of the logical network configuration:
- Use of routing protocols and static routing
 - Use of first-hop redundancy protocols, such as the HSRP, VRRP, and GLBP, including a mapping of the active routers for all relevant VLANs
 - The DHCP servers that are used for all the relevant VLANs present in the logical topology diagrams
 - Any access lists that are used to filter traffic on the network
- Step 4** Document anything that you deem noteworthy about the logical configuration of the devices.

Student Notes

Use this notes section write down any logical configuration details that you think are important to document for future troubleshooting.

This image shows a single page from a notebook or ledger. It features approximately 20 evenly spaced, thin grey horizontal lines running across the width of the page. The background is plain white, and there are no margins, text, or other markings present.

Student Notes

Use this notes section to write down any physical configuration details that you think are important to document for future troubleshooting.

[illegible]

Activity Verification

You have completed this task when you attain these results:

- You have discovered and documented the host part of the IP addresses of all devices in the logical network diagrams.
- You have reviewed and documented the use of routing protocols and static routing in the network.
- You have reviewed the use of DHCP and FHRP in the network and documented the roles of the relevant devices for each subnet.
- You have documented all other noteworthy aspects of the logical structure of your lab pod.

Task 4: Review Troubleshooting and Maintenance Tools

In this task, you will review the lab topology and verify the operation of the core protocols that were implemented in the lab.

Activity Procedure

Complete these steps:

- Step 1** Review the configurations of your assigned devices for features that support troubleshooting and maintenance, such as the use of syslog, SNMP, and other network management features.
- Step 2** Document the features and the corresponding servers and applications or tools in the following table and in the lab diagrams. A sample entry for switch ASW1 has been provided as an example.

Device	Configured feature	Target server	Target tool or application
ASW1	Syslog	SRV1	Syslog server
	DNS	SRV1	DNS server
	Configuration archive	SRV1	TFTP server
	SNMP traps	SRV1	
	NTP	IRO1, IRO2	NTP server
CSW1			
CSW2			
IRO1			
IRO2			

Device	Configured feature	Target server	Target tool or application
CRO1			
CRO2			
BRO1			
BRO2			
BSW1			

Step 3 Discuss your findings with your teammates to ensure that all team members know which maintenance and troubleshooting tools are available in the network.

Step 4 Document anything that you deem noteworthy about the implementation of the tools and services.

Note This is your final chance to document the lab network and create a baseline of it before starting the troubleshooting exercises. Ask your instructor for clarification of any aspects of the network design and configurations that are unclear to you.

Reference Configurations

Your supervisor has provided you with the baseline configurations of all devices for reference during troubleshooting and support tasks.

Switch ASW1

```
!
version 12.2
no service pad
service timestamps debug datetime msec localtime show-timezone
service timestamps log datetime localtime show-timezone
no service password-encryption
!
hostname ASW1
!
boot-start-marker
boot-end-marker
!
logging buffered 16384
logging console warnings
enable secret 5 $1$4mWg$eqWTmSlBAh3uQorxo3HYG/
!
username admin secret 5 $1$0up3$16kazjjd6SgWRRQJqWLSd1
aaa new-model
!
!
aaa authentication login default local
aaa authentication login CONSOLE none
aaa authorization exec default local
!
!
!
aaa session-id common
clock timezone PST -8
clock summer-time PDT recurring
system mtu routing 1500
vtp domain TSHOOT
vtp mode transparent
ip subnet-zero
ip domain-name mgmt.tshoot.local
ip name-server 10.1.152.1
!
!
!
!
crypto pki trustpoint TP-self-signed-2041165184
  enrollment selfsigned
  subject-name cn=IOS-Self-Signed-Certificate-2041165184
  revocation-check none
  rsakeypair TP-self-signed-2041165184
!
!
crypto pki certificate chain TP-self-signed-2041165184
  certificate self-signed 01
    3082024E 308201B7 A0030201 02020101 300D0609 2A864886 F70D0101 04050030
    31312F30 2D060355 04031326 494F532D 53656C66 2D536967 6E65642D 43657274
    69666963 6174652D 32303431 31363531 3834301E 170D3933 30333031 30303030
    34365A17 0D323030 31303130 30303030 305A3031 312F302D 06035504 03132649
    4F532D53 656C662D 5369676E 65642D43 65727469 66696361 74652D32 30343131
    36353138 3430819F 300D0609 2A864886 F70D0101 01050003 818D0030 81890281
    8100DA6B AB52830F 1452A4E8 5293FCD3 D4E86C3F 02C554FA 9254689B 7573618E
    7F9D33C0 54AED542 0ECBFEB0 ED25571C C88F3CD5 6A19D6AF 867D2DD6 B452D058
    124448BD CF02A47F 9C3BA738 390C3C8C CAE69736 50F4EAB1 F93DD618 3F33C52D
    A1A412FF B7402D5C 8116640B 63C5E84E 4700E2EB B363B432 E0D6DCDC 6D89AC83
```

```

DFAD0203 010001A3 76307430 0F060355 1D130101 FF040530 030101FF 30210603
551D1104 1A301882 16415357 312E6D67 6D742E74 73686F6F 742E6C6F 63616C30
1F060355 1D230418 30168014 D7FBBE07 F3148F8D 0B325DDB AA79D569 320738E2
301D0603 551D0E04 160414D7 FBBE07F3 148F8D0B 325DDBAA 79D56932 0738E230
0D06092A 864886F7 0D010104 05000381 8100AF98 3A9C3A39 1566C605 8065B77E
803B66E8 F6BE0B89 348985AB 0D9D1C61 99599006 544CEC0F E6CD44E8 5125B0A7
C2E12429 08AD340D 6B254A5E 66337BF5 BCD1CD75 B6CC4A04 247C1558 69071BDE
8B53B2E6 0D8773FD F29864E8 3F5930B5 8EF65B87 596C9EB4 7F4288D6 B68B4489
09EA4539 7746F761 06E2CBD1 ABB46D42 8726
quit
!
!
!
!
!
archive
log config
logging enable
logging size 50
notify syslog contenttype plaintext
hidekeys
path tftp://srv1.mgmt.tshoot.local/$h-archive-config
write-memory
file prompt quiet
!
spanning-tree mode rapid-pvst
spanning-tree portfast default
spanning-tree extend system-id
!
vlan internal allocation policy ascending
!
vlan 17
name F1S1-OFFICE
!
vlan 18
name F1S1-VOICE
!
vlan 19
name F1S1-GUEST
!
vlan 128
name MGMT
!
vlan 1000
name NATIVE
!
vlan 1001
name UNUSED
!
ip telnet source-interface Vlan128
ip ssh source-interface Vlan128
!
!
!
interface Port-channel1
description Channel to CSW1
switchport trunk encapsulation dot1q
switchport trunk native vlan 1000
switchport trunk allowed vlan 17-19,128
switchport mode trunk
switchport nonegotiate
!
interface Port-channel2
description Channel to CSW2
switchport trunk encapsulation dot1q
switchport trunk native vlan 1000
switchport trunk allowed vlan 17-19,128
switchport mode trunk

```

```

switchport nonegotiate
!
interface FastEthernet0/1
description Channel to CSW1
switchport trunk encapsulation dot1q
switchport trunk native vlan 1000
switchport trunk allowed vlan 17-19,128
switchport mode trunk
switchport nonegotiate
channel-group 1 mode on
!
interface FastEthernet0/2
description Channel to CSW1
switchport trunk encapsulation dot1q
switchport trunk native vlan 1000
switchport trunk allowed vlan 17-19,128
switchport mode trunk
switchport nonegotiate
channel-group 1 mode on
!
interface FastEthernet0/3
description Unused
switchport access vlan 1001
switchport mode access
switchport nonegotiate
shutdown
!
interface FastEthernet0/4
description Unused
switchport access vlan 1001
switchport mode access
switchport nonegotiate
shutdown
!
interface FastEthernet0/5
description Channel to CSW2
switchport trunk encapsulation dot1q
switchport trunk native vlan 1000
switchport trunk allowed vlan 17-19,128
switchport mode trunk
switchport nonegotiate
channel-group 2 mode on
!
interface FastEthernet0/6
description Channel to CSW2
switchport trunk encapsulation dot1q
switchport trunk native vlan 1000
switchport trunk allowed vlan 17-19,128
switchport mode trunk
switchport nonegotiate
channel-group 2 mode on
!
interface FastEthernet0/7
description To CLT1
switchport access vlan 17
switchport mode access
switchport nonegotiate
switchport voice vlan 18
spanning-tree portfast
!
interface FastEthernet0/8
description Unused
switchport access vlan 1001
switchport mode access
switchport nonegotiate
shutdown
!
interface FastEthernet0/9

```

```

description Unused
switchport access vlan 1001
switchport mode access
switchport nonegotiate
shutdown
!
interface FastEthernet0/10
description Unused
switchport access vlan 1001
switchport mode access
switchport nonegotiate
shutdown
!
interface FastEthernet0/11
description Unused
switchport access vlan 1001
switchport mode access
switchport nonegotiate
shutdown
!
interface FastEthernet0/12
description Unused
switchport access vlan 1001
switchport mode access
switchport nonegotiate
shutdown
!
interface FastEthernet0/13
description Unused
switchport access vlan 1001
switchport mode access
switchport nonegotiate
shutdown
!
interface FastEthernet0/14
description Unused
switchport access vlan 1001
switchport mode access
switchport nonegotiate
shutdown
!
interface FastEthernet0/15
description Unused
switchport access vlan 1001
switchport mode access
switchport nonegotiate
shutdown
!
interface FastEthernet0/16
description Unused
switchport access vlan 1001
switchport mode access
switchport nonegotiate
shutdown
!
interface FastEthernet0/17
description Unused
switchport access vlan 1001
switchport mode access
switchport nonegotiate
shutdown
!
interface FastEthernet0/18
description Unused
switchport access vlan 1001
switchport mode access
switchport nonegotiate
shutdown

```

```

!
interface FastEthernet0/19
description Unused
switchport access vlan 1001
switchport mode access
switchport nonegotiate
shutdown
!
interface FastEthernet0/20
description Unused
switchport access vlan 1001
switchport mode access
switchport nonegotiate
shutdown
!
interface FastEthernet0/21
description Unused
switchport access vlan 1001
switchport mode access
switchport nonegotiate
shutdown
!
interface FastEthernet0/22
description Unused
switchport access vlan 1001
switchport mode access
switchport nonegotiate
shutdown
!
interface FastEthernet0/23
description Unused
switchport access vlan 1001
switchport mode access
switchport nonegotiate
shutdown
!
interface FastEthernet0/24
description Unused
switchport access vlan 1001
switchport mode access
switchport nonegotiate
shutdown
!
interface GigabitEthernet0/1
!
interface GigabitEthernet0/2
!
interface Vlan1
no ip address
no ip route-cache
no ip mroute-cache
shutdown
!
interface Vlan128
ip address 10.1.156.1 255.255.252.0
no ip route-cache
no ip mroute-cache
!
ip default-gateway 10.1.159.254
ip classless
ip http server
ip http secure-server
!
!
logging source-interface Vlan128
logging 10.1.152.1
snmp-server community cisco RO
snmp-server community san-fran RW

```

```

snmp-server trap-source Vlan128
snmp-server location TSHOOT Lab Facility
snmp-server contact support@mgmt.tshoot.local
snmp-server enable traps snmp authentication linkdown linkup coldstart
warmstart
snmp-server enable traps transceiver all
snmp-server enable traps tty
snmp-server enable traps eigrp
snmp-server enable traps ospf state-change
snmp-server enable traps ospf errors
snmp-server enable traps ospf retransmit
snmp-server enable traps ospf lsa
snmp-server enable traps ospf cisco-specific state-change nssa-trans-change
snmp-server enable traps ospf cisco-specific state-change shamlink interface-
old
snmp-server enable traps ospf cisco-specific state-change shamlink neighbor
snmp-server enable traps ospf cisco-specific errors
snmp-server enable traps ospf cisco-specific retransmit
snmp-server enable traps ospf cisco-specific lsa
snmp-server enable traps cluster
snmp-server enable traps entity
snmp-server enable traps cpu threshold
snmp-server enable traps power-ethernet group 1
snmp-server enable traps power-ethernet police
snmp-server enable traps rep
snmp-server enable traps vtp
snmp-server enable traps vlancreate
snmp-server enable traps vlandelete
snmp-server enable traps flash insertion removal
snmp-server enable traps port-security
snmp-server enable traps dot1x auth-fail-vlan guest-vlan no-auth-fail-vlan no-
guest-vlan
snmp-server enable traps envmon fan shutdown supply temperature status
snmp-server enable traps bgp
snmp-server enable traps cef resource-failure peer-state-change peer-fib-
state-change inconsistency
snmp-server enable traps config-copy
snmp-server enable traps config
snmp-server enable traps config-ctid
snmp-server enable traps event-manager
snmp-server enable traps hsrp
snmp-server enable traps ipmulticast
snmp-server enable traps msdp
snmp-server enable traps pim neighbor-change rp-mapping-change invalid-pim-
message
snmp-server enable traps bridge newroot topologychange
snmp-server enable traps stpx inconsistency root-inconsistency loop-
inconsistency
snmp-server enable traps syslog
snmp-server enable traps rtr
snmp-server enable traps mvpn
snmp-server enable traps mac-notification change move threshold
snmp-server enable traps vlan-membership
snmp-server enable traps errdisable
snmp-server host 10.1.152.1 version 2c cisco
snmp ifmib ifindex persist
!
control-plane
!
line con 0
  exec-timeout 60 0
  login authentication CONSOLE
line vty 0 4
  transport input telnet ssh
line vty 5 15
  transport input telnet ssh
!
ntp source Vlan128

```

```
ntp server 10.1.220.4
ntp server 10.1.220.3
end
```

Router BRO1

```
!
version 12.4
service timestamps debug datetime msec localtime show-timezone
service timestamps log datetime localtime show-timezone
no service password-encryption
!
hostname BRO1
!
boot-start-marker
boot-end-marker
!
logging buffered 16384 debugging
logging console warnings
enable secret 5 $1$4mWg$eqWTmSlBAh3uQorxo3HYG/
!
aaa new-model
!
!
aaa authentication login default local
aaa authentication login CONSOLE none
aaa authorization exec default local
!
aaa session-id common
clock timezone PST -8
clock summer-time PDT recurring
ip cef
!
!
no ip dhcp use vrf connected
ip dhcp excluded-address 10.1.160.124 10.1.160.126
ip dhcp excluded-address 10.1.160.252 10.1.160.254
!
ip dhcp pool B1S1-OFFICE
    network 10.1.160.64 255.255.255.192
    default-router 10.1.160.126
    dns-server 10.1.152.1
    domain-name tshoot.local
!
ip dhcp pool B1S1-GUEST
    network 10.1.160.192 255.255.255.192
    default-router 10.1.160.254
    dns-server 10.1.152.1
    domain-name tshoot.local
!
!
ip domain name mgmt.tshoot.local
ip name-server 10.1.152.1
ip auth-proxy max-nodata-conns 3
ip admission max-nodata-conns 3
!
!
!
file prompt quiet
username admin secret 5 $1$0up3$16kazjjd6SgWRRQJqWLSd1
archive
    log config
    logging enable
    logging size 50
    notify syslog
    hidekeys
    path tftp://srv1.mgmt.tshoot.local/$h-archive-config
write-memory
```

```

!
!
ip telnet source-interface Loopback0
ip ssh source-interface Loopback0
!
!
!
!
interface Loopback0
 ip address 10.1.221.1 255.255.255.255
!
interface FastEthernet0/0
 no ip address
 shutdown
 duplex auto
 speed auto
!
interface FastEthernet0/1
 description FE to BSW1
 no ip address
 speed 100
 full-duplex
!
interface FastEthernet0/1.16
 description Vlan B1S1-SERVER
 encapsulation dot1Q 16
 ip address 10.1.160.60 255.255.255.192
 glbp 16 ip 10.1.160.62
 glbp 16 priority 110
 glbp 16 preempt
!
interface FastEthernet0/1.17
 description Vlan B1S1-OFFICE
 encapsulation dot1Q 17
 ip address 10.1.160.124 255.255.255.192
 glbp 17 ip 10.1.160.126
 glbp 17 preempt
!
interface FastEthernet0/1.18
 description Vlan B1S1-VOICE
 encapsulation dot1Q 18
 ip address 10.1.160.188 255.255.255.192
 glbp 18 ip 10.1.160.190
 glbp 18 priority 110
 glbp 18 preempt
!
interface FastEthernet0/1.19
 description Vlan B1S1-GUEST
 encapsulation dot1Q 19
 ip address 10.1.160.252 255.255.255.192
 ip access-group LIMIT-GUEST-ACCESS in
 glbp 19 ip 10.1.160.254
 glbp 19 preempt
!
interface FastEthernet0/1.30
 description Vlan TRANSIT to BRO2
 encapsulation dot1Q 30
 ip address 10.1.163.129 255.255.255.252
!
interface FastEthernet0/1.128
 description Vlan MGMT
 encapsulation dot1Q 128
 ip address 10.1.163.252 255.255.255.192
 glbp 128 ip 10.1.163.254
 glbp 128 priority 110
 glbp 128 preempt
!
interface Serial0/0/0

```



```

description 128 kbps to Frame Relay
bandwidth 128
no ip address
encapsulation frame-relay
no fair-queue
!
interface Serial0/0/0.111 point-to-point
description PVC to CRO1
bandwidth 64
ip address 10.1.194.2 255.255.255.252
frame-relay interface-dlci 111
!
interface Serial0/0/0.112 point-to-point
description PVC to CRO2
bandwidth 64
ip address 10.1.194.10 255.255.255.252
frame-relay interface-dlci 112
!
interface Serial0/0/1
description 128kbps leased line to CRO1
bandwidth 128
ip address 10.1.193.2 255.255.255.252
encapsulation ppp
!
router eigrp 1
passive-interface default
no passive-interface FastEthernet0/1.30
no passive-interface Serial0/0/0.111
no passive-interface Serial0/0/0.112
no passive-interface Serial0/0/1
network 10.1.160.0 0.0.3.255
network 10.1.193.2 0.0.0.0
network 10.1.194.2 0.0.0.0
network 10.1.194.10 0.0.0.0
network 10.1.221.1 0.0.0.0
no auto-summary
!
ip forward-protocol nd
!
!
ip http server
no ip http secure-server
!
ip access-list extended LIMIT-GUEST-ACCESS
permit icmp any any
permit udp any any eq 3222
permit udp any host 10.1.152.1 eq domain
permit tcp any host 10.1.152.1 eq domain
permit udp any any eq bootpc
deny ip any 10.1.128.0 0.0.127.255
permit ip any any
!
logging source-interface Loopback0
logging 10.1.152.1
snmp-server community cisco RO
snmp-server community san-fran RW
snmp-server ifindex persist
snmp-server trap-source Loopback0
snmp-server location TSHOOT Lab Facility
snmp-server contact support@mgmt.tsshoot.local
snmp-server enable traps snmp authentication linkdown linkup coldstart
warmstart
snmp-server enable traps vrrp
snmp-server enable traps ds1
snmp-server enable traps tty
snmp-server enable traps eigrp
snmp-server enable traps envmon
snmp-server enable traps flash insertion removal

```

```

snmp-server enable traps icsudsu
snmp-server enable traps isdn call-information
snmp-server enable traps isdn layer2
snmp-server enable traps isdn chan-not-avail
snmp-server enable traps isdn ietf
snmp-server enable traps ds0-busyout
snmp-server enable traps ds1-loopback
snmp-server enable traps aaa_server
snmp-server enable traps atm subif
snmp-server enable traps bgp
snmp-server enable traps bulkstat collection transfer
snmp-server enable traps cnpd
snmp-server enable traps config-copy
snmp-server enable traps config
snmp-server enable traps dsp card-status
snmp-server enable traps entity
snmp-server enable traps event-manager
snmp-server enable traps frame-relay
snmp-server enable traps frame-relay subif
snmp-server enable traps hsrp
snmp-server enable traps ipmobile
snmp-server enable traps ipmulticast
snmp-server enable traps msdp
snmp-server enable traps mvpn
snmp-server enable traps ospf state-change
snmp-server enable traps ospf errors
snmp-server enable traps ospf retransmit
snmp-server enable traps ospf lsa
snmp-server enable traps ospf cisco-specific state-change nssa-trans-change
snmp-server enable traps ospf cisco-specific state-change shamlink interface-
old
snmp-server enable traps ospf cisco-specific state-change shamlink neighbor
snmp-server enable traps ospf cisco-specific errors
snmp-server enable traps ospf cisco-specific retransmit
snmp-server enable traps ospf cisco-specific lsa
snmp-server enable traps pim neighbor-change rp-mapping-change invalid-pim-
message
snmp-server enable traps pppoe
snmp-server enable traps cpu threshold
snmp-server enable traps rsvp
snmp-server enable traps rtr
snmp-server enable traps syslog
snmp-server enable traps l2tun session
snmp-server enable traps vtp
snmp-server enable traps isakmp policy add
snmp-server enable traps isakmp policy delete
snmp-server enable traps isakmp tunnel start
snmp-server enable traps isakmp tunnel stop
snmp-server enable traps ipsec cryptomap add
snmp-server enable traps ipsec cryptomap delete
snmp-server enable traps ipsec cryptomap attach
snmp-server enable traps ipsec cryptomap detach
snmp-server enable traps ipsec tunnel start
snmp-server enable traps ipsec tunnel stop
snmp-server enable traps ipsec too-many-sas
snmp-server host 10.1.152.1 version 2c cisco
!
!
!
!
control-plane
!
!
line con 0
  exec-timeout 60 0
  login authentication CONSOLE
line aux 0
line vty 0 4

```

```

transport input telnet ssh
line vty 5 15
transport input telnet ssh
!
scheduler allocate 20000 1000
ntp source Loopback0
ntp update-calendar
ntp server 10.1.220.4
ntp server 10.1.220.3
end

```

Router BRO2

```

!
version 12.4
service timestamps debug datetime msec localtime show-timezone
service timestamps log datetime localtime show-timezone
no service password-encryption
!
hostname BRO2
!
boot-start-marker
boot-end-marker
!
logging buffered 16384 debugging
logging console warnings
enable secret 5 $1$4mWg$eqWTmSlBAh3uQorxo3HYG/
!
aaa new-model
!
!
aaa authentication login default local
aaa authentication login CONSOLE none
aaa authorization exec default local
!
aaa session-id common
clock timezone PST -8
clock summer-time PDT recurring
ip cef
!
!
no ip dhcp use vrf connected
ip dhcp excluded-address 10.1.160.188 10.1.160.190
!
ip dhcp pool B1S1-VOICE
    network 10.1.160.128 255.255.255.192
    default-router 10.1.160.190
    dns-server 10.1.152.1
    domain-name tshoot.local
!
!
ip domain name mgmt.tshoot.local
ip name-server 10.1.152.1
ip auth-proxy max-nodata-conns 3
ip admission max-nodata-conns 3
!
!
!
file prompt quiet
username admin secret 5 $1$0up3$16kazjjd6SgWRRQJqWLSd1
archive
    log config
    logging enable
    logging size 50
    notify syslog
    hidekeys
    path tftp://srv1.mgmt.tshoot.local/$h-archive-config
write-memory

```

```

!
!
ip telnet source-interface Loopback0
ip ssh source-interface Loopback0
!
!
!
!
interface Loopback0
 ip address 10.1.221.2 255.255.255.255
!
interface FastEthernet0/0
 no ip address
 shutdown
 duplex auto
 speed auto
!
interface FastEthernet0/1
 description FE to BSW1
 no ip address
 speed 100
 full-duplex
!
interface FastEthernet0/1.16
 description Vlan B1S1-SERVER
 encapsulation dot1Q 16
 ip address 10.1.160.61 255.255.255.192
 glbp 16 ip 10.1.160.62
 glbp 16 preempt
!
interface FastEthernet0/1.17
 description Vlan B1S1-OFFICE
 encapsulation dot1Q 17
 ip address 10.1.160.125 255.255.255.192
 glbp 17 ip 10.1.160.126
 glbp 17 priority 110
 glbp 17 preempt
!
interface FastEthernet0/1.18
 description Vlan B1S1-VOICE
 encapsulation dot1Q 18
 ip address 10.1.160.189 255.255.255.192
 glbp 18 ip 10.1.160.190
 glbp 18 preempt
!
interface FastEthernet0/1.19
 description Vlan B1S1-GUEST
 encapsulation dot1Q 19
 ip address 10.1.160.253 255.255.255.192
 ip access-group LIMIT-GUEST-ACCESS in
 glbp 19 ip 10.1.160.254
 glbp 19 priority 110
 glbp 19 preempt
!
interface FastEthernet0/1.30
 description Vlan TRANSIT to BRO1
 encapsulation dot1Q 30
 ip address 10.1.163.130 255.255.255.252
!
interface FastEthernet0/1.128
 description Vlan MGMT
 encapsulation dot1Q 128
 ip address 10.1.163.253 255.255.255.192
 glbp 128 ip 10.1.163.254
 glbp 128 preempt
!
interface Serial0/0/0
 description 128 kbps to Frame Relay

```

```

bandwidth 128
no ip address
encapsulation frame-relay
no fair-queue
!
interface Serial0/0/0.111 point-to-point
description PVC to CRO1
bandwidth 64
ip address 10.1.194.6 255.255.255.252
frame-relay interface-dlci 111
!
interface Serial0/0/0.112 point-to-point
description PVC to CRO2
bandwidth 64
ip address 10.1.194.14 255.255.255.252
frame-relay interface-dlci 112
!
interface Serial0/0/1
description 128 kbps leased line to CRO2
bandwidth 128
ip address 10.1.193.6 255.255.255.252
encapsulation ppp
!
router eigrp 1
passive-interface default
no passive-interface FastEthernet0/1.30
no passive-interface Serial0/0/0.111
no passive-interface Serial0/0/0.112
no passive-interface Serial0/0/1
network 10.1.160.0 0.0.3.255
network 10.1.193.6 0.0.0.0
network 10.1.194.6 0.0.0.0
network 10.1.194.14 0.0.0.0
network 10.1.221.2 0.0.0.0
no auto-summary
!
ip forward-protocol nd
!
!
ip http server
no ip http secure-server
!
ip access-list extended LIMIT-GUEST-ACCESS
permit icmp any any
permit udp any any eq 3222
permit udp any host 10.1.152.1 eq domain
permit tcp any host 10.1.152.1 eq domain
permit udp any any eq bootpc
deny ip any 10.1.128.0 0.0.127.255
permit ip any any
!
logging source-interface Loopback0
logging 10.1.152.1
snmp-server community cisco RO
snmp-server community san-fran RW
snmp-server ifindex persist
snmp-server trap-source Loopback0
snmp-server location TSHOOT Lab Facility
snmp-server contact support@mgmt.tshoot.local
snmp-server enable traps snmp authentication linkdown linkup coldstart
warmstart
snmp-server enable traps vrrp
snmp-server enable traps dsl
snmp-server enable traps tty
snmp-server enable traps eigrp
snmp-server enable traps envmon
snmp-server enable traps flash insertion removal
snmp-server enable traps icsudsu

```

```

snmp-server enable traps isdn call-information
snmp-server enable traps isdn layer2
snmp-server enable traps isdn chan-not-avail
snmp-server enable traps isdn ietf
snmp-server enable traps ds0-busyout
snmp-server enable traps ds1-loopback
snmp-server enable traps aaa_server
snmp-server enable traps atm subif
snmp-server enable traps bgp
snmp-server enable traps bulkstat collection transfer
snmp-server enable traps cnpd
snmp-server enable traps config-copy
snmp-server enable traps config
snmp-server enable traps dsp card-status
snmp-server enable traps entity
snmp-server enable traps event-manager
snmp-server enable traps frame-relay
snmp-server enable traps frame-relay subif
snmp-server enable traps hsrp
snmp-server enable traps ipmobile
snmp-server enable traps ipmulticast
snmp-server enable traps msdp
snmp-server enable traps mvpn
snmp-server enable traps ospf state-change
snmp-server enable traps ospf errors
snmp-server enable traps ospf retransmit
snmp-server enable traps ospf lsa
snmp-server enable traps ospf cisco-specific state-change nssa-trans-change
snmp-server enable traps ospf cisco-specific state-change shamlink interface-
old
snmp-server enable traps ospf cisco-specific state-change shamlink neighbor
snmp-server enable traps ospf cisco-specific errors
snmp-server enable traps ospf cisco-specific retransmit
snmp-server enable traps ospf cisco-specific lsa
snmp-server enable traps pim neighbor-change rp-mapping-change invalid-pim-
message
snmp-server enable traps pppoe
snmp-server enable traps cpu threshold
snmp-server enable traps rsvp
snmp-server enable traps rtr
snmp-server enable traps syslog
snmp-server enable traps l2tun session
snmp-server enable traps vtp
snmp-server enable traps isakmp policy add
snmp-server enable traps isakmp policy delete
snmp-server enable traps isakmp tunnel start
snmp-server enable traps isakmp tunnel stop
snmp-server enable traps ipsec cryptomap add
snmp-server enable traps ipsec cryptomap delete
snmp-server enable traps ipsec cryptomap attach
snmp-server enable traps ipsec cryptomap detach
snmp-server enable traps ipsec tunnel start
snmp-server enable traps ipsec tunnel stop
snmp-server enable traps ipsec too-many-sas
snmp-server host 10.1.152.1 version 2c cisco
!
!
!
!
control-plane
!
!
line con 0
  exec-timeout 60 0
  login authentication CONSOLE
line aux 0
line vty 0 4
  transport input telnet ssh

```

```

line vty 5 15
  transport input telnet ssh
!
scheduler allocate 20000 1000
ntp source Loopback0
ntp update-calendar
ntp server 10.1.220.4
ntp server 10.1.220.3
end

```

Switch BSW1

```

!
version 12.2
no service pad
service timestamps debug datetime msec localtime show-timezone
service timestamps log datetime localtime show-timezone
no service password-encryption
!
hostname BSW1
!
boot-start-marker
boot-end-marker
!
logging buffered 16384
logging console warnings
enable secret 5 $1$4mWg$eqWTmSlBAh3uQorxo3HYG/
!
username admin secret 5 $1$0up3$16kazjjd6SgWRRQJqWLSd1
aaa new-model
!
!
aaa authentication login default local
aaa authentication login CONSOLE none
aaa authorization exec default local
!
!
!
aaa session-id common
clock timezone PST -8
clock summer-time PDT recurring
system mtu routing 1500
vtp domain TSHOOT
vtp mode transparent
ip subnet-zero
ip domain-name mgmt.tshoot.local
ip name-server 10.1.152.1
!
!
!
!
crypto pki trustpoint TP-self-signed-656508160
  enrollment selfsigned
  subject-name cn=IOS-Self-Signed-Certificate-656508160
  revocation-check none
  rsa-keypair TP-self-signed-656508160
!
!
crypto pki certificate chain TP-self-signed-656508160
  certificate self-signed 01
    3082024C 308201B5 A0030201 02020101 300D0609 2A864886 F70D0101 04050030
    30312E30 2C060355 04031325 494F532D 53656C66 2D536967 6E65642D 43657274
    69666963 6174652D 36353635 30383136 30301E17 0D393330 33303130 30303034
    355A170D 32303031 30313030 30303030 5A303031 2E302C06 03550403 1325494F
    532D5365 6C662D53 69676E65 642D4365 72746966 69636174 652D3635 36353038
    31363030 819F300D 06092A86 4886F70D 01010105 0003818D 00308189 02818100
    EAB798C8 5F70F991 FE8DFC36 F2C54F09 4D6D0115 EC6D4F67 BF0BA9BF E71BBA58
    C453F6EB 94133D9B E98B6176 70294CA0 5024DF3E 9401BD57 988DD527 1DE3BA23

```

```

097F72DF DB73C3FF B937A97F 0D087BFC 94E62215 2605E001 3026BB29 8E3543C2
459217A5 6AA4E376 56ABA350 7C075259 DD48479C E94B91F8 749ECF76 B8216FFB
02030100 01A37630 74300F06 03551D13 0101FF04 05300301 01FF3021 0603551D
11041A30 18821642 5357312E 6D676D74 2E747368 6F6F742E 6C6F6361 6C301F06
03551D23 04183016 8014BD44 BE1649F8 3E88B070 6535BE96 153FA046 7D9E301D
0603551D 0E041604 14BD44BE 1649F83E 88B07065 35BE9615 3FA0467D 9E300D06
092A8648 86F70D01 01040500 03818100 A54A7CF7 A1D18399 2801A531 35E3C102
54802487 C6202E52 0D0684B2 022F1583 38E6C544 EFA6B71D D9EE8219 F794B2C6
1F924A0A 55E52276 4BBA6840 59C8DBA7 E3F94CB6 EE3D1514 1C2FFB3C 72A37BA2
B00B46E0 737B2380 6B53206F 74BD5F00 86D379DF D4EBA967 943932BE 9F8CB4CE
F9DD444F AD563232 4D5B1AE0 0D0EC529
quit
!
!
!
!
!
archive
log config
logging enable
logging size 50
notify syslog contenttype plaintext
hidekeys
path tftp://srv1.mgmt.tshoot.local/$h-archive-config
write-memory
file prompt quiet
!
spanning-tree mode rapid-pvst
spanning-tree portfast default
spanning-tree extend system-id
!
vlan internal allocation policy ascending
!
vlan 16
name B1S1-SERVER
!
vlan 17
name B1S1-OFFICE
!
vlan 18
name B1S1-VOICE
!
vlan 19
name B1S1-GUEST
!
vlan 30
name TRANSIT
!
vlan 128
name MGMT
!
vlan 1000
name NATIVE
!
vlan 1001
name UNUSED
!
ip telnet source-interface Vlan128
ip ssh source-interface Vlan128
!
!
!
interface FastEthernet0/1
description Unused
switchport access vlan 1001
switchport mode access
switchport nonegotiate
shutdown

```



```

!
interface FastEthernet0/2
description Unused
switchport access vlan 1001
switchport mode access
switchport nonegotiate
shutdown
!
interface FastEthernet0/3
description FE to BR01
switchport trunk encapsulation dot1q
switchport trunk native vlan 1000
switchport trunk allowed vlan 16-19,30,128
switchport mode trunk
switchport nonegotiate
speed 100
duplex full
spanning-tree portfast trunk
!
interface FastEthernet0/4
description FE to BR02
switchport trunk encapsulation dot1q
switchport trunk native vlan 1000
switchport trunk allowed vlan 16-19,30,128
switchport mode trunk
switchport nonegotiate
speed 100
duplex full
spanning-tree portfast trunk
!
interface FastEthernet0/5
description Unused
switchport access vlan 1001
switchport mode access
switchport nonegotiate
shutdown
!
interface FastEthernet0/6
description Unused
switchport access vlan 1001
switchport mode access
switchport nonegotiate
shutdown
!
interface FastEthernet0/7
description To CLT2
switchport access vlan 17
switchport mode access
switchport nonegotiate
switchport voice vlan 18
spanning-tree portfast
!
interface FastEthernet0/8
description To CLT3
switchport access vlan 19
switchport mode access
switchport nonegotiate
switchport voice vlan 18
spanning-tree portfast
!
interface FastEthernet0/9
description Unused
switchport access vlan 1001
switchport mode access
switchport nonegotiate
shutdown
!
interface FastEthernet0/10

```

```

description Unused
switchport access vlan 1001
switchport mode access
switchport nonegotiate
shutdown
!
interface FastEthernet0/11
description Unused
switchport access vlan 1001
switchport mode access
switchport nonegotiate
shutdown
!
interface FastEthernet0/12
description Unused
switchport access vlan 1001
switchport mode access
switchport nonegotiate
shutdown
!
interface FastEthernet0/13
description Unused
switchport access vlan 1001
switchport mode access
switchport nonegotiate
shutdown
!
interface FastEthernet0/14
description Unused
switchport access vlan 1001
switchport mode access
switchport nonegotiate
shutdown
!
interface FastEthernet0/15
description Unused
switchport access vlan 1001
switchport mode access
switchport nonegotiate
shutdown
!
interface FastEthernet0/16
description Unused
switchport access vlan 1001
switchport mode access
switchport nonegotiate
shutdown
!
interface FastEthernet0/17
description Unused
switchport access vlan 1001
switchport mode access
switchport nonegotiate
shutdown
!
interface FastEthernet0/18
description Unused
switchport access vlan 1001
switchport mode access
switchport nonegotiate
shutdown
!
interface FastEthernet0/19
description Unused
switchport access vlan 1001
switchport mode access
switchport nonegotiate
shutdown

```

```

!
interface FastEthernet0/20
description Unused
switchport access vlan 1001
switchport mode access
switchport nonegotiate
shutdown
!
interface FastEthernet0/21
description Unused
switchport access vlan 1001
switchport mode access
switchport nonegotiate
shutdown
!
interface FastEthernet0/22
description Unused
switchport access vlan 1001
switchport mode access
switchport nonegotiate
shutdown
!
interface FastEthernet0/23
description Unused
switchport access vlan 1001
switchport mode access
switchport nonegotiate
shutdown
!
interface FastEthernet0/24
description Unused
switchport access vlan 1001
switchport mode access
switchport nonegotiate
shutdown
!
interface GigabitEthernet0/1
!
interface GigabitEthernet0/2
!
interface Vlan1
no ip address
shutdown
!
interface Vlan128
ip address 10.1.163.193 255.255.255.192
!
ip default-gateway 10.1.163.254
ip classless
ip http server
ip http secure-server
!
!
logging source-interface Vlan128
logging 10.1.152.1
snmp-server community cisco RO
snmp-server community san-fran RW
snmp-server trap-source Vlan128
snmp-server location TSHOOT Lab Facility
snmp-server contact support@mgmt.tshoot.local
snmp-server enable traps snmp authentication linkdown linkup coldstart
warmstart
snmp-server enable traps transceiver all
snmp-server enable traps tty
snmp-server enable traps eigrp
snmp-server enable traps ospf state-change
snmp-server enable traps ospf errors
snmp-server enable traps ospf retransmit

```

```

snmp-server enable traps ospf lsa
snmp-server enable traps ospf cisco-specific state-change nssa-trans-change
snmp-server enable traps ospf cisco-specific state-change shamlink interface-
old
snmp-server enable traps ospf cisco-specific state-change shamlink neighbor
snmp-server enable traps ospf cisco-specific errors
snmp-server enable traps ospf cisco-specific retransmit
snmp-server enable traps ospf cisco-specific lsa
snmp-server enable traps cluster
snmp-server enable traps entity
snmp-server enable traps cpu threshold
snmp-server enable traps power-ethernet group 1
snmp-server enable traps power-ethernet police
snmp-server enable traps rep
snmp-server enable traps vtp
snmp-server enable traps vlancreate
snmp-server enable traps vlandelete
snmp-server enable traps flash insertion removal
snmp-server enable traps port-security
snmp-server enable traps dot1x auth-fail-vlan guest-vlan no-auth-fail-vlan no-
guest-vlan
snmp-server enable traps envmon fan shutdown supply temperature status
snmp-server enable traps bgp
snmp-server enable traps cef resource-failure peer-state-change peer-fib-
state-change inconsistency
snmp-server enable traps config-copy
snmp-server enable traps config
snmp-server enable traps config-ctid
snmp-server enable traps event-manager
snmp-server enable traps hsrp
snmp-server enable traps ipmulticast
snmp-server enable traps msdp
snmp-server enable traps pim neighbor-change rp-mapping-change invalid-pim-
message
snmp-server enable traps bridge newroot topologychange
snmp-server enable traps stpx inconsistency root-inconsistency loop-
inconsistency
snmp-server enable traps syslog
snmp-server enable traps rtr
snmp-server enable traps mvpn
snmp-server enable traps mac-notification change move threshold
snmp-server enable traps vlan-membership
snmp-server enable traps errdisable
snmp-server host 10.1.152.1 version 2c cisco
snmp ifmib ifindex persist
!
control-plane
!
line con 0
  exec-timeout 60 0
  login authentication CONSOLE
line vty 0 4
  transport input telnet ssh
line vty 5 15
  transport input telnet ssh
!
ntp source Vlan128
ntp server 10.1.220.4
ntp server 10.1.220.3
end

```

Router CRO1

```

!
version 12.4
service timestamps debug datetime msec localtime show-timezone
service timestamps log datetime localtime show-timezone
no service password-encryption

```

```

!
hostname CR01
!
boot-start-marker
boot-end-marker
!
logging buffered 16384 debugging
logging console warnings
enable secret 5 $1$4mWg$eqWTmS1BAh3uQorxo3HYG/
!
aaa new-model
!
!
aaa authentication login default local
aaa authentication login CONSOLE none
aaa authorization exec default local
!
aaa session-id common
clock timezone PST -8
clock summer-time PDT recurring
ip cef
!
!
!
!
ip domain name mgmt.tshoot.local
ip name-server 10.1.152.1
ip auth-proxy max-nodata-conns 3
ip admission max-nodata-conns 3
!
!
!
file prompt quiet
username admin secret 5 $1$0up3$16kazjdd6SgWRRQJqWLSd1
archive
  log config
  logging enable
  logging size 50
  notify syslog
  hidekeys
  path tftp://srv1.mgmt.tshoot.local/$h-archive-config
  write-memory
!
!
ip telnet source-interface Loopback0
ip ssh source-interface Loopback0
!
!
!
!
interface Loopback0
  ip address 10.1.220.1 255.255.255.255
!
interface FastEthernet0/0
  description FE to CSW1
  ip address 10.1.192.2 255.255.255.252
  ip flow ingress
  speed 100
  full-duplex
!
interface FastEthernet0/1
  description FE to CSW2
  ip address 10.1.192.10 255.255.255.252
  ip flow ingress
  speed 100
  full-duplex
!
interface Serial0/0/0

```

```

description 128 kbps to Frame Relay
bandwidth 128
no ip address
ip flow ingress
encapsulation frame-relay
no fair-queue
!
interface Serial0/0/0.121 point-to-point
description PVC to BRO1
bandwidth 64
ip address 10.1.194.1 255.255.255.252
frame-relay interface-dlci 121
!
interface Serial0/0/0.122 point-to-point
description PVC to BRO2
bandwidth 64
ip address 10.1.194.5 255.255.255.252
frame-relay interface-dlci 122
!
interface Serial0/0/1
description 128 kbps leased line to BRO1
bandwidth 128
ip address 10.1.193.1 255.255.255.252
ip flow ingress
encapsulation ppp
clock rate 128000
!
router eigrp 1
network 10.1.192.2 0.0.0.0
network 10.1.192.10 0.0.0.0
network 10.1.193.1 0.0.0.0
network 10.1.194.1 0.0.0.0
network 10.1.194.5 0.0.0.0
network 10.1.220.1 0.0.0.0
no auto-summary
!
ip forward-protocol nd
!
ip flow-export source Loopback0
ip flow-export version 5
ip flow-export destination 10.1.152.1 9996
!
ip http server
no ip http secure-server
!
logging source-interface Loopback0
logging 10.1.152.1
snmp-server community cisco RO
snmp-server community san-fran RW
snmp-server ifindex persist
snmp-server trap-source Loopback0
snmp-server location TSHOOT Lab Facility
snmp-server contact support@mgmt.tshoot.local
snmp-server enable traps snmp authentication linkdown linkup coldstart
warmstart
snmp-server enable traps vrrp
snmp-server enable traps ds1
snmp-server enable traps tty
snmp-server enable traps eigrp
snmp-server enable traps envmon
snmp-server enable traps flash insertion removal
snmp-server enable traps icsudsu
snmp-server enable traps isdn call-information
snmp-server enable traps isdn layer2
snmp-server enable traps isdn chan-not-avail
snmp-server enable traps isdn ietf
snmp-server enable traps ds0-busyout
snmp-server enable traps ds1-loopback

```

```

snmp-server enable traps aaa_server
snmp-server enable traps atm subif
snmp-server enable traps bgp
snmp-server enable traps bulkstat collection transfer
snmp-server enable traps cnpd
snmp-server enable traps config-copy
snmp-server enable traps config
snmp-server enable traps dsp card-status
snmp-server enable traps entity
snmp-server enable traps event-manager
snmp-server enable traps frame-relay
snmp-server enable traps frame-relay subif
snmp-server enable traps hsrp
snmp-server enable traps ipmobile
snmp-server enable traps ipmulticast
snmp-server enable traps msdp
snmp-server enable traps mvpn
snmp-server enable traps ospf state-change
snmp-server enable traps ospf errors
snmp-server enable traps ospf retransmit
snmp-server enable traps ospf lsa
snmp-server enable traps ospf cisco-specific state-change nssa-trans-change
snmp-server enable traps ospf cisco-specific state-change shamlink interface-
old
snmp-server enable traps ospf cisco-specific state-change shamlink neighbor
snmp-server enable traps ospf cisco-specific errors
snmp-server enable traps ospf cisco-specific retransmit
snmp-server enable traps ospf cisco-specific lsa
snmp-server enable traps pim neighbor-change rp-mapping-change invalid-pim-
message
snmp-server enable traps pppoe
snmp-server enable traps cpu threshold
snmp-server enable traps rsvp
snmp-server enable traps rtr
snmp-server enable traps syslog
snmp-server enable traps l2tun session
snmp-server enable traps vtp
snmp-server enable traps isakmp policy add
snmp-server enable traps isakmp policy delete
snmp-server enable traps isakmp tunnel start
snmp-server enable traps isakmp tunnel stop
snmp-server enable traps ipsec cryptomap add
snmp-server enable traps ipsec cryptomap delete
snmp-server enable traps ipsec cryptomap attach
snmp-server enable traps ipsec cryptomap detach
snmp-server enable traps ipsec tunnel start
snmp-server enable traps ipsec tunnel stop
snmp-server enable traps ipsec too-many-sas
snmp-server host 10.1.152.1 version 2c cisco
!
!
!
!
control-plane
!
!
line con 0
  exec-timeout 60 0
  login authentication CONSOLE
line aux 0
line vty 0 4
  transport input telnet ssh
line vty 5 15
  transport input telnet ssh
!
scheduler allocate 20000 1000
ntp source Loopback0
ntp update-calendar

```

```

ntp server 10.1.220.4
ntp server 10.1.220.3
end

```

Router CRO2

```

!
version 12.4
service timestamps debug datetime msec localtime show-timezone
service timestamps log datetime localtime show-timezone
no service password-encryption
!
hostname CRO2
!
boot-start-marker
boot-end-marker
!
logging buffered 16384 debugging
logging console warnings
enable secret 5 $1$4mWg$eqWTmSlBAh3uQorxo3HYG/
!
aaa new-model
!
!
aaa authentication login default local
aaa authentication login CONSOLE none
aaa authorization exec default local
!
aaa session-id common
clock timezone PST -8
clock summer-time PDT recurring
ip cef
!
!
!
!
ip domain name mgmt.tshoot.local
ip name-server 10.1.152.1
ip auth-proxy max-nodata-conns 3
ip admission max-nodata-conns 3
!
!
!
file prompt quiet
username admin secret 5 $1$0up3$16kazjld6SgWRRQJqWLSd1
archive
  log config
    logging enable
    logging size 50
    notify syslog
    hidekeys
  path tftp://srv1.mgmt.tshoot.local/$h-archive-config
  write-memory
!
!
ip telnet source-interface Loopback0
ip ssh source-interface Loopback0
!
!
!
interface Loopback0
  ip address 10.1.220.2 255.255.255.255
!
interface FastEthernet0/0
  description FE to CSW1
  ip address 10.1.192.6 255.255.255.252
  ip flow ingress

```



```

speed 100
full-duplex
!
interface FastEthernet0/1
description FE to CSW2
ip address 10.1.192.14 255.255.255.252
ip flow ingress
speed 100
full-duplex
!
interface Serial0/0/0
description 128 kbps to Frame Relay
bandwidth 128
no ip address
ip flow ingress
encapsulation frame-relay
no fair-queue
!
interface Serial0/0/0.121 point-to-point
description PVC to BRO1
bandwidth 64
ip address 10.1.194.9 255.255.255.252
frame-relay interface-dlci 121
!
interface Serial0/0/0.122 point-to-point
description PVC to BRO2
bandwidth 64
ip address 10.1.194.13 255.255.255.252
frame-relay interface-dlci 122
!
interface Serial0/0/1
description 128 kbps leased line to BRO2
bandwidth 128
ip address 10.1.193.5 255.255.255.252
ip flow ingress
encapsulation ppp
clock rate 128000
!
router eigrp 1
network 10.1.192.6 0.0.0.0
network 10.1.192.14 0.0.0.0
network 10.1.193.5 0.0.0.0
network 10.1.194.9 0.0.0.0
network 10.1.194.13 0.0.0.0
network 10.1.220.2 0.0.0.0
no auto-summary
!
ip forward-protocol nd
!
ip flow-export source Loopback0
ip flow-export version 5
ip flow-export destination 10.1.152.1 9996
!
ip http server
no ip http secure-server
!
logging source-interface Loopback0
logging 10.1.152.1
snmp-server community cisco RO
snmp-server community san-fran RW
snmp-server ifindex persist
snmp-server trap-source Loopback0
snmp-server location TSHOOT Lab Facility
snmp-server contact support@mgmt.tshoot.local
snmp-server enable traps snmp authentication linkdown linkup coldstart
warmstart
snmp-server enable traps vrrp
snmp-server enable traps ds1

```

```

snmp-server enable traps tty
snmp-server enable traps eigrp
snmp-server enable traps envmon
snmp-server enable traps flash insertion removal
snmp-server enable traps icsudsu
snmp-server enable traps isdn call-information
snmp-server enable traps isdn layer2
snmp-server enable traps isdn chan-not-avail
snmp-server enable traps isdn ietf
snmp-server enable traps ds0-busyout
snmp-server enable traps ds1-loopback
snmp-server enable traps aaa_server
snmp-server enable traps atm_subif
snmp-server enable traps bgp
snmp-server enable traps bulkstat collection transfer
snmp-server enable traps cnpd
snmp-server enable traps config-copy
snmp-server enable traps config
snmp-server enable traps dsp card-status
snmp-server enable traps entity
snmp-server enable traps event-manager
snmp-server enable traps frame-relay
snmp-server enable traps frame-relay subif
snmp-server enable traps hsrp
snmp-server enable traps ipmobile
snmp-server enable traps ipmulticast
snmp-server enable traps msdp
snmp-server enable traps mvpn
snmp-server enable traps ospf state-change
snmp-server enable traps ospf errors
snmp-server enable traps ospf retransmit
snmp-server enable traps ospf lsa
snmp-server enable traps ospf cisco-specific state-change nssa-trans-change
snmp-server enable traps ospf cisco-specific state-change shamlink interface-
old
snmp-server enable traps ospf cisco-specific state-change shamlink neighbor
snmp-server enable traps ospf cisco-specific errors
snmp-server enable traps ospf cisco-specific retransmit
snmp-server enable traps ospf cisco-specific lsa
snmp-server enable traps pim neighbor-change rp-mapping-change invalid-pim-
message
snmp-server enable traps pppoe
snmp-server enable traps cpu threshold
snmp-server enable traps rsvp
snmp-server enable traps rtr
snmp-server enable traps syslog
snmp-server enable traps l2tun session
snmp-server enable traps vtp
snmp-server enable traps isakmp policy add
snmp-server enable traps isakmp policy delete
snmp-server enable traps isakmp tunnel start
snmp-server enable traps isakmp tunnel stop
snmp-server enable traps ipsec cryptomap add
snmp-server enable traps ipsec cryptomap delete
snmp-server enable traps ipsec cryptomap attach
snmp-server enable traps ipsec cryptomap detach
snmp-server enable traps ipsec tunnel start
snmp-server enable traps ipsec tunnel stop
snmp-server enable traps ipsec too-many-sas
snmp-server host 10.1.152.1 version 2c cisco
!
!
!
!
control-plane
!
!
line con 0

```

```

exec-timeout 60 0
login authentication CONSOLE
line aux 0
line vty 0 4
transport input telnet ssh
line vty 5 15
transport input telnet ssh
!
scheduler allocate 20000 1000
ntp source Loopback0
ntp update-calendar
ntp server 10.1.220.4
ntp server 10.1.220.3
end

```

Switch CSW1

```

!
version 12.2
no service pad
service timestamps debug datetime msec localtime show-timezone
service timestamps log datetime localtime show-timezone
no service password-encryption
!
hostname CSW1
!
boot-start-marker
boot-end-marker
!
logging buffered 16384
logging console warnings
enable secret 5 $1$4mWg$eqWTmSlBAh3uQorxo3HYG/
!
username admin secret 5 $1$0up3$16kazjjd6SgWRRQJqWLSd1
aaa new-model
!
!
aaa authentication login default local
aaa authentication login CONSOLE none
aaa authorization exec default local
!
!
!
aaa session-id common
clock timezone PST -8
clock summer-time PDT recurring
system mtu routing 1500
vtp domain TSHOOT
vtp mode transparent
ip subnet-zero
ip routing
ip domain-name mgmt.tshoot.local
ip name-server 10.1.152.1
no ip dhcp use vrf connected
ip dhcp excluded-address 10.1.128.124 10.1.160.126
ip dhcp excluded-address 10.1.128.188 10.1.160.190
ip dhcp excluded-address 10.1.128.252 10.1.160.254
ip dhcp excluded-address 10.1.129.124 10.1.160.126
ip dhcp excluded-address 10.1.129.188 10.1.160.190
ip dhcp excluded-address 10.1.129.252 10.1.160.254
ip dhcp excluded-address 10.1.130.124 10.1.160.126
ip dhcp excluded-address 10.1.130.188 10.1.160.190
ip dhcp excluded-address 10.1.130.252 10.1.160.254
!

```

```

ip dhcp pool F1S1-OFFICE
  network 10.1.128.64 255.255.255.192
  default-router 10.1.128.126
  dns-server 10.1.152.1
  domain-name tshoot.local
!
ip dhcp pool F1S1-VOICE
  network 10.1.128.128 255.255.255.192
  default-router 10.1.128.190
  dns-server 10.1.152.1
  domain-name tshoot.local
!
ip dhcp pool F1S1-GUEST
  network 10.1.128.192 255.255.255.192
  default-router 10.1.128.254
  dns-server 10.1.152.1
  domain-name tshoot.local
!
ip dhcp pool F1S2-OFFICE
  network 10.1.129.64 255.255.255.192
  default-router 10.1.129.126
  dns-server 10.1.152.1
  domain-name tshoot.local
!
ip dhcp pool F1S2-VOICE
  network 10.1.129.128 255.255.255.192
  default-router 10.1.129.190
  dns-server 10.1.152.1
  domain-name tshoot.local
!
ip dhcp pool F1S2-GUEST
  network 10.1.129.192 255.255.255.192
  default-router 10.1.129.254
  dns-server 10.1.152.1
  domain-name tshoot.local
!
ip dhcp pool F1S3-OFFICE
  network 10.1.130.64 255.255.255.192
  default-router 10.1.130.126
  dns-server 10.1.152.1
  domain-name tshoot.local
!
ip dhcp pool F1S3-VOICE
  network 10.1.130.128 255.255.255.192
  default-router 10.1.130.190
  dns-server 10.1.152.1
  domain-name tshoot.local
!
ip dhcp pool F1S3-GUEST
  network 10.1.130.192 255.255.255.192
  default-router 10.1.130.254
  dns-server 10.1.152.1
  domain-name tshoot.local
!
!
!
!
crypto pki trustpoint TP-self-signed-4156290944
  enrollment selfsigned
  subject-name cn=IOS-Self-Signed-Certificate-4156290944
  revocation-check none
  rsakeypair TP-self-signed-4156290944
!
!
crypto pki certificate chain TP-self-signed-4156290944
  certificate self-signed 01
    3082024E 308201B7 A0030201 02020101 300D0609 2A864886 F70D0101 04050030
    31312F30 2D060355 04031326 494F532D 53656C66 2D536967 6E65642D 43657274

```

```

69666963 6174652D 34313536 32393039 3434301E 170D3933 30333031 30303030
34355A17 0D323030 31303130 30303030 305A3031 312F302D 06035504 03132649
4F532D53 656C662D 5369676E 65642D43 65727469 66696361 74652D34 31353632
39303934 3430819F 300D0609 2A864886 F70D0101 01050003 818D0030 81890281
8100C056 3990CEA3 C012811F 6CEC2151 6E9DDB6E 2024D281 4BE06263 B7BBF2AE
750163FF 3278ED3C 7BF6626F 1680FB2B 1AD380C6 E2CF75C9 0775C9C9 D4EA512F
D610DA5F B843969F 9453E644 51DD84E8 82A0B709 A33297FC 7FF5C8B9 45C885B4
4A9D4B6F F0E6B13E 23E20453 920BEC94 7BAFD193 49013DEC 2979E7B6 05547CEA
1C730203 010001A3 76307430 0F060355 1D130101 FF040530 030101FF 30210603
551D1104 1A301882 16435357 312E6D67 6D742E74 73686F6F 742E6C6F 63616C30
1F060355 1D230418 30168014 C498B4CF A4F171B8 6DCF1A2E 40B4EF38 C87D7F53
301D0603 551D0E04 160414C4 98B4CFA4 F171B86D CF1A2E40 B4EF38C8 7D7F5330
0D06092A 864886F7 0D010104 05000381 8100404B C9C7CFA7 A10120EE 07D933DB
4E3528FA FE93F860 94377B80 E2BFB70B 6B0480D9 52A370A1 84FD6514 B130EA67
9F31D865 5D768472 71EF7B63 D1E03C71 227D3772 C112A272 F16EA386 9357E0E5
8ECCC82D 9BD7898D 66CB9806 C093640D B3A0F498 E83E0B61 04CFB7EB 3E59555C
61B1AFFB 47A9A027 CD705957 57FFB32 F566
quit
!
!
!
errdisable recovery cause bpduguard
!
!
archive
log config
logging enable
logging size 50
notify syslog contenttype plaintext
hidekeys
path tftp://srv1.mgmt.tshoot.local/$h-archive-config
write-memory
file prompt quiet
!
spanning-tree mode rapid-pvst
spanning-tree extend system-id
spanning-tree vlan 1-15,32-47,64-79,96-111,128-129 priority 24576
spanning-tree vlan 16-31,48-63,80-95,112-127 priority 28672
!
vlan internal allocation policy ascending
!
vlan 11
name ISP1
!
vlan 17
name F1S1-OFFICE
!
vlan 18
name F1S1-VOICE
!
vlan 19
name F1S1-GUEST
!
vlan 21
name F1S2-OFFICE
!
vlan 22
name F1S2-VOICE
!
vlan 23
name F1S2-GUEST
!
vlan 25
name F1S3-OFFICE
!
vlan 26
name F1S3-VOICE
!

```

```

vlan 27
  name F1S3-GUEST
!
vlan 33
  name F2S1-OFFICE
!
vlan 34
  name F2S1-VOICE
!
vlan 35
  name F2S1-GUEST
!
vlan 37
  name F2S2-OFFICE
!
vlan 38
  name F2S2-VOICE
!
vlan 39
  name F2S2-GUEST
!
vlan 41
  name F2S3-OFFICE
!
vlan 42
  name F2S3-VOICE
!
vlan 43
  name F2S3-GUEST
!
vlan 112
  name INT-SERVER
!
vlan 128
  name MGMT
!
vlan 129
  name TRANSIT
!
vlan 1000
  name NATIVE
!
vlan 1001
  name UNUSED
!
ip telnet source-interface Vlan128
ip ssh source-interface Vlan128
!
!
!
interface Port-channel1
  description Channel to ASW1
  switchport trunk encapsulation dot1q
  switchport trunk native vlan 1000
  switchport trunk allowed vlan 17-19,128
  switchport mode trunk
  switchport nonegotiate
!
interface Port-channel10
  description Channel to CSW2
  switchport trunk encapsulation dot1q
  switchport trunk native vlan 1000
  switchport trunk allowed vlan 17-19,21-23,25-27,33-35,37-39,41-43,112,128,129
  switchport mode trunk
  switchport nonegotiate
!
interface FastEthernet0/1
  description Channel to ASW1

```

```

switchport trunk encapsulation dot1q
switchport trunk native vlan 1000
switchport trunk allowed vlan 17-19,128
switchport mode trunk
switchport nonegotiate
channel-group 1 mode on
!
interface FastEthernet0/2
description Channel to ASW1
switchport trunk encapsulation dot1q
switchport trunk native vlan 1000
switchport trunk allowed vlan 17-19,128
switchport mode trunk
switchport nonegotiate
channel-group 1 mode on
!
interface FastEthernet0/3
description Channel to CSW2
switchport trunk encapsulation dot1q
switchport trunk native vlan 1000
switchport trunk allowed vlan 17-19,21-23,25-27,33-35,37-39,41-43,112,128,129
switchport mode trunk
switchport nonegotiate
channel-group 10 mode on
!
interface FastEthernet0/4
description Channel to CSW2
switchport trunk encapsulation dot1q
switchport trunk native vlan 1000
switchport trunk allowed vlan 17-19,21-23,25-27,33-35,37-39,41-43,112,128,129
switchport mode trunk
switchport nonegotiate
channel-group 10 mode on
!
interface FastEthernet0/5
description Unused
switchport access vlan 1001
switchport mode access
switchport nonegotiate
shutdown
!
interface FastEthernet0/6
description Unused
switchport access vlan 1001
switchport mode access
switchport nonegotiate
shutdown
!
interface FastEthernet0/7
description Unused
switchport access vlan 1001
switchport mode access
switchport nonegotiate
shutdown
spanning-tree portfast
!
interface FastEthernet0/8
description Unused
switchport access vlan 1001
switchport mode access
switchport nonegotiate
shutdown
!
interface FastEthernet0/9
description Unused
switchport access vlan 1001
switchport mode access
switchport nonegotiate

```

```

        shutdown
    !
interface FastEthernet0/10
    description FE to SRV1
    switchport access vlan 112
    switchport mode access
    switchport nonegotiate
    spanning-tree portfast
    !
interface FastEthernet0/11
    description FE to CRO1
    no switchport
    ip address 10.1.192.1 255.255.255.252
    speed 100
    duplex full
    !
interface FastEthernet0/12
    description FE to CRO2
    no switchport
    ip address 10.1.192.5 255.255.255.252
    speed 100
    duplex full
    !
interface FastEthernet0/13
    description Unused
    switchport access vlan 1001
    switchport mode access
    switchport nonegotiate
    shutdown
    !
interface FastEthernet0/14
    description Unused
    switchport access vlan 1001
    switchport mode access
    switchport nonegotiate
    shutdown
    !
interface FastEthernet0/15
    description Vlan ISP1 to IRO1
    switchport access vlan 11
    switchport mode access
    switchport nonegotiate
    speed 100
    duplex full
    spanning-tree portfast
    !
interface FastEthernet0/16
    description Vlan TRANSIT to IRO2
    switchport access vlan 129
    switchport mode access
    switchport nonegotiate
    speed 100
    duplex full
    spanning-tree portfast
    !
interface FastEthernet0/17
    description Unused
    switchport access vlan 1001
    switchport mode access
    switchport nonegotiate
    shutdown
    !
interface FastEthernet0/18
    description Unused
    switchport access vlan 1001
    switchport mode access
    switchport nonegotiate
    shutdown

```



```

!
interface FastEthernet0/19
description Unused
switchport access vlan 1001
switchport mode access
switchport nonegotiate
shutdown
!
interface FastEthernet0/20
description Unused
switchport access vlan 1001
switchport mode access
switchport nonegotiate
shutdown
!
interface FastEthernet0/21
description Unused
switchport access vlan 1001
switchport mode access
switchport nonegotiate
shutdown
!
interface FastEthernet0/22
description Unused
switchport access vlan 1001
switchport mode access
switchport nonegotiate
shutdown
!
interface FastEthernet0/23
description Unused
switchport access vlan 1001
switchport mode access
switchport nonegotiate
shutdown
!
interface FastEthernet0/24
description Metro FE to ISP1
switchport access vlan 11
switchport mode access
switchport nonegotiate
no cdp enable
spanning-tree portfast
spanning-tree bpduguard enable
!
interface GigabitEthernet0/1
!
interface GigabitEthernet0/2
!
interface Vlan1
no ip address
shutdown
!
interface Vlan17
ip address 10.1.128.124 255.255.255.192
standby 17 ip 10.1.128.126
standby 17 preempt
!
interface Vlan18
ip address 10.1.128.188 255.255.255.192
standby 18 ip 10.1.128.190
standby 18 preempt
!
interface Vlan19
ip address 10.1.128.252 255.255.255.192
standby 19 ip 10.1.128.254
standby 19 preempt
!

```

```

interface Vlan21
 ip address 10.1.129.124 255.255.255.192
 standby 21 ip 10.1.129.126
 standby 21 preempt
!
interface Vlan22
 ip address 10.1.129.188 255.255.255.192
 standby 22 ip 10.1.129.190
 standby 22 preempt
!
interface Vlan23
 ip address 10.1.129.252 255.255.255.192
 standby 23 ip 10.1.129.254
 standby 23 preempt
!
interface Vlan25
 ip address 10.1.130.124 255.255.255.192
 standby 25 ip 10.1.130.126
 standby 25 preempt
!
interface Vlan26
 ip address 10.1.130.188 255.255.255.192
 standby 26 ip 10.1.130.190
 standby 26 preempt
!
interface Vlan27
 ip address 10.1.130.252 255.255.255.192
 standby 27 ip 10.1.130.254
 standby 27 preempt
!
interface Vlan33
 ip address 10.1.132.124 255.255.255.192
 standby 33 ip 10.1.132.126
 standby 33 priority 110
 standby 33 preempt
!
interface Vlan34
 ip address 10.1.132.188 255.255.255.192
 standby 34 ip 10.1.132.190
 standby 34 priority 110
 standby 34 preempt
!
interface Vlan35
 ip address 10.1.132.252 255.255.255.192
 standby 35 ip 10.1.132.254
 standby 35 priority 110
 standby 35 preempt
!
interface Vlan37
 ip address 10.1.133.124 255.255.255.192
 standby 37 ip 10.1.133.126
 standby 37 priority 110
 standby 37 preempt
!
interface Vlan38
 ip address 10.1.133.188 255.255.255.192
 standby 38 ip 10.1.133.190
 standby 38 priority 110
 standby 38 preempt
!
interface Vlan39
 ip address 10.1.133.252 255.255.255.192
 standby 39 ip 10.1.133.254
 standby 39 priority 110
 standby 39 preempt
!
interface Vlan41
 ip address 10.1.134.124 255.255.255.192

```

```

standby 41 ip 10.1.134.126
standby 41 priority 110
standby 41 preempt
!
interface Vlan42
ip address 10.1.134.188 255.255.255.192
standby 42 ip 10.1.134.190
standby 42 priority 110
standby 42 preempt
!
interface Vlan43
ip address 10.1.134.252 255.255.255.192
standby 43 ip 10.1.134.254
standby 43 priority 110
standby 43 preempt
!
interface Vlan112
ip address 10.1.152.252 255.255.255.0
standby 112 ip 10.1.152.254
standby 112 preempt
!
interface Vlan128
ip address 10.1.159.252 255.255.252.0
standby 128 ip 10.1.159.254
standby 128 priority 110
standby 128 preempt
!
interface Vlan129
ip address 10.1.192.17 255.255.255.248
!
!
router eigrp 1
passive-interface default
no passive-interface Vlan129
no passive-interface FastEthernet0/11
no passive-interface FastEthernet0/12
no auto-summary
network 10.1.0.0 0.0.255.255
!
ip classless
ip http server
ip http secure-server
!
!
logging source-interface Vlan128
logging 10.1.152.1
snmp-server community cisco RO
snmp-server community san-fran RW
snmp-server trap-source Vlan128
snmp-server location TSHOOT Lab Facility
snmp-server contact support@mgmt.tshoot.local
snmp-server enable traps snmp authentication linkdown linkup coldstart
warmstart
snmp-server enable traps transceiver all
snmp-server enable traps tty
snmp-server enable traps eigrp
snmp-server enable traps ospf state-change
snmp-server enable traps ospf errors
snmp-server enable traps ospf retransmit
snmp-server enable traps ospf lsa
snmp-server enable traps ospf cisco-specific state-change nssa-trans-change
snmp-server enable traps ospf cisco-specific state-change shamlink interface-
old
snmp-server enable traps ospf cisco-specific state-change shamlink neighbor
snmp-server enable traps ospf cisco-specific errors
snmp-server enable traps ospf cisco-specific retransmit
snmp-server enable traps ospf cisco-specific lsa
snmp-server enable traps cluster

```

```

snmp-server enable traps entity
snmp-server enable traps cpu threshold
snmp-server enable traps power-ethernet group 1
snmp-server enable traps power-ethernet police
snmp-server enable traps rep
snmp-server enable traps vtp
snmp-server enable traps vlancreate
snmp-server enable traps vlandelete
snmp-server enable traps flash insertion removal
snmp-server enable traps port-security
snmp-server enable traps dot1x auth-fail-vlan guest-vlan no-auth-fail-vlan no-
guest-vlan
snmp-server enable traps envmon fan shutdown supply temperature status
snmp-server enable traps bgp
snmp-server enable traps cef resource-failure peer-state-change peer-fib-
state-change inconsistency
snmp-server enable traps config-copy
snmp-server enable traps config
snmp-server enable traps config-ctid
snmp-server enable traps event-manager
snmp-server enable traps hsrp
snmp-server enable traps ipmulticast
snmp-server enable traps msdp
snmp-server enable traps pim neighbor-change rp-mapping-change invalid-pim-
message
snmp-server enable traps bridge newroot topologychange
snmp-server enable traps stpx inconsistency root-inconsistency loop-
inconsistency
snmp-server enable traps syslog
snmp-server enable traps rtr
snmp-server enable traps mvpn
snmp-server enable traps mac-notification change move threshold
snmp-server enable traps vlan-membership
snmp-server enable traps errdisable
snmp-server host 10.1.152.1 version 2c cisco
snmp ifmib ifindex persist
!
control-plane
!
line con 0
  exec-timeout 60 0
  login authentication CONSOLE
line vty 0 4
  transport input telnet ssh
line vty 5 15
  transport input telnet ssh
!
ntp source Vlan128
ntp server 10.1.220.4
ntp server 10.1.220.3
end

```

Switch CSW2

```

!
version 12.2
no service pad
service timestamps debug datetime msec localtime show-timezone
service timestamps log datetime localtime show-timezone
no service password-encryption
!
hostname CSW2
!
boot-start-marker
boot-end-marker
!
logging buffered 16384
logging console warnings

```

```

enable secret 5 $1$4mWg$eqWTmSlBAh3uQorxo3HYG/
!
username admin secret 5 $1$0up3$16kazjld6SgWRRQJqWLSd1
aaa new-model
!
!
aaa authentication login default local
aaa authentication login CONSOLE none
aaa authorization exec default local
!
!
!
aaa session-id common
clock timezone PST -8
clock summer-time PDT recurring
system mtu routing 1500
vtp domain TSHOOT
vtp mode transparent
ip subnet-zero
ip routing
ip domain-name mgmt.tshoot.local
ip name-server 10.1.152.1
no ip dhcp use vrf connected
ip dhcp excluded-address 10.1.132.124 10.1.160.126
ip dhcp excluded-address 10.1.132.188 10.1.160.190
ip dhcp excluded-address 10.1.132.252 10.1.160.254
ip dhcp excluded-address 10.1.133.124 10.1.160.126
ip dhcp excluded-address 10.1.133.188 10.1.160.190
ip dhcp excluded-address 10.1.133.252 10.1.160.254
ip dhcp excluded-address 10.1.134.124 10.1.160.126
ip dhcp excluded-address 10.1.134.188 10.1.160.190
ip dhcp excluded-address 10.1.134.252 10.1.160.254
!
ip dhcp pool F2S1-OFFICE
    network 10.1.132.64 255.255.255.192
    default-router 10.1.132.126
    dns-server 10.1.152.1
    domain-name tshoot.local
!
ip dhcp pool F2S1-VOICE
    network 10.1.132.128 255.255.255.192
    default-router 10.1.132.190
    dns-server 10.1.152.1
    domain-name tshoot.local
!
ip dhcp pool F2S1-GUEST
    network 10.1.132.192 255.255.255.192
    default-router 10.1.132.254
    dns-server 10.1.152.1
    domain-name tshoot.local
!
ip dhcp pool F2S2-OFFICE
    network 10.1.133.64 255.255.255.192
    default-router 10.1.133.126
    dns-server 10.1.152.1
    domain-name tshoot.local
!
ip dhcp pool F2S2-VOICE
    network 10.1.133.128 255.255.255.192
    default-router 10.1.133.190
    dns-server 10.1.152.1
    domain-name tshoot.local
!
ip dhcp pool F2S2-GUEST
    network 10.1.133.192 255.255.255.192
    default-router 10.1.133.254
    dns-server 10.1.152.1
    domain-name tshoot.local

```

```

!
ip dhcp pool F2S3-OFFICE
 network 10.1.134.64 255.255.255.192
 default-router 10.1.134.126
 dns-server 10.1.152.1
 domain-name tshoot.local
!
ip dhcp pool F2S3-VOICE
 network 10.1.134.128 255.255.255.192
 default-router 10.1.134.190
 dns-server 10.1.152.1
 domain-name tshoot.local
!
ip dhcp pool F2S3-GUEST
 network 10.1.134.192 255.255.255.192
 default-router 10.1.134.254
 dns-server 10.1.152.1
 domain-name tshoot.local
!
!
!
!
crypto pki trustpoint TP-self-signed-656507904
 enrollment selfsigned
 subject-name cn=IOS-Self-Signed-Certificate-656507904
 revocation-check none
 rsakeypair TP-self-signed-656507904
!
!
crypto pki certificate chain TP-self-signed-656507904
 certificate self-signed 01
  3082024C 308201B5 A0030201 02020101 300D0609 2A864886 F70D0101 04050030
  30312E30 2C060355 04031325 494F532D 53656C66 2D536967 6E65642D 43657274
  69666963 6174652D 36353635 30373930 34301E17 0D393330 33303130 30303034
  355A170D 32303031 30313030 30303030 5A303031 2E302C06 03550403 1325494F
  532D5365 6C662D53 69676E65 642D4365 72746966 69636174 652D3635 36353037
  39303430 819F300D 06092A86 4886F70D 01010105 0003818D 00308189 02818100
  C80D243D E03F0AFE 741422DE 2918138D 8D992578 22BA27C7 DAD315CB 79848E2A
  8B5A6804 6FBEE422 9D2E91D4 CE087496 CD87F485 88B9132E 3DF86DC5 095CCA11
  3C991558 DB9AD6B9 7934600F 4BC124CC AE1E3CC1 19BB3A0A 3DEA534D 0237C1A3
  04719956 DB240A6B 8B5E3994 E1A6D097 7C54EDC2 5A82DBD0 41CD943A 7280A279
  02030100 01A37630 74300F06 03551D13 0101FF04 05300301 01FF3021 0603551D
  11041A30 18821643 5357322E 6D676D74 2E747368 6F6F742E 6C6F6361 6C301F06
  03551D23 04183016 8014E098 76669115 647F8B72 EE9B0447 B75F049D 1DBE301D
  0603551D 0E041604 14E09876 66911564 7F8B72EE 9B0447B7 5F049D1D BE300D06
  092A8648 86F70D01 01040500 03818100 8DF2B893 33B3BDFF 6E1A8FA9 7360851E
  9AD5C7E8 5158EEDF B108792E 53391743 6FCF9652 43309673 38793C55 68761E66
  48ED38FB FDC8117D 5247EC3D 8A20F14D 300A3189 AC02E7FA C1A0B271 B511E2ED
  E8709E96 EFF9962C 1D38F392 B7C28590 4CE26038 63AA0384 C2D05915 C112CDD7
  39C15F4B 6FE95A29 193E4415 40713238
 quit
!
!
!
errdisable recovery cause bpduguard
!
!
archive
 log config
  logging enable
  logging size 50
  notify syslog contenttype plaintext
  hidekeys
 path tftp://srv1.mgmt.tshoot.local/$h-archive-config
 write-memory
 file prompt quiet
!
spanning-tree mode rapid-pvst

```

```

spanning-tree extend system-id
spanning-tree vlan 1-15,32-47,64-79,96-111,128-129 priority 28672
spanning-tree vlan 16-31,48-63,80-95,112-127 priority 24576
!
vlan internal allocation policy ascending
!
vlan 12
    name ISP2
!
vlan 17
    name F1S1-OFFICE
!
vlan 18
    name F1S1-VOICE
!
vlan 19
    name F1S1-GUEST
!
vlan 21
    name F1S2-OFFICE
!
vlan 22
    name F1S2-VOICE
!
vlan 23
    name F1S2-GUEST
!
vlan 25
    name F1S3-OFFICE
!
vlan 26
    name F1S3-VOICE
!
vlan 27
    name F1S3-GUEST
!
vlan 33
    name F2S1-OFFICE
!
vlan 34
    name F2S1-VOICE
!
vlan 35
    name F2S1-GUEST
!
vlan 37
    name F2S2-OFFICE
!
vlan 38
    name F2S2-VOICE
!
vlan 39
    name F2S2-GUEST
!
vlan 41
    name F2S3-OFFICE
!
vlan 42
    name F2S3-VOICE
!
vlan 43
    name F2S3-GUEST
!
vlan 112
    name INT-SERVER
!
vlan 128
    name MGMT

```

```

!
vlan 129
  name TRANSIT
!
vlan 1000
  name NATIVE
!
vlan 1001
  name UNUSED
!
ip telnet source-interface Vlan128
ip ssh source-interface Vlan128
!
!
!
interface Port-channel1
  description Channel to ASW1
  switchport trunk encapsulation dot1q
  switchport trunk native vlan 1000
  switchport trunk allowed vlan 17-19,128
  switchport mode trunk
  switchport nonegotiate
!
interface Port-channel10
  description Channel to CSW1
  switchport trunk encapsulation dot1q
  switchport trunk native vlan 1000
  switchport trunk allowed vlan 17-19,21-23,25-27,33-35,37-39,41-43,112,128,129
  switchport mode trunk
  switchport nonegotiate
!
interface FastEthernet0/1
  description Unused
  switchport access vlan 1001
  switchport mode access
  switchport nonegotiate
  shutdown
!
interface FastEthernet0/2
  description Unused
  switchport access vlan 1001
  switchport mode access
  switchport nonegotiate
  shutdown
!
interface FastEthernet0/3
  description Channel to CSW1
  switchport trunk encapsulation dot1q
  switchport trunk native vlan 1000
  switchport trunk allowed vlan 17-19,21-23,25-27,33-35,37-39,41-43,112,128,129
  switchport mode trunk
  switchport nonegotiate
  channel-group 10 mode on
!
interface FastEthernet0/4
  description Channel to CSW1
  switchport trunk encapsulation dot1q
  switchport trunk native vlan 1000
  switchport trunk allowed vlan 17-19,21-23,25-27,33-35,37-39,41-43,112,128,129
  switchport mode trunk
  switchport nonegotiate
  channel-group 10 mode on
!
interface FastEthernet0/5
  description Channel to ASW1
  switchport trunk encapsulation dot1q
  switchport trunk native vlan 1000
  switchport trunk allowed vlan 17-19,128

```



```

switchport mode trunk
switchport nonegotiate
channel-group 1 mode on
!
interface FastEthernet0/6
description Channel to ASW1
switchport trunk encapsulation dot1q
switchport trunk native vlan 1000
switchport trunk allowed vlan 17-19,128
switchport mode trunk
switchport nonegotiate
channel-group 1 mode on
!
interface FastEthernet0/7
description Unused
switchport access vlan 1001
switchport mode access
switchport nonegotiate
shutdown
!
interface FastEthernet0/8
description Unused
switchport access vlan 1001
switchport mode access
switchport nonegotiate
shutdown
!
interface FastEthernet0/9
description Unused
switchport access vlan 1001
switchport mode access
switchport nonegotiate
shutdown
!
interface FastEthernet0/10
description Unused
switchport access vlan 1001
switchport mode access
switchport nonegotiate
shutdown
!
interface FastEthernet0/11
description FE to CRO1
no switchport
ip address 10.1.192.9 255.255.255.252
speed 100
duplex full
!
interface FastEthernet0/12
description FE to CRO2
no switchport
ip address 10.1.192.13 255.255.255.252
speed 100
duplex full
!
interface FastEthernet0/13
description Unused
switchport access vlan 1001
switchport mode access
switchport nonegotiate
shutdown
!
interface FastEthernet0/14
description Unused
switchport access vlan 1001
switchport mode access
switchport nonegotiate
shutdown

```

```

!
interface FastEthernet0/15
  description Vlan TRANSIT to IRO1
  switchport access vlan 129
  switchport mode access
  switchport nonegotiate
  speed 100
  duplex full
  spanning-tree portfast
!
interface FastEthernet0/16
  description Vlan ISP2 to IRO2
  switchport access vlan 12
  switchport mode access
  switchport nonegotiate
  speed 100
  duplex full
  spanning-tree portfast
!
interface FastEthernet0/17
  description Unused
  switchport access vlan 1001
  switchport mode access
  switchport nonegotiate
  shutdown
!
interface FastEthernet0/18
  description Unused
  switchport access vlan 1001
  switchport mode access
  switchport nonegotiate
  shutdown
!
interface FastEthernet0/19
  description Unused
  switchport access vlan 1001
  switchport mode access
  switchport nonegotiate
  shutdown
!
interface FastEthernet0/20
  description Unused
  switchport access vlan 1001
  switchport mode access
  switchport nonegotiate
  shutdown
!
interface FastEthernet0/21
  description Unused
  switchport access vlan 1001
  switchport mode access
  switchport nonegotiate
  shutdown
!
interface FastEthernet0/22
  description Unused
  switchport access vlan 1001
  switchport mode access
  switchport nonegotiate
  shutdown
!
interface FastEthernet0/23
  description Unused
  switchport access vlan 1001
  switchport mode access
  switchport nonegotiate
  shutdown
!

```

```

interface FastEthernet0/24
description Metro FE to ISP2
switchport access vlan 12
switchport mode access
switchport nonegotiate
no cdp enable
spanning-tree portfast
spanning-tree bpduguard enable
!
interface GigabitEthernet0/1
!
interface GigabitEthernet0/2
!
interface Vlan1
no ip address
!
interface Vlan17
ip address 10.1.128.125 255.255.255.192
standby 17 ip 10.1.128.126
standby 17 priority 110
standby 17 preempt
!
interface Vlan18
ip address 10.1.128.189 255.255.255.192
standby 18 ip 10.1.128.190
standby 18 priority 110
standby 18 preempt
!
interface Vlan19
ip address 10.1.128.253 255.255.255.192
standby 19 ip 10.1.128.254
standby 19 priority 110
standby 19 preempt
!
interface Vlan21
ip address 10.1.129.125 255.255.255.192
standby 21 ip 10.1.129.126
standby 21 priority 110
standby 21 preempt
!
interface Vlan22
ip address 10.1.129.189 255.255.255.192
standby 22 ip 10.1.129.190
standby 22 priority 110
standby 22 preempt
!
interface Vlan23
ip address 10.1.129.253 255.255.255.192
standby 23 ip 10.1.129.254
standby 23 priority 110
standby 23 preempt
!
interface Vlan25
ip address 10.1.130.125 255.255.255.192
standby 25 ip 10.1.130.126
standby 25 priority 110
standby 25 preempt
!
interface Vlan26
ip address 10.1.130.189 255.255.255.192
standby 26 ip 10.1.130.190
standby 26 priority 110
standby 26 preempt
!
interface Vlan27
ip address 10.1.130.253 255.255.255.192
standby 27 ip 10.1.130.254
standby 27 priority 110

```

```

    standby 27 preempt
!
interface Vlan33
 ip address 10.1.132.125 255.255.255.192
 standby 33 ip 10.1.132.126
 standby 33 preempt
!
interface Vlan34
 ip address 10.1.132.189 255.255.255.192
 standby 34 ip 10.1.132.190
 standby 34 preempt
!
interface Vlan35
 ip address 10.1.132.253 255.255.255.192
 standby 35 ip 10.1.132.254
 standby 35 preempt
!
interface Vlan37
 ip address 10.1.133.125 255.255.255.192
 standby 37 ip 10.1.133.126
 standby 37 preempt
!
interface Vlan38
 ip address 10.1.133.189 255.255.255.192
 standby 38 ip 10.1.133.190
 standby 38 preempt
!
interface Vlan39
 ip address 10.1.133.253 255.255.255.192
 standby 39 ip 10.1.133.254
 standby 39 preempt
!
interface Vlan41
 ip address 10.1.134.125 255.255.255.192
 standby 41 ip 10.1.134.126
 standby 41 preempt
!
interface Vlan42
 ip address 10.1.134.189 255.255.255.192
 standby 42 ip 10.1.134.190
 standby 42 preempt
!
interface Vlan43
 ip address 10.1.134.253 255.255.255.192
 standby 43 ip 10.1.134.254
 standby 43 preempt
!
interface Vlan112
 ip address 10.1.152.253 255.255.255.0
 standby 112 ip 10.1.152.254
 standby 112 priority 110
 standby 112 preempt
!
interface Vlan128
 ip address 10.1.159.253 255.255.252.0
 standby 128 ip 10.1.159.254
 standby 128 preempt
!
interface Vlan129
 ip address 10.1.192.18 255.255.255.248
!
!
router eigrp 1
 passive-interface default
 no passive-interface Vlan129
 no passive-interface FastEthernet0/11
 no passive-interface FastEthernet0/12
 no auto-summary

```

```

network 10.1.0.0 0.0.255.255
!
ip classless
ip http server
ip http secure-server
!
!
logging source-interface Vlan128
logging 10.1.152.1
snmp-server community cisco RO
snmp-server community san-fran RW
snmp-server trap-source Vlan128
snmp-server location TSHOOT Lab Facility
snmp-server contact support@mgmt.tshoot.local
snmp-server enable traps snmp authentication linkdown linkup coldstart
warmstart
snmp-server enable traps transceiver all
snmp-server enable traps tty
snmp-server enable traps eigrp
snmp-server enable traps ospf state-change
snmp-server enable traps ospf errors
snmp-server enable traps ospf retransmit
snmp-server enable traps ospf lsa
snmp-server enable traps ospf cisco-specific state-change nssa-trans-change
snmp-server enable traps ospf cisco-specific state-change shamlink interface-
old
snmp-server enable traps ospf cisco-specific state-change shamlink neighbor
snmp-server enable traps ospf cisco-specific errors
snmp-server enable traps ospf cisco-specific retransmit
snmp-server enable traps ospf cisco-specific lsa
snmp-server enable traps cluster
snmp-server enable traps entity
snmp-server enable traps cpu threshold
snmp-server enable traps power-ethernet group 1
snmp-server enable traps power-ethernet police
snmp-server enable traps rep
snmp-server enable traps vtp
snmp-server enable traps vlancreate
snmp-server enable traps vlandelete
snmp-server enable traps flash insertion removal
snmp-server enable traps port-security
snmp-server enable traps dot1x auth-fail-vlan guest-vlan no-auth-fail-vlan no-
guest-vlan
snmp-server enable traps envmon fan shutdown supply temperature status
snmp-server enable traps bgp
snmp-server enable traps cef resource-failure peer-state-change peer-fib-
state-change inconsistency
snmp-server enable traps config-copy
snmp-server enable traps config
snmp-server enable traps config-ctid
snmp-server enable traps event-manager
snmp-server enable traps hsrp
snmp-server enable traps ipmulticast
snmp-server enable traps msdp
snmp-server enable traps pim neighbor-change rp-mapping-change invalid-pim-
message
snmp-server enable traps bridge newroot topologychange
snmp-server enable traps stpx inconsistency root-inconsistency loop-
inconsistency
snmp-server enable traps syslog
snmp-server enable traps rtr
snmp-server enable traps mvpn
snmp-server enable traps mac-notification change move threshold
snmp-server enable traps vlan-membership
snmp-server enable traps errdisable
snmp-server host 10.1.152.1 version 2c cisco
snmp ifmib ifindex persist
!

```

```

control-plane
!
line con 0
  exec-timeout 60 0
  login authentication CONSOLE
line vty 0 4
  transport input telnet ssh
line vty 5 15
  transport input telnet ssh
!
ntp source Vlan128
ntp server 10.1.220.4
ntp server 10.1.220.3
end

```

Router IRO1

```

!
version 12.4
service timestamps debug datetime msec localtime show-timezone
service timestamps log datetime localtime show-timezone
no service password-encryption
!
hostname IRO1
!
boot-start-marker
boot-end-marker
!
logging buffered 16384 debugging
logging console warnings
enable secret 5 $1$4mWg$eqWTmSlBAh3uQorxo3HYG/
!
aaa new-model
!
!
aaa authentication login default local
aaa authentication login CONSOLE none
aaa authorization exec default local
!
aaa session-id common
clock timezone PST -8
clock summer-time PDT recurring
ip cef
!
!
!
!
ip domain name mgmt.tshoot.local
ip name-server 10.1.152.1
ip auth-proxy max-nodata-conns 3
ip admission max-nodata-conns 3
ip sla monitor 11
  type echo protocol ipIcmpEcho 192.168.224.254 source-interface
FastEthernet0/0
  frequency 10
ip sla monitor schedule 11 life forever start-time now
!
!
!
file prompt quiet
username admin secret 5 $1$0up3$16kazjdd6SgWRRQJqWLSd1
archive
  log config
  logging enable
  logging size 50
  notify syslog
  hidekeys
  path tftp://srv1.mgmt.tshoot.local/$h-archive-config

```

```

write-memory
!
!
ip telnet source-interface Loopback0
ip ssh source-interface Loopback0
!
track 1 rtr 11
!
!
!
!
interface Loopback0
 ip address 10.1.220.3 255.255.255.255
 ip nat inside
 ip virtual-reassembly
!
interface FastEthernet0/0
 description FE to ISP1 via CSW1
 ip address 192.168.224.241 255.255.255.240
 ip flow ingress
 ip nat outside
 ip virtual-reassembly
 speed 100
 full-duplex
!
interface FastEthernet0/1
 description Vlan TRANSIT
 ip address 10.1.192.19 255.255.255.248
 ip flow ingress
 ip nat inside
 ip virtual-reassembly
 speed 100
 full-duplex
!
interface Serial0/0/0
 no ip address
 shutdown
 no fair-queue
!
router eigrp 1
 redistribute static metric 10000 100 255 1 1500
 passive-interface default
 no passive-interface FastEthernet0/0
 no passive-interface FastEthernet0/1
 network 10.1.192.19 0.0.0.0
 network 10.1.220.3 0.0.0.0
 network 192.168.224.241 0.0.0.0
 no auto-summary
!
ip forward-protocol nd
ip route 0.0.0.0 0.0.0.0 192.168.224.254 track 1
!
ip flow-export source Loopback0
ip flow-export version 5
ip flow-export destination 10.1.152.1 9996
!
ip http server
no ip http secure-server
ip nat inside source list INTERNAL-NETWORKS interface FastEthernet0/0 overload
ip nat inside source static 10.1.152.1 192.168.224.242
!
ip access-list standard INTERNAL-NETWORKS
 permit 10.1.128.0 0.0.127.255
!
logging source-interface Loopback0
logging 10.1.152.1
snmp-server community cisco RO
snmp-server community san-fran RW

```

```

snmp-server ifindex persist
snmp-server trap-source Loopback0
snmp-server location TSHOOT Lab Facility
snmp-server contact support@mgmt.tshoot.local
snmp-server enable traps snmp authentication linkdown linkup coldstart
warmstart
snmp-server enable traps vrrp
snmp-server enable traps ds1
snmp-server enable traps tty
snmp-server enable traps eigrp
snmp-server enable traps envmon
snmp-server enable traps flash insertion removal
snmp-server enable traps icsudsu
snmp-server enable traps isdn call-information
snmp-server enable traps isdn layer2
snmp-server enable traps isdn chan-not-avail
snmp-server enable traps isdn ietf
snmp-server enable traps ds0-busyout
snmp-server enable traps ds1-loopback
snmp-server enable traps aaa_server
snmp-server enable traps atm subif
snmp-server enable traps bgp
snmp-server enable traps bulkstat collection transfer
snmp-server enable traps cnpd
snmp-server enable traps config-copy
snmp-server enable traps config
snmp-server enable traps dsp card-status
snmp-server enable traps entity
snmp-server enable traps event-manager
snmp-server enable traps frame-relay
snmp-server enable traps frame-relay subif
snmp-server enable traps hsrp
snmp-server enable traps ipmobile
snmp-server enable traps ipmulticast
snmp-server enable traps msdp
snmp-server enable traps mvpn
snmp-server enable traps ospf state-change
snmp-server enable traps ospf errors
snmp-server enable traps ospf retransmit
snmp-server enable traps ospf lsa
snmp-server enable traps ospf cisco-specific state-change nssa-trans-change
snmp-server enable traps ospf cisco-specific state-change shamlink interface-
old
snmp-server enable traps ospf cisco-specific state-change shamlink neighbor
snmp-server enable traps ospf cisco-specific errors
snmp-server enable traps ospf cisco-specific retransmit
snmp-server enable traps ospf cisco-specific lsa
snmp-server enable traps pim neighbor-change rp-mapping-change invalid-pim-
message
snmp-server enable traps pppoe
snmp-server enable traps cpu threshold
snmp-server enable traps rsvp
snmp-server enable traps rtr
snmp-server enable traps syslog
snmp-server enable traps l2tun session
snmp-server enable traps vtp
snmp-server enable traps isakmp policy add
snmp-server enable traps isakmp policy delete
snmp-server enable traps isakmp tunnel start
snmp-server enable traps isakmp tunnel stop
snmp-server enable traps ipsec cryptomap add
snmp-server enable traps ipsec cryptomap delete
snmp-server enable traps ipsec cryptomap attach
snmp-server enable traps ipsec cryptomap detach
snmp-server enable traps ipsec tunnel start
snmp-server enable traps ipsec tunnel stop
snmp-server enable traps ipsec too-many-sas
snmp-server host 10.1.152.1 version 2c cisco

```



```

!
!
!
!
control-plane
!
!
line con 0
  exec-timeout 60 0
  login authentication CONSOLE
line aux 0
line vty 0 4
  transport input telnet ssh
line vty 5 15
  transport input telnet ssh
!
scheduler allocate 20000 1000
ntp update-calendar
ntp server 192.168.224.1
ntp server 172.24.240.1
end

```

Router IRO2

```

!
version 12.4
service timestamps debug datetime msec localtime show-timezone
service timestamps log datetime localtime show-timezone
no service password-encryption
!
hostname IRO2
!
boot-start-marker
boot-end-marker
!
logging buffered 16384 debugging
logging console warnings
enable secret 5 $1$4mWg$eqWTmSlBAh3uQorxo3HYG/
!
aaa new-model
!
!
aaa authentication login default local
aaa authentication login CONSOLE none
aaa authorization exec default local
!
aaa session-id common
clock timezone PST -8
clock summer-time PDT recurring
ip cef
!
!
!
!
ip domain name mgmt.tshoot.local
ip name-server 10.1.152.1
ip auth-proxy max-nodata-conns 3
ip admission max-nodata-conns 3
ip sla monitor 12
  type echo protocol ipIcmpEcho 172.24.244.86 source-interface FastEthernet0/1
  frequency 10
ip sla monitor schedule 12 life forever start-time now
!
!
!
file prompt quiet
username admin secret 5 $1$0up3$16kazjjd6SgWRRQJqWLSd1
archive

```

```

log config
  logging enable
  logging size 50
  notify syslog
  hidekeys
path tftp://srv1.mgmt.tshoot.local/$h-archive-config
write-memory
!
!
ip telnet source-interface Loopback0
ip ssh source-interface Loopback0
!
track 1 rtr 12
!
!
!
!
interface Loopback0
  ip address 10.1.220.4 255.255.255.255
!
interface FastEthernet0/0
  description Vlan TRANSIT
  ip address 10.1.192.20 255.255.255.248
  ip flow ingress
  ip nat inside
  ip virtual-reassembly
  speed 100
  full-duplex
!
interface FastEthernet0/1
  description FE to ISP2 via CSW2
  ip address 172.24.244.81 255.255.255.248
  ip flow ingress
  ip nat outside
  ip virtual-reassembly
  speed 100
  full-duplex
!
interface Serial0/0/0
  no ip address
  shutdown
!
router eigrp 1
  redistribute static metric 10000 100 255 1 1500
  passive-interface default
  no passive-interface FastEthernet0/0
  no passive-interface FastEthernet0/1
  network 10.1.192.20 0.0.0.0
  network 10.1.220.4 0.0.0.0
  network 172.24.244.81 0.0.0.0
  no auto-summary
!
ip forward-protocol nd
ip route 0.0.0.0 0.0.0.0 172.24.244.86 track 1
!
ip flow-export source Loopback0
ip flow-export version 5
ip flow-export destination 10.1.152.1 9996
!
ip http server
no ip http secure-server
ip nat inside source list INTERNAL-NETWORKS interface FastEthernet0/1 overload
ip nat inside source static 10.1.152.1 172.24.244.82
!
ip access-list standard INTERNAL-NETWORKS
  permit 10.1.128.0 0.0.127.255
!
logging source-interface Loopback0

```

```

logging 10.1.152.1
snmp-server community cisco RO
snmp-server community san-fran RW
snmp-server ifindex persist
snmp-server trap-source Loopback0
snmp-server location TSHOOT Lab Facility
snmp-server contact support@mgmt.tshoot.local
snmp-server enable traps snmp authentication linkdown linkup coldstart
warmstart
snmp-server enable traps vrrp
snmp-server enable traps ds1
snmp-server enable traps tty
snmp-server enable traps eigrp
snmp-server enable traps envmon
snmp-server enable traps flash insertion removal
snmp-server enable traps icsudsu
snmp-server enable traps isdn call-information
snmp-server enable traps isdn layer2
snmp-server enable traps isdn chan-not-avail
snmp-server enable traps isdn ietf
snmp-server enable traps ds0-busyout
snmp-server enable traps ds1-loopback
snmp-server enable traps aaa_server
snmp-server enable traps atm subif
snmp-server enable traps bgp
snmp-server enable traps bulkstat collection transfer
snmp-server enable traps cnpd
snmp-server enable traps config-copy
snmp-server enable traps config
snmp-server enable traps dsp card-status
snmp-server enable traps entity
snmp-server enable traps event-manager
snmp-server enable traps frame-relay
snmp-server enable traps frame-relay subif
snmp-server enable traps hsrp
snmp-server enable traps ipmobile
snmp-server enable traps ipmulticast
snmp-server enable traps msdp
snmp-server enable traps mvpn
snmp-server enable traps ospf state-change
snmp-server enable traps ospf errors
snmp-server enable traps ospf retransmit
snmp-server enable traps ospf lsa
snmp-server enable traps ospf cisco-specific state-change nssa-trans-change
snmp-server enable traps ospf cisco-specific state-change shamlink interface-
old
snmp-server enable traps ospf cisco-specific state-change shamlink neighbor
snmp-server enable traps ospf cisco-specific errors
snmp-server enable traps ospf cisco-specific retransmit
snmp-server enable traps ospf cisco-specific lsa
snmp-server enable traps pim neighbor-change rp-mapping-change invalid-pim-
message
snmp-server enable traps pppoe
snmp-server enable traps cpu threshold
snmp-server enable traps rsvp
snmp-server enable traps rtr
snmp-server enable traps syslog
snmp-server enable traps l2tun session
snmp-server enable traps vtp
snmp-server enable traps isakmp policy add
snmp-server enable traps isakmp policy delete
snmp-server enable traps isakmp tunnel start
snmp-server enable traps isakmp tunnel stop
snmp-server enable traps ipsec cryptomap add
snmp-server enable traps ipsec cryptomap delete
snmp-server enable traps ipsec cryptomap attach
snmp-server enable traps ipsec cryptomap detach
snmp-server enable traps ipsec tunnel start

```

```
snmp-server enable traps ipsec tunnel stop
snmp-server enable traps ipsec too-many-sas
snmp-server host 10.1.152.1 version 2c cisco
!
!
!
!
control-plane
!
!
line con 0
  exec-timeout 60 0
  login authentication CONSOLE
line aux 0
line vty 0 4
  transport input telnet ssh
line vty 5 15
  transport input telnet ssh
!
scheduler allocate 20000 1000
ntp update-calendar
ntp server 192.168.224.1
ntp server 172.24.240.1
end
```

Student Notes

Use this notes section to write down any logical configuration details that you think are important to document for future troubleshooting.

[illegible]

Student Notes

Use this notes section to write down any physical configuration details that you think are important to document for future troubleshooting.

This image shows a single sheet of white paper with horizontal ruling lines. The lines are evenly spaced and run across the width of the page. There are no margins, text, or other markings on the paper.

Activity Verification

You have completed this task when you attain these results:

- You have used the provided table to identify and document the available network maintenance services, tools, and applications that are needed to support your troubleshooting process.
- You have clarified any questions that you might have about the design and configuration of your lab pod with your instructor.

Lab Debrief Notes

Use these notes sections to write down the primary learning points that are discussed during the Lab Debrief.

Lab 3-1: Alternate Solutions

[illegible]

Lab 3-1: Alternate Methods and Processes

[illegible]

[illegible]

Lab 3-1: Important Commands and Tools

[illegible]

Lab 4-1: Layer 2 Connectivity and Spanning Tree

Complete this lab activity to practice what you learned in the related module.

Activity Objective


In this activity, you will troubleshoot various Layer 2 and spanning-tree problems. After completing this activity, you will be able to meet these objectives:

- Diagnose and resolve Layer 2 connectivity problems
- Diagnose and resolve spanning-tree problems
- Document troubleshooting progress, configuration changes, and problem resolution

Information Packet

The figure illustrates what you will accomplish in this activity.

Visual Objective



Required Resources

These are the resources and equipment that are required to complete this activity:

- One PC with Internet access per team member
- Four Cisco Catalyst 3560 Series Switches per team
- Six Cisco 1841 Integrated Services Routers per team
- Three client PCs per team
- One server per team

Job Aids

These job aids are available to help you complete the lab activity:

- Trouble tickets
- Troubleshooting log

Trouble Ticket A: Switch Replacement Gone Bad

Late yesterday afternoon, access switch ASW1 failed and you quickly concluded that the power supply had gone bad and that the switch needed to be replaced. Luckily, you still had a comparable switch on the shelf and you tasked a couple of your junior colleagues (who have only been with the company for two weeks) with the replacement of this switch so that you could evaluate their skill level.

This morning, when you came in and asked them how things went, they told you that they stayed late trying to restore ASW1, but in the end, they could not, so they ask you to have a look because they are out of ideas. When you ask them what the exact problem is, they tell you that they do not know and that it “simply does not work.”

Users on the first floor have already started to complain that they cannot get access to the network and they had expected this problem to be fixed today.

Your task is to diagnose the issues and restore switch ASW1 as a fully functional access switch on the network.

Trouble Ticket B: Guest Access Problem in Branch

This morning, there was a call from one of the branch offices: An external consultant came in today and needs access to the Internet and email. His PC, CLT3, was plugged into one of the outlets that are patched to the guest VLAN on switch BSW1. However, he has not been able to get an IP address and cannot get onto the network.

Your task is to diagnose and solve this problem, making sure that the consultant gets Internet access.

Trouble Ticket C: Internet Service Provider 1 Seems to Be Down

The network management system has reported that the connection to Internet Service Provider 1 is down. The connection to Internet Service Provider 1 is tracked by pinging the IP address of their router. This issue does not cause any immediate problems because all traffic is routed via Internet Service Provider 2, but the issue needs to be researched and either solved or escalated to Internet Service Provider 1.

Your task is to research this issue and then to either resolve the problem, or if it cannot be resolved on your side, to escalate it to Internet Service Provider 1 with a clear report of why you think that the problem is on their end.

Instructions

Together with your team members, create a troubleshooting plan to divide the work, assign each team member appropriate roles, and coordinate device access between team members. Together, work on Trouble Tickets A, B, and C to resolve the issues. Document your progress in the following Troubleshooting Logs in order to help facilitate efficient communication within the team and to have an overview of your troubleshooting process for reference during the Lab Debrief discussions.

You are allowed a total of two hours to complete as many of the trouble tickets as you can. After two hours, the instructor will debrief the lab and review all trouble tickets and their solutions. The main objective for the troubleshooting labs in this course is to give you an opportunity to practice structured troubleshooting. Fixing the problems is secondary to practicing proper processes and procedures.

Lab Setup

The instructor will provide you with directions to prepare the lab equipment for this lab. After the instructor indicates that the lab is fully prepared, you are ready to start troubleshooting.

Troubleshooting Log

Use this log to document your actions and results during the troubleshooting process.

Trouble Ticket A

Your task is to diagnose the issues and restore switch ASW1 as a fully functional access switch on the network.

Note Refer to the Activity Verification items at the end of this log to verify that you have successfully completed this task.

[illegible]

Troubleshooting Log

Use this log to document your actions and results during the troubleshooting process.

Trouble Ticket B

Your task is to diagnose and solve this problem, making sure that the consultant who is using client PC CLT3 has Internet access.

Note Refer to the Activity Verification items at the end of this log to verify that you have successfully completed this task.

[illegible]

Troubleshooting Log

Use this log to document your actions and results during the troubleshooting process.

Trouble Ticket C


Your task is to research the issue of the failing access to router, ISP1, and then to either resolve the problem, or if it cannot be resolved on your side, to escalate it to Internet Service Provider 1 with a clear report of why you think that the problem is on their end.

Note Refer to the Activity Verification items at the end of this log to verify that you have successfully completed this task.

[illegible]

Lab 4-1: Sample Troubleshooting Flows

The figure illustrates an example of a method that you could follow to diagnose and resolve Layer 2 problems.



Usually, you would start troubleshooting the Layer 2 connectivity between devices because you have discovered that there is no Layer 3 connectivity between two adjacent Layer 2 hosts, such as two hosts in the same VLAN or a host and its default gateway. The following issues are typical symptoms that could lead you to start examining Layer 2 connectivity:

- Failing pings between adjacent devices. (Keep in mind, though, that this problem may also be caused by a host-based firewall that is blocking pings.)
- ARP failures. After clearing the ARP cache and triggering a connection attempt (for instance, via the **ping** command), ARP entries show up as “incomplete” or are missing.
- Use of a packet sniffer on the receiving host shows that packets are not being received.

Confirm Layer 3 Connectivity

Incomplete ARP entries indicate Layer 2 connectivity problems.

```
CSW1#ping 10.1.128.75
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.128.75, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
CSW1#
CSW1#show arp

```

Protocol	Address	Age (min)	Hardware Addr	Type	Interface
Internet	10.1.255.252	-	001b.2b93.b1c7	ARPA	Vlan100
Internet	10.1.255.254	-	0000.0c07.ac64	ARPA	Vlan100
Internet	10.1.128.190	-	0000.0c07.ac0b	ARPA	Vlan11
Internet	10.1.128.188	-	001b.2b93.b1c2	ARPA	Vlan11
Internet	10.1.128.254	-	0000.0c07.ac0c	ARPA	Vlan12
Internet	10.1.128.252	-	001b.2b93.b1c3	ARPA	Vlan12
Internet	10.1.128.75	0	Incomplete	ARPA	
Internet	10.1.128.65	4	000c.2905.19bd	ARPA	Vlan10
Internet	10.1.252.2	160	001b.2a44.3941	ARPA	Vlan100
Internet	10.1.128.126	-	0000.0c07.ac0a	ARPA	Vlan10
Internet	10.1.252.1	228	001b.0c81.2741	ARPA	Vlan100
Internet	10.1.128.124	-	001b.2b93.b1c1	ARPA	Vlan10

© 2009 Cisco Systems, Inc. All rights reserved.


TSHOOT v1.0-15

The most relevant fields in the output are the IP address, hardware address, and interface fields, because these give you the essential information that you are usually looking for when you issue the **show arp** command.

The age field is also relevant. By default, ARP entries are cached for four hours, so to make sure that you are looking at current information, you can use the **clear arp-cache** command to flush existing entries from the cache.

If there is a “-” in the age field instead of a number, this entry is local to the router. In other words, these entries represent locally configured IP and MAC addresses and the router will respond to ARP requests for these entries.

Sample Layer 2 Troubleshooting Flow



© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0-16

When you have determined that the problem is most likely a Layer 2 or Layer 1 problem, you need to reduce the scope of the potential failures. You can diagnose Layer 2 problems with this common troubleshooting method:

- Determine the Layer 2 path. Based on documentation, baselines, and knowledge of your network in general, the first step is to determine the path that you would expect frames to follow between the affected hosts. Determining the expected traffic path beforehand helps you in two ways: It gives you a starting point for gathering information about what is actually happening on the network and it makes it easier to spot abnormal behavior. The second step in determining the Layer 2 path is to follow the expected path and verify that the links on the expected path are actually up and forwarding traffic. If the actual traffic path is different from your expected path, this step may give you clues about the particular links or protocols that are failing and the cause of these failures.
- Track the flow of traffic across the Layer 2 path. By following the expected Layer 2 path and verifying that frames actually flow along that path, you are likely find the exact spot where the connectivity is failing.
- When you have found the spot where the connectivity is failing, examine the link or links where the path is broken. Now you can apply targeted troubleshooting commands to find the root cause of the problem. Even if you cannot find the underlying cause of the problem yourself, by reducing the scope of the problem, you now have a better-defined problem that can be escalated to the next level of support.

Although there are many different approaches to troubleshooting Layer 2 problems, the elements mentioned here would most likely be part of any methodical approach. These elements are not necessarily executed in the presented order. Determining the expected path and verifying the actual path often must be done together.

To determine the traffic path between the affected hosts, you can combine knowledge from the following sources:

- **Documentation and baselines:** Documentation that was written during design and implementation should usually contain information about the intended traffic paths between hosts. If the documentation does not provide this information, you can usually reconstruct the expected flow of traffic by analyzing network diagrams and configurations.
- **Link status across the path:** After you have determined the expected path of the traffic, a very straightforward check you can do is to verify that all ports and links in the path are operational.
- **Spanning-tree topology:** Specifically, in Layer 2 networks that have a level of redundancy built in to the topology, you should analyze the operation of the STP to determine which of the available links will actually be used.

Verify Link Status

The **show interface status** lists essential Layer 2 information in a brief overview.

```
DSW1#show interface status
```

Port	Name	Status	Vlan	Duplex	Speed	Type
Fa0/1	FE to ASW1	connected	trunk	a-full	a-100	10/100BaseTX
Fa0/2	FE to ASW1	connected	trunk	a-full	a-100	10/100BaseTX
Fa0/3	FE to ASW2	connected	trunk	a-full	a-100	10/100BaseTX
Fa0/4	FE to ASW2	connected	trunk	a-full	a-100	10/100BaseTX
Fa0/5	FE to DSW2	connected	trunk	a-full	a-100	10/100BaseTX
Fa0/6	FE to DSW2	connected	trunk	a-full	a-100	10/100BaseTX
Fa0/7		notconnect	1	auto	auto	10/100BaseTX
Fa0/8		notconnect	1	auto	auto	10/100BaseTX
Fa0/9	Reserved for later	disabled	1	auto	auto	10/100BaseTX
Fa0/10	Reserved for later	disabled	1	auto	auto	10/100BaseTX
Fa0/11		notconnect	1	auto	auto	10/100BaseTX
Fa0/12		notconnect	1	auto	auto	10/100BaseTX

© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.9-17

To determine link status on switches, the **show interface status** command is very useful because it gives a brief overview of all the interfaces on the switch, yet contains essential elements, such as link status, speed, duplex, trunk or VLAN membership, and interface descriptions.

Verify Link Status (Cont.)

If Cisco Discovery Protocol is enabled between network devices, **show cdp neighbor** can help confirm correct Layer 2 operation and pinpoint potential cabling problems.

```
CSW1#show cdp neighbors
Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater, P - Phone

Device ID         Local Intrfce   Holdtme    Capability  Platform  Port ID
CRO1.mgmt.tshoot.local
Fas 0/11          151           R S I      1841        Fas 0/0
CRO2.mgmt.tshoot.local
Fas 0/12          152           R S I      1841        Fas 0/0
IRO2.mgmt.tshoot.local
Fas 0/16          152           R S I      1841        Fas 0/0
IRO1.mgmt.tshoot.local
Fas 0/15          150           R S I      1841        Fas 0/0
CSW2.mgmt.tshoot.local
Fas 0/4           132           R S I      WS-C3560-   Fas 0/4
CSW2.mgmt.tshoot.local
Fas 0/3           132           R S I      WS-C3560-   Fas 0/3
```

© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0--18

If the Cisco Discovery Protocol is enabled between your switches and routers, the **show cdp neighbor** command can be very useful in helping you to confirm that a link is operational at the data link layer in both directions. This command is also essential for use in uncovering cabling problems because it records both the sending and receiving ports, as you can see in the **show** command output.

Analyze Spanning Tree

The **show spanning tree vlan *vlan-id*** command is a good starting point for spanning tree analysis.

```
ASW1#show spanning-tree vlan 17

VLAN0017
  Spanning tree enabled protocol rstp
  Root ID    Priority    24593
            Address    001f.2721.8400
            Cost        12
            Port        64 (Port-channel2)
            Hello Time  2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    32785 (priority 32768 sys-id-ext 17)
            Address    001e.79a9.b580
            Hello Time  2 sec  Max Age 20 sec  Forward Delay 15 sec
            Aging Time  300

Interface                Role Sts Cost      Prio.Nbr Type
-----
Fa0/7                    Desg FWD 19       128.9    P2p Edge
Po1                      Altn BLK 12       128.56   P2p
Po2                      Root FWD 12       128.64   P2p
```

© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0-18

To analyze the spanning tree topology and the consequences that the spanning tree protocol has for the Layer 2 path, **show spanning-tree vlan *vlan-id*** is a good starting point. The output from this command lists all essential parameters that affect the topology, such as root port, designated ports, port state, and port type.


Typical values for the port status field are BLK (blocking) and FWD (forwarding). You might also see LTN (listening) or LRN (learning) while the STP is converging.

The states LBK (loopback), DWN (down) or BKN* (broken) typically indicate problems. In the case of a broken (BKN) port status, the type field gives an additional indication of what is causing the broken status. Possible values could be “*ROOT_Inc,” “*LOOP_Inc,” “*PVID_Inc,” “*TYPE_Inc,” or “*PVST_Inc.” To get a more detailed description of the type of inconsistency and what might be causing it, you can examine the output of the **show spanning-tree inconsistentports** command.

Typical values for the type field are as follows:

- “P2p” or “Shr” to indicate the link type (typically, based on duplex status).
- “Edge” for edge (portfast) ports.
- “Bound” for boundary ports, in the case where this switch is running 802.1s (MST) and the other switch is running a different spanning tree variety. The output also indicates which other type of STP was detected on the port.
- “Peer” for peer ports, in the case where this switch is running PVST+ or PVRST+ and the other switch is running a different standard variety of the Spanning Tree Protocol (802.1D or 802.1s MST).

Sample Layer 2 Troubleshooting Flow



© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0-20

When you have determined the Layer 2 path between the two affected hosts, you can start tracking the traffic between the hosts as it is being switched along the path. The most direct approach to tracking the traffic is to capture packets at set points along the path by using a packet sniffer. Tracking packets in real time is a fairly intensive procedure and you may find that there are technical limitations that restrict the links where traffic captures could be collected. However, tracking packets yields the most definitive proof that traffic is or is not flowing along specific paths and links. A less labor-intensive method that you can use is to track the flow of traffic by analyzing MAC address tables or traffic statistics. These methods are less direct, since you are not looking at the actual traffic itself, but at traces left by the passing of frames.

In a network that has not gone into production yet, packet statistics may help you see where traffic is flowing. On live networks, the test traffic that you are generating is, in most cases, lost against the background of the live traffic patterns. However, if the switches that you are using have the capability to track packet statistics for access lists, you may be able to write an access list that matches the specific traffic that you are interested in and isolate the traffic statistics for that type of traffic.

A method of tracing traffic that you can use under all circumstances is to analyze the process of MAC address learning along the Layer 2 path. When a switch receives a frame on a particular port and for a particular VLAN, it records the source MAC address of that frame together with the port and VLAN in the MAC address table. Therefore, if the MAC address of the source host is recorded in a switch, but the address is not on the next switch in the path, the missing address indicates a communication problem between these switches for the VLAN concerned. This existence of this situation indicates that you should do a detailed examination of the link between these switches.

Analyze MAC Address Tables

- The **show mac-address-table address *mac-address*** command is useful for verifying that a particular MAC address is being learned and to see its port and VLAN association.

```
CSW1#show mac address-table address 0050.5684.44b6
Mac Address Table
-----
Vlan    Mac Address      Type    Ports
----
17      0050.5684.44b6   DYNAMIC Po10
Total Mac Addresses for this criterion: 1
```


- Another useful option is the **show mac-address-table interface *intf-id*** command to determine if any MAC addresses are being learned on a particular interface.

The **show mac-address-table** command can be used to check the content of the MAC address table. Since this table usually contains hundreds to thousands of entries, you have to use command options to narrow the results to find what you are looking for.

In many cases, you are looking for the MAC address of a specific host. To select a specific MAC address entry in the table, you can use the **show mac-address-table address *mac-address*** option.

Another useful option you can use is the **show mac-address-table interface *intf-id*** option, which allows you to see which MAC addresses were learned on a specific port.

Sample Layer 2 Troubleshooting Flow



© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0--22

When you have found the spot in the Layer 2 path where one switch is learning the source MAC address and the next switch is not, you should examine the link between those two switches carefully.

What could cause the MAC address not to be learned on the next switch?

Does the VLAN exist on the next switch? Is there an operational trunk between the two switches? Is the VLAN allowed on the trunk between the switches? If there is an EtherChannel between the switches, is that EtherChannel operational?

Verify VLAN Existence

The **show vlan brief** command can be used to see a summary of all existing VLANs.

```
ASW1#show vlan brief
```

VLAN	Name	Status	Ports
1	default	active	Gi0/1, Gi0/2
17	FIS1-OFFICE	active	Fa0/7
18	FIS1-VOICE	active	Fa0/7
19	FIS1-GUEST	active	
128	MGMT	active	
1000	NATIVE	active	
1001	UNUSED	active	Fa0/3, Fa0/4, Fa0/8, Fa0/9 Fa0/10, Fa0/11, Fa0/12, Fa0/13 Fa0/14, Fa0/15, Fa0/16, Fa0/17 Fa0/18, Fa0/19, Fa0/20, Fa0/21 Fa0/22, Fa0/23, Fa0/24
1002	fddi-default	act/unsup	
1003	token-ring-default	act/unsup	
1004	fddinet-default	act/unsup	
1005	trnet-default	act/unsup	

© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0-23

To get a quick overview of all existing VLANs, you can use the **show vlan brief** command. It is important for you to note that in the output of this command, trunk ports are not listed. For instance, in the sample output in the figure you can see that FastEthernet 0/7 is listed as the only port in VLAN 17.

Verify VLAN Existence (Cont.)

- The **show vlan id *vlan-id*** is a good way to verify the existence of a particular VLAN on a switch.
- This commands lists all ports that are associated with a VLAN including trunk ports.

```
ASW1#show vlan id 17
```

VLAN Name	Status	Ports
17 F1S1-OFFICE	active	Fa0/7, Po1, Po2

VLAN Type	SAID	MTU	Parent	RingNo	BridgeNo	Stp	BrdgMode	Trans1	Trans2
17 enet	100017	1500	-	-	-	-	-	0	0

Remote SPAN VLAN

Disabled

Primary	Secondary	Type	Ports
-----	-----	-----	-----

© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0-24

To verify the existence of a particular VLAN on a switch, you can use the **show vlan id *vlan-id*** command. This command shows you whether the VLAN exists and, if so, which ports are assigned to it. Note that this command includes trunk ports on which the VLAN is allowed. For the same VLAN 17 that was referenced in the previous figure, you now see that interface Port-channel 1 and Port-channel 2 are also listed as ports that are associated with VLAN 17.

Verify Trunk Operation

The **show interface trunk** command is useful to get an overview of trunk operation for all ports configured as trunks.

```
ASW1#show interfaces trunk

Port      Mode      Encapsulation  Status        Native vlan
Po1       on        802.1q         trunking      1000
Po2       on        802.1q         trunking      1000

Port      Vlans allowed on trunk
Po1       17-19,128
Po2       17-19,128

Port      Vlans allowed and active in management domain
Po1       17-19,128
Po2       17-19,128

Port      Vlans in spanning tree forwarding state and not pruned
Po1       128
Po2       17-19
```

© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0-25

The easiest way you can get an overview of trunk operation is to use of the **show interface trunk** command. Not only does it list trunk status, trunk encapsulation, and native VLAN, but it also displays the list of allowed VLANs, the list of active VLANs, and the list of VLANs that are in the spanning tree forwarding state for the trunk. The last list can be very helpful in determining whether frames for a particular VLAN will be forwarded on a trunk.

For instance, in the example in the figure, you can see that both interface Port-channel 1 and Port-channel 2 allow VLANs 17 to 19 and 128, but VLAN 128 is forwarded on Port-channel 1 while VLANs 17 to 19 are forwarded on Port-channel 2.

Verify VLAN Port Status

The **show interface *intf-id* switchport** command can be used to get an overview of all VLAN-related parameters on an interface.

```
ASW1#show interfaces FastEthernet 0/7 switchport
Name: Fa0/7
Switchport: Enabled
Administrative Mode: static access
Operational Mode: static access
Administrative Trunking Encapsulation: negotiate
Operational Trunking Encapsulation: native
Negotiation of Trunking: Off
Access Mode VLAN: 17 (F1S1-OFFICE)
Trunking Native Mode VLAN: 1 (default)
Administrative Native VLAN tagging: enabled
Voice VLAN: 18 (F1S1-VOICE)
Administrative private-vlan host-association: none
Administrative private-vlan mapping: none
Administrative private-vlan trunk native VLAN: none
Administrative private-vlan trunk Native VLAN tagging: enabled
Administrative private-vlan trunk encapsulation: dot1q
Administrative private-vlan trunk normal VLANs: none
Administrative private-vlan trunk associations: none
Administrative private-vlan trunk mappings: none
<...further output omitted...>
```

© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0-26

The **show interface *intf-id* switchport** command is useful for checking all VLAN-related parameters for a specific interface. You can use this command for checking access ports and trunk ports. For instance, in the example in the figure, you can see that the port is configured as a static access port in VLAN 17 and VLAN 18 is assigned to the port as a voice VLAN.

Verify EtherChannel Operation

The **show etherchannel summary** gives you a quick overview of all existing EtherChannel groups on a switch, the relation between Port channel interfaces and physical interfaces, and the status of the EtherChannels.

```
ASW1#show etherchannel summary
Flags:  D - down          P - bundled in port-channel
        I - stand-alone  S - suspended
        H - Hot-standby (LACP only)
        R - Layer3       S - Layer2
        U - in use       f - failed to allocate aggregator

        M - not in use, minimum links not met
        u - unsuitable for bundling
        w - waiting to be aggregated
        d - default port

Number of channel-groups in use: 2
Number of aggregators:          2

Group  Port-channel  Protocol    Ports
-----+-----+-----+-----
1      Po1 (SU)      -           Fa0/1 (P)  Fa0/2 (P)
2      Po2 (SU)      -           Fa0/5 (P)  Fa0/6 (P)
```

© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0-27

When an EtherChannel is configured between the switches and you suspect that EtherChannel operation may be causing the communication failure between the switches, you can verify this fact by using the **show etherchannel summary** command. Although the command output is fairly self-explanatory, the typical things that you should look for are the flag “(s)”, which indicates that a (physical) interface is suspended because of incompatibility with the other ports in the channel, or the flag “(D)” which indicates that an interface (physical or port channel) is down.

Lab Debrief Notes

Use these notes sections to write down the primary learning points that are discussed during the Lab Debrief.

Lab 4-1: Alternate Solutions

This image shows a blank sheet of white paper with horizontal ruling lines. The lines are evenly spaced and run across the width of the page. There are no margins, text, or other markings on the paper.

[illegible]

Lab 4-1: Procedure and Communication Improvements

[illegible]

[illegible]

Lab 4-1: References

If you need more information on the commands and their options, you can go to the following sections of <http://www.cisco.com>.

- Cisco Systems, Inc. Command References for Cisco Catalyst LAN Switches: Go to Product Support (<http://www.cisco.com/web/psa/products/index.html>), select Switches, select LAN Switches and then the product family that you are working with. The Command References can then be found under the “Reference Guides” section.
- Cisco Systems, Inc. Virtual LANs/VLAN Trunking Protocol (VLANs/VTP) Troubleshooting TechNotes:
http://www.cisco.com/en/US/tech/tk389/tk689/tsd_technology_support_troubleshooting_technotes_list.html
- Cisco Systems, Inc. Spanning Tree Protocol Troubleshooting TechNotes:
http://www.cisco.com/en/US/tech/tk389/tk621/tsd_technology_support_troubleshooting_technotes_list.html
- Cisco Systems, Inc. EtherChannel Troubleshooting TechNotes:
http://www.cisco.com/en/US/tech/tk389/tk213/tsd_technology_support_troubleshooting_technotes_list.html

Lab 4-2: Layer 3 Switching and First-Hop Redundancy

Complete this lab activity to practice what you learned in the related module.

Activity Objective


In this activity, you will troubleshoot various problems related to Layer 3 switching and FHRPs, such as the HSRP and the GLBP. After completing this activity, you will be able to meet these objectives:

- Diagnose and resolve problems related to SVIs and multilayer switching
- Diagnose and resolve problems related to FHRPs such as HSRP and GLBP
- Document troubleshooting progress, configuration changes, and problem resolution

Information Packet

The figure illustrates what you will accomplish in this activity.

Visual Objective



Required Resources

These are the resources and equipment that are required to complete this activity:

- One PC with Internet access per team member
- Four Cisco Catalyst 3560 Series Switches per team
- Six Cisco 1841 Integrated Services Routers per team
- Three client PCs per team
- One server per team

Job Aids

These job aids are available to help you complete the lab activity.

- Trouble tickets
- Troubleshooting log

Trouble Ticket D: Switch ASW1 Cannot Be Managed from Server SRV1

When you come into the office this morning, you find the following ticket in the system:

“Switch ASW1 has been showing CRC errors on a group of eight ports for several days. Hardware was suspected to be the cause. During the maintenance window yesterday evening, the switch was swapped with a similar switch from our lab. The configuration was pasted in on the console. After this replacement, clients could connect and no errors were shown on the ports. However, making a backup to server SRV1 did not work, nor is the switch reachable via Telnet or SSH from server SRV1. Unfortunately, there was no time for additional research yesterday, but because there is no impact to the users, it was decided to leave the switch and pick up this issue the next day. Please follow up.”

Your task is to diagnose the issue and restore connectivity between switch ASW1 and server SRV1. After resolving the problem, make a backup of the configuration to server SRV1.

Trouble Ticket E: Failover not Functioning as Expected

During the maintenance window last Friday, a series of failover tests between headquarters and the branch offices were executed. As a result of these tests, it was discovered that, during a reboot of router BRO1, connectivity between clients in the VLAN B1S1-OFFICE and hosts in the LANs at headquarters is lost. After router BRO1 comes back online, the clients regain connectivity. In addition, management connectivity between server SRV1 and switch BSW1 on VLAN 128 is also lost during the failover. This behavior is not the expected behavior, because the network is fully redundant and both a routing protocol (EIGRP) and first-hop redundancy protocols (HSRP at headquarters and GLBP in the branch office) have been configured to ensure correct failover during outages.

Most of the users in the branch office are out of the office to attend training, so although it is not an official maintenance window, you have been authorized to run necessary failover tests during office hours. However, the disruption to the remaining branch office users should be kept to a minimum.

Your task is to diagnose this issue and restore the functionality of the failover mechanisms, as intended in the design.

Trouble Ticket F: Verify HSRP Authentication

Several weeks ago, an external company performed a security audit on the network. One of the exposed attack vectors—or weaknesses— was that a DoS attack could be launched against the HSRP protocol. The recommended solution was to use MD5-based authentication between the HSRP routers. One of your colleagues has been too busy to implement this solution in a test-VLAN in the LAN (VLAN 44) at headquarters before rolling it out on all LANs.

Yesterday, just before this colleague left for a two-week vacation, she asked you to see if somebody else could finalize the tests and to guarantee that it can be rolled out as soon as she returns.

Your task is to review and verify the implementation of HSRP authentication in VLAN 44 and fix any issues that may remain.

Trouble Ticket G: HSRP and GLBP Comparison

The failover tests that were executed last Friday (as mentioned in trouble ticket E) have caused another scenario to be implemented and tested. One of the network engineers who works at Branch Office 1 has always said that it would be better to use HSRP instead of GLBP. The fact that the failover tests did not work out as expected has now caused him to push for a good comparative test of the failover behavior of the two protocols and revert to HSRP, unless it can be proven that GLBP functions at least as well as HSRP where failover is concerned. You receive a phone call from him in which he asks you to look at the configuration because it is frustrating him. Somehow, he cannot get HSRP to work in his test-VLAN (VLAN 1000) and now that he has pushed for this test, he has to make it work. You offer to look and help him run his tests.

Your task is to diagnose and resolve the problems with HSRP in the newly configured VLAN 1000 on routers BRO1 and BRO2, and to execute failover tests to compare the behavior of GLBP and HSRP. To minimize the disruption on the network, these tests should be coordinated with the rest of the team, specifically with the team members that are working on Trouble Ticket D.

Note	You are allowed to assign PC CLT3 to the test-VLAN to test the HSRP failover. Make sure that you reassign the PC to the guest VLAN and verify proper operation after you have finished your tests.

Instructions

Together with your team members, create a troubleshooting plan to divide the work, assign each team member appropriate roles, and coordinate device access between team members. Together, work on Trouble Tickets D, E, F, and G to resolve the issues. Document your progress in the following Troubleshooting Logs in order to help facilitate efficient communication within the team and to have an overview of your troubleshooting process for reference during the Lab Debrief discussions.

You are allowed a total of two hours to complete as many of the trouble tickets as you can. After two hours, the instructor will debrief the lab and review all trouble tickets and their solutions. The main objective for the troubleshooting labs in this course is to give you an opportunity to practice structured troubleshooting. Fixing the problems is secondary to practicing proper processes and procedures.

Lab Setup

The instructor will provide you with directions to prepare the lab equipment for this lab. After the instructor indicates that the lab is fully prepared, you are ready to start troubleshooting.

Troubleshooting Log

Use this log to document your actions and results during the troubleshooting process.

Trouble Ticket D

Your task is to diagnose the issue and restore connectivity between switch ASW1 and server SRV1. After resolving the problem, make a backup of the configuration to server SRV1.

Note Refer to the Activity Verification items at the end of this log to verify that you have successfully completed this task.

[illegible]

Troubleshooting Log

Use this log to document your actions and results during the troubleshooting process.

Trouble Ticket E

Your task is to diagnose the redundancy issues between the headquarters and the branch office and restore the functionality of the failover mechanisms, as intended in the design.

Note Refer to the Activity Verification items at the end of this log to verify that you have successfully completed this task.

[illegible]

Troubleshooting Log

Use this log to document your actions and results during the troubleshooting process.

Trouble Ticket F

Your task is to review and verify the implementation of HSRP authentication in VLAN 44 and fix any issues that may remain.

Note Refer to the Activity Verification items at the end of this log to verify that you have successfully completed this task.

[illegible]

Troubleshooting Log

Use this log to document your actions and results during the troubleshooting process.

Trouble Ticket G

Your task is to diagnose and resolve the problems with HSRP in the newly configured VLAN 1000 on routers BRO1 and BRO2, and to execute failover tests to compare the behavior of GLBP and HSRP. To minimize the disruption on the network, these tests should be coordinated with the rest of the team, specifically with the team members that are working on Trouble Ticket D.


Note Refer to the Activity Verification items at the end of this log to verify that you have successfully completed this task.

[illegible]

Lab 4-2: Sample Troubleshooting Flows

Troubleshooting Multilayer Switching


The figure illustrates an example of a method that you could follow to diagnose and resolve problems related to multilayer switching.



What is multilayer switching? In essence, a multilayer switch is a switch that is capable of switching Ethernet frames based on information in the Layer 2 and Layer 3 headers. Troubleshooting Layer 2 switching was covered in the previous lab exercise; therefore, this troubleshooting flow focuses on troubleshooting the process of switching Ethernet frames based on Layer 3 information.

Under what circumstances would you start troubleshooting the multilayer switching process? Troubleshooting multilayer switching is just one of the steps in the bigger picture of troubleshooting network connectivity along a Layer 3 path. You would start troubleshooting multilayer switches when you have determined—by using tools like **traceroute** or **ping** or through analysis of packet captures—that a particular hop in the Layer 3 path seems to be the point where packets start to get dropped, and that hop turns out to be a multilayer switch. At that point, start tracing and verifying the Layer 3 forwarding behavior of the multilayer switch that you suspect is causing the problem. When you are troubleshooting performance problems and you want to find the exact physical links on which packets travel, you would use the same method.

Sample Multilayer Switching Troubleshooting Flow



© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0-30

Layer 3 packet switching generally consists of three major steps:

- Receiving the packet on a Layer 3 interface. This interface can either be a routed port or a SVI.
- Performing a lookup in the hardware packet switching data structures. Multilayer switches store packet forwarding information in special TCAM data structures. The information contained in these data structures is compiled from the Cisco Express Forwarding data structures in the main memory of the route processor, and these data structures are derived, in turn, from control plane tables, such as the routing table and the ARP cache.
- Rewriting the frame and switch it to the outbound interface based on the information that is found in the TCAM.

Consequently, a straightforward approach that you can use to troubleshoot a Layer 3 switching problem is to verify the components that are involved in this process. First, verify the ingress Layer 3 interface, then the control plane data structures, and, subsequently, the packet forwarding data structures. (Alternatively, these steps could be taken in the reverse order).

If the ingress interface is a routed port, the first step in this process is simple because the Layer 3 and Layer 2 ports are identical. You can determine the status of the Layer 3 ingress interface just by verifying the physical interface status and the configured IP address and subnet mask for that interface. However, if the ingress interface is a SVI, its status is not directly related to any particular physical interface.

Verify SVI Status

VLAN interfaces are down unless at least one interface in the VLAN is in spanning tree forwarding state.

One of the probable causes for an SVI being down is:

- Missing VLANs

```
CSW1#show vlan id 100
VLAN id 100 not found in current VLAN database

CSW1#show interface vlan 100
Vlan100 is down, line protocol is down
Hardware is EtherSVI, address is 001e.f7bb.f7d6 (bia 001e.f7bb.f7d6)
Internet address is 10.100.100.100/24
<... further output omitted ...>
```

A VLAN interface or SVI is up if at least one interface for that VLAN is in the spanning tree forwarding state. This status implies that, if an SVI is down, you should verify the existence of the VLAN, VLAN port assignments, and spanning tree state for the SVI.

In this figure, you can see that a missing VLAN results in a VLAN interface that is in state down, line protocol down.

Verify SVI Status (Cont.)

Another probably cause is: VLANs without ports assigned

```
CSW1#show vlan id 100

VLAN Name                Status    Ports
-----
100  TEST                    active

VLAN Type  SAID      MTU    Parent RingNo BridgeNo Stp  BrdgMode Trans1 Trans2
-----
100  enet  100100    1500    -     -     -     -     -         0       0

Remote SPAN VLAN
-----
Disabled

Primary Secondary Type      Ports
-----

CSW1#show interface vlan 100
Vlan100 is up, line protocol is down
  Hardware is EtherSVI, address is 001e.f7bb.f7d6 (bia 001e.f7bb.f7d6)
  Internet address is 10.100.100.100/24
<... further output omitted ...>
```

© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0-32

When the VLAN exists, but no ports are assigned to that VLAN, the status of the SVI changes to up, line protocol down.

Verify SVI Status (Cont.)

Another probably cause is: VLANs with assigned ports, but no ports in spanning tree forwarding state

```
CSW1#show spanning-tree vlan 100
<... output omitted ...>

Interface                Role Sts Cost      Prio.Nbr Type
-----
Po1                      Desg LRN 12        128.56   P2p
Po10                     Desg LRN 12        128.128  P2p


CSW1#sh int vlan 100
Vlan100 is up, line protocol is down
  Hardware is EtherSVI, address is 001e.f7bb.f7d6 (bia 001e.f7bb.f7d6)
  Internet address is 10.100.100.100/24
<... further output omitted ...>
```

Conclusion: If an SVI is down, there may be Layer 2 problems in the associated VLAN.

Finally, if ports are assigned to the VLAN and at least one of these physical ports (trunk or access port) is up, one more condition needs to be met: The spanning tree state for at least one of the ports needs to be the forwarding state. Under normal circumstances, if at least one interface is assigned to a VLAN, then there is at least one interface that is in the spanning tree forwarding state. Either the switch is the root for the VLAN and all the ports assigned to the VLAN are designated ports and therefore are in a forwarding state or the switch is not the root and therefore has a root port that is in the forwarding state.

As a result, when you are troubleshooting a multilayer switching problem and you find that the ingress interface is an SVI and the SVI is down, you know that there is an underlying Layer 2 problem for that VLAN and that you need to initiate a Layer 2 troubleshooting process.

Sample Multilayer Switching Troubleshooting Flow



© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0—34

The next step in this process is to verify that the control plane information that is needed to forward the packets is present. The two control plane data structures that are relevant to multilayer switching are the routing table and the ARP cache.

In this sample troubleshooting flow, you can verify the multilayer switching data structures for an ICMP echo request traveling from source IP address 10.1.128.65 to destination IP address 10.1.160.65 by using various **show** commands.

Verify the Routing Table and ARP Cache

- The routing table shows that a route is available for the destination IP address and it lists the next hop and outbound interface.
- The ARP cache shows that the MAC address for the next hop is known.

```
CSW1#show ip route 10.1.160.65
Routing entry for 10.1.160.64/26
  Known via "eigrp 1", distance 90, metric 20517120, type internal
  Redistributing via eigrp 1
  Last update from 10.1.192.2 on FastEthernet0/11, 00:01:07 ago
  Routing Descriptor Blocks:
  * 10.1.192.2, from 10.1.192.2, 00:01:07 ago, via FastEthernet0/11
    Route metric is 20517120, traffic share count is 1
    Total delay is 20200 microseconds, minimum bandwidth is 128 Kbit
    Reliability 255/255, minimum MTU 1500 bytes
    Loading 1/255, Hops 2

CSW1#show ip arp 10.1.192.2
Protocol Address      Age (min)  Hardware Addr  Type   Interface
Internet 10.1.192.2      211       0019.562c.8fb4  ARPA   FastEthernet0/11

CSW1#show interface FastEthernet 0/11 | include bia
Hardware is Fast Ethernet, address is 001e.f7bb.f7c2 (bia 001e.f7bb.f7c2)
```

© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0-35


In the figure, you can see that a route is found in the routing table for the destination IP address 10.1.160.65 and the next hop and outbound interface for packets with that destination are listed.

If the routing table does not contain an entry (specific prefix or default route) for the destination, the problem is not a packet switching problem, but a routing problem, and you should initiate a process to troubleshoot the routing operation on the control plane.

The ARP cache provides the destination MAC address for the next hop. If an ARP entry for the destination is missing or listed as “incomplete,” either the next hop listed in the route is not valid or there is a Layer 2 problem between the multilayer switch and the next hop. In both cases, the problem is not really a multilayer switching problem, and you should investigate the routing operation on the control plane and the Layer 2 connectivity to the next hop first.

The final element that the router needs in order to rewrite a frame and switch it out is the source MAC address of the frame, which corresponds to the MAC address of the outbound Layer 3 interface.

Sample Multilayer Switching Troubleshooting Flow



© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0-36

When the control plane data structures have been verified, the next step in the multilayer switching troubleshooting process is to verify the data structures in software and in hardware that are used to forward packets.

All recent Layer 3 switches use the Cisco Express Forwarding technology as the foundation for the multilayer switching process. This means that they will combine the information from the control plane data structure, such as the routing table and the ARP cache, into two different data structures: the FIB and the adjacency table. These two data structures are stored in the main memory of the route processor and they are only used to forward packets that are not handled in hardware.

However, based on the information in the FIB and adjacency table, the hardware TCAM will be populated and the resulting TCAM information is what is eventually used to forward frames in hardware.

So to verify the correct operation of the multilayer switching process, you should first verify that the control plane information is accurately reflected in the software FIB and adjacency table and, next, that the information from the FIB and adjacency table is correctly compiled into the TCAM.

Verify the FIB and Adjacency Table

- The Cisco Express Forwarding FIB contains the next-hop and outbound interface information also found in the routing table.
- The adjacency table contains frame rewrite information, which combines information from the ARP cache and egress interface.

```
CSW1#show ip cef 10.1.160.65 detail
10.1.160.64/26, epoch 2
  nexthop 10.1.192.2 FastEthernet0/11

CSW1#show adjacency FastEthernet 0/11 detail
Protocol Interface Address
IP        FastEthernet0/11 10.1.192.2 (24)
          0 packets, 0 bytes
          epoch 0
          sourced in sev-epoch 5
          Encap length 14
          0019562C8FB4001EF7BBF7C20800
          L2 destination address byte offset 0
          L2 destination address byte length 6
          Link-type after encap: ip
          ARP
```

MAC address of egress interface FastEthernet 0/11

MAC address of 10.1.192.2 as listed in ARP cache

The **show ip cef** command can be used in a way that is similar to the way the **show ip route** command is used. When you specify a destination IP address as an option to the command, it lists the entry in the Cisco Express Forwarding FIB that matches that IP address and shows the next-hop IP address and egress interface, which serve as a pointer to the adjacency table.

The command **show adjacency** can be used to display the information contained in the adjacency table. You can specify the next-hop IP address or interface to select specific adjacencies. Adding the **detail** keyword to the command, allows you to see the complete frame rewrite information for packets that will be switched through that adjacency. The frame rewrite information lists the complete Ethernet header. For the example in the figure, the header consists of the destination MAC address 0019562C8FB4 (which is the same MAC address that was listed as the MAC address of next-hop 10.1.192.2 in the ARP cache) followed by the source MAC address 001EF7BBF7C2 (which equals the MAC address of the egress interface Fa 0/11) and, finally, the Ethertype 0x0800 (which indicates that the protocol contained in the Ethernet frame is IP version 4).

The information displayed in these **show** commands should accurately reflect the information in the routing table and ARP cache.

Verify the Hardware TCAM Information

The **show platform forward** command can be used to determine the exact forwarding behavior for a Layer 2 or Layer 3 switched frame.

```
CSW1#show platform forward FastEthernet 0/1 vlan 17 0050.5684.44b6
001e.f7bb.f7c4 ip 10.1.128.65 10.1.160.65 icmp 8 0
Ingress:
Global Port Number: 57, lpn: 2 Asic Number: 0
Source Vlan Id: Real 17, Mapped 3. L2EncapType 0, L3EncapType 0

<...output omitted...>

=====
Egress: Asic 0, switch 1
Source Vlan Id: Real 17, Mapped 3. L2EncapType 0, L3EncapType 0
portMap 0x1000, non-SPAN portMap 0x1000

<...output omitted...>

Port      Vlan    SrcMac      DstMac      Cos  Dscp
Fa0/11    1006    001e.f7bb.f7c2  0019.562c.8fb4
```

© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0-38

Note The **show platform forward** command shown in this figure is specific to the Cisco Catalyst 3560 and 3750 Series Switches. Consult the documentation for the platform that you are working with to find similar commands that can be used to examine the content of the hardware forwarding data structures for the platform.

The **show platform forward** command consults the hardware TCAM information and displays the exact forwarding behavior for a Layer 2 or Layer 3 switched frame.

This command displays the exact forwarding behavior for a packet, taking into account all the features that affect packet forwarding, including Cisco Express Forwarding load balancing, EtherChannel load balancing, and packet filtering using ACLs. Therefore, you have to specify the exact content of all the relevant fields in the header of the packet.

In the example in the figure, you can see that the following fields are specified:

- **Ingress interface:** In the example interface, FastEthernet 0/1 is specified as the ingress interface for the packet.
- **Ingress VLAN:** It is not necessary for you to specify this parameter if the port is an access port, but for trunk ports, you have to specify the VLAN that the frame is tagged with when it enters the ingress interface. In the example, VLAN 17 is specified as the ingress VLAN.
- **Source MAC address:** You need to specify the source MAC address of the frame when it enters the switch. In the example, the address is 0050.5684.44b6. This is the MAC address of the egress interface of the previous hop.
- **Destination MAC address:** You need to specify the destination MAC address of the frame when it enters the switch. In the example, the address is 001e.f7bb.f7c4. For a Layer 3 switched packet, this address is the MAC address of the ingress Layer 3 interface (routed port or SVI).

- **Protocol:** This field is not necessary for Layer 2 switched frames, but for Layer 3 switching, you need to specify the Layer 3 protocol that is being used and the major fields in that protocol's header. In the example, IP is listed as the protocol.
- **Source IP address:** When the IP is specified as the Layer 3 protocol, you need to specify the source IP address of the packet. In the example, it is 10.1.128.65.
- **Destination IP address:** When IP is specified as the Layer 3 protocol, you need to specify the destination IP address of the packet. In the example, it is 10.1.160.65.
- **IP protocol:** When IP is specified as the Layer 3 protocol, you need to specify the IP protocol in the IP header, for example, TCP, UDP, or ICMP. In the example, ICMP is specified because the example represents an ICMP echo request packet.
- **ICMP type and code:** When ICMP is specified as the IP protocol, you need to specify the ICMP type and code values. When TCP or UDP are specified as the protocol, you need to specify additional header fields that are appropriate for those protocols, such as source and destination port numbers. In the example, ICMP type 8 and code 0 are specified to represent an echo request packet.


This command is very powerful because it shows you exactly how frames will be forwarded based on all features that affect forwarding behavior, such as load balancing, EtherChannel, and ACLs. In addition, if a frame would be dropped instead of forwarded, the command lists the reason why the frame will be dropped.

What should you do if somewhere in this chain of verifying the control plane, the software packet forwarding data structures, and the hardware packet forwarding data structures, you find an inconsistency between these data structures?

The process of building the FIB and adjacency table from the routing table and ARP cache, and subsequently populating the TCAM based on the FIB and adjacency table, is a process that is internal to the Cisco IOS Software and not configurable. The lack of configurability means that whenever you find information in these data structures that is not consistent, you should open a case with the Cisco Technical Assistance Center (provided that you have a valid support contract for your device) to investigate and resolve the issue. As a workaround, you can try to clear the control plane data structures, such as the routing table and the ARP cache, for the particular entries that you are troubleshooting. This workaround triggers both the control plane and the packet forwarding data structures to be repopulated for those entries, and in certain cases, this workaround may resolve the inconsistencies. However, this solution is only a workaround, not a real solution, because it only addresses the symptoms of the problem and not the underlying cause.

Troubleshooting First-Hop Redundancy Protocols

The figure illustrates an example of a method that you could follow to diagnose and resolve problems related to first hop redundancy protocols, such as the HSRP, VRRP, and GLBP.



The most common reason for you to start troubleshooting FHRP behavior is that, during an outage or a test, network connectivity is lost for longer than expected when a redundant device or link is (temporarily) disabled. In redundantly configured IP networks, usually, a number of different protocols need to reconverge to recover from a failure, and the FHRP that is used is just one of the protocols that could be the cause of the loss of connectivity. Other protocols that need to converge—and could be the cause of the problem—are routing protocols and the STP.

So how do you determine if the FHRP is the problem?

If you have the opportunity to execute failover tests (for instance, during a scheduled maintenance window), a good way to determine if the problem is caused by the FHRP or by another protocol is by using the following method: Start multiple continuous pings from a client that is using the virtual router as its default gateway. Ping to the virtual and real IP addresses of the routers that participate in the FHRP, and ping to an IP address of a host that is one or more router hops removed from the client. Observe and compare the behavior of the pings while you force a failover by disabling a device or a link.


Based on the observed differences between the ping responses, you can draw conclusions about the likelihood that the problem is related to the FHRP or to any of the other protocols that are involved in the convergence. Here are a few examples:

- If you observe that the pings to the real IP address of the redundant router and the virtual IP address of the FHRP both fail at the same time and resume at the same time when you disable the primary router, it is safe to assume that the problem is not related to the FHRP (because the FHRP does not affect the pings to the real IP address). The most likely cause in this scenario is the Layer 2 convergence for the VLAN, so you should start a Layer 2 troubleshooting procedure.
- If you observe that the pings to the real IP address of the redundant router do not suffer any packet loss, but pings to the virtual IP address fail, this result strongly suggests that there is a problem with the FHRP.
- If you observe that the pings to the real IP address of the redundant router and to the virtual IP address do not suffer packet loss, but the ping to the host further out in the network fails, this result may indicate an issue with the routing protocol. (Alternatively, it could indicate that the client is using the primary router address as its default gateway rather than the virtual IP address.)

There are too many possible scenarios, combinations of ping results, and conclusions to list, but, in any scenario, you can gain important clues by comparing the differences between several pings during a failover.

If you have to troubleshoot without the opportunity to force failover for testing purposes, you may need to simply assume that the FHRP is the cause of the problem and carefully verify its implementation and operation, even if you cannot determine beforehand if this protocol might be the cause of the problem.

Sample First-Hop Redundancy Troubleshooting Flow




© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0-40

Before you even start to troubleshoot the FHRP itself, you should verify if the client is correctly using the virtual IP address and MAC address of the FHRP as its default gateway. This process involves verifying the default gateway configuration (whether statically configured or learned via DHCP) and the ARP cache on the client, to verify that both the virtual IP address and the virtual MAC address on the client match the expected values for the FHRP that is in use.

Sample First-Hop Redundancy Troubleshooting Flow



© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0-41

Many problems with first hop redundancy protocols are caused by underlying problems in the Layer 3 connectivity between the routers. Therefore, a good next-step in the troubleshooting process is to verify that there is Layer 3 connectivity between all routers that are participating in the first hop redundancy protocol. Ping from each of the participating routers to the IP addresses of the other participating routers. If one of these pings fails, you should start a troubleshooting process to diagnose and resolve the Layer 3 connectivity issues between the routers before further investigating the FHRP.

When you have confirmed that there is Layer 3 connectivity between the participating routers in general, you need to verify the proper transmission and reception of FHRP packets. To limit potential disruption, you should always use **show** commands to gather information before you consider using **debug** commands.

Verify Reception of HSRP Messages

The **show standby brief** command lists the active and standby router.

```
CSW1#show standby vlan 112 brief
P indicates configured to preempt.
Interface  Grp  Pri P State  Active      Standby      Virtual IP
Vl112      112  100 P Standby 10.1.152.253 local        10.1.152.254
```

```
CSW2#show standby vlan 112 brief
P indicates configured to preempt.
Interface  Grp  Pri P State  Active      Standby      Virtual IP
Vl112      112  110 P Active  local       10.1.152.252 10.1.152.254
```

Execute the command on all routers in the group to confirm consistency between the parameters as well as proper reception of HSRP messages.

© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0-42

This example shows how to confirm proper transmission and reception of HSRP messages. For GLBP or VRRP, the procedure is similar although the command output is slightly different.

To confirm the proper reception of HSRP messages on all routers in the group, you should verify that all routers list an active and a standby router and that these roles are listed in a consistent way across all the routers. The **show standby brief** command is concise and still shows the most relevant information. As you can see in the example, switch CSW1 lists the IP address of switch CSW2 as the active router, and as the standby router, it lists “local” to indicate that it considers itself the standby router. On switch CSW2, the situation is the exact opposite: The address of switch CSW1 is listed as the standby address, while the active router is listed as “local.” While you are verifying these roles, you can also use this opportunity to confirm that both the standby group number and the virtual IP address are configured in a consistent manner. Misconfiguration of these parameters is a common cause of HSRP problems.

Verify Reception of FHRP Messages (Cont.)

- The **debug standby packets** command displays all HSRP packets that are being transmitted and received in real time.
- A good way to limit the impact of this command is to log to buffers in RAM instead of to a terminal. Messages in the log can be filtered using regular expressions.

```
CSW1#debug standby packets
HSRP Packets debugging is on
CSW1#show logging | include Grp 112
Mar 30 03:31:22.920 PDT: HSRP: V1112 Grp 112 Hello out 10.1.152.252 Standby
pri 100 vIP 10.1.152.254
Mar 30 03:31:25.202 PDT: HSRP: V1112 Grp 112 Hello in 10.1.152.253 Active
pri 110 vIP 10.1.152.254
Mar 30 03:31:25.789 PDT: HSRP: V1112 Grp 112 Hello out 10.1.152.252 Standby
pri 100 vIP 10.1.152.254
Mar 30 03:31:27.861 PDT: HSRP: V1112 Grp 112 Hello in 10.1.152.253 Active
pri 110 vIP 10.1.152.254
```

© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0-43

Inconsistencies in the output of the **show standby brief** commands, such as a missing standby router on one of the routers or multiple routers claiming the active or standby router for a group, strongly suggests that there is a problem with the reception or interpretation of the HSRP messages on the routers. You can now use a **debug** command to investigate the transmission and reception of HSRP messages in order to gather more clues about the failure.


Before enabling a debug, you should first verify that the CPU of the device is not running at such high levels that adding the load of a debug would risk overloading the CPU. Secondly, it is always good to have a fallback plan to stop the debug when it unexpectedly starts to affect the performance of the device. For instance, you could open a second connection to the device and before you enable the debug in your primary session, type the **undebug all** command in the secondary session, but do not press the Enter key to confirm the command, yet. Another fallback scenario you could follow is to schedule a timed reload within a short time by using the **reload in** command. If you lose your connection to the device because of your debug, you can be assured that it will reload shortly and you will be able to reconnect to it. Finally, you should always refer to the policies of your organization before executing any commands on a device that put the operation of the network at risk.

The **debug standby packets** command displays all HSRP packets sent or received by the device. This command can quickly generate a lot of output, especially if you have configured many different HSRP groups or if you have tuned the hello timer to be shorter than the default value of three seconds. To make it easier to select the packets that you are interested in, you could use the technique shown in the figure. Instead of logging the debug output to the console or virtual terminal session, you can capture the output in a buffer in the device's RAM and then display the content of the buffer by using the **show logging** command. The output of the command can then be filtered by using a regular expression to select the HSRP group that you are interested in.

In the example in the figure, the output reveals that hellos are sent by this router and received from the other router. Just like the **show** commands in the previous figure, you should execute the **debug** command on both routers to spot possible differences in behavior between the devices.

Do not forget to disable the debug by using the **no debug** command after you have gathered the information that you were interested in.

If these debugs reveal that HSRP protocol packets are not properly received on any of the routers, check to see if access lists are blocking the packets. Given that you have already verified the Layer 3 connectivity between the devices, this problem should be on a higher layer.



When you have established that FHRP messages are sent and received properly on all routers and still the FHRP does not perform as expected, the problem must be related to the role selection and transferring roles between routers during failover. You may need to verify two potential problem areas.

If the FHRP is using authentication and a mismatch between the authentication parameters exists, then the devices will not accept each other's messages as valid messages when they are received. A typical symptom of this situation is that there will be more than one router that considers itself to be the active router for a group.

For all FHRPs, role selection is influenced by two parameters: priority and preemption. Tracking objects such as interfaces and routes can further alter these priorities. If an unexpected router is selected for the primary role at any point in the process, you should carefully analyze the priorities configured on the different devices and determine how they are affected by potential tracking options. However, to determine properly how properties behave during a failover, you will need to be able to force a failover, which means that you may need to postpone this type of testing until a regularly scheduled maintenance interval.

Lab Debrief Notes

Use these notes sections to write down the primary learning points that are discussed during the Lab Debrief.

Lab 4-2: Alternate Solutions

[illegible]

Lab 4-2: Alternate Methods and Processes

[illegible]

[illegible]

Lab 4-2: Important Commands and Tools

[illegible]

Lab 4-2: References

If you need more information on the commands and their options, you can go to the following sections of <http://www.cisco.com>:

- Cisco Systems, Inc. Command References for Cisco Catalyst LAN Switches: Go to Product Support (<http://www.cisco.com/web/psa/products/index.html>), select Switches, select LAN Switches and then the product family that you are working with. The Command References can then be found under the “Reference Guides” section.
- Cisco Systems, Inc. Virtual LANs/VLAN Trunking Protocol (VLANs/VTP) Troubleshooting TechNotes:
http://www.cisco.com/en/US/tech/tk389/tk689/tsd_technology_support_troubleshooting_technotes_list.html
- Cisco Systems, Inc. Layer-Three Switching and Forwarding Troubleshooting TechNotes:
http://www.cisco.com/en/US/tech/tk389/tk815/tsd_technology_support_troubleshooting_technotes_list.html

Lab 5-1: Layer 3 Connectivity and EIGRP

Complete this lab activity to practice what you learned in the related module.

Activity Objective


In this activity, you will troubleshoot various problems related to Layer 3 connectivity in general and routing problems related to the EIGRP. After completing this activity, you will be able to meet these objectives:

- Diagnose and resolve problems related to network layer connectivity
- Diagnose and resolve problems related to the EIGRP routing protocol
- Document troubleshooting progress, configuration changes, and problem resolution

Information Packet

The figure illustrates what you will accomplish in this activity:

Visual Objective



Required Resources

These are the resources and equipment that are required to complete this activity:

- One PC with Internet access per team member
- Four Cisco Catalyst 3560 Series Switches per team
- Six Cisco 1841 Integrated Services Routers per team

- Three client PCs per team
- One server per team

Job Aids

These job aids are available to help you complete the lab activity.

- Trouble tickets
- Troubleshooting log

Trouble Ticket H: Preparation for CCTV Pilot

Your company is interested in implementing an IP-based closed-circuit television (CCTV) solution. Currently, different solutions and vendors are being evaluated and one of the vendors has offered to implement a small pilot to show the capabilities of their solution. Although most of the video will be stored locally, there needs to be some communication between the central server at headquarters and the servers at the branch locations. To keep the traffic associated with the CCTV solution separate from the rest of the traffic, the CCTV solution will be implemented using two new VLANs, one at headquarters (VLAN 115 and subnet 10.1.155.0/24) and one at the branch office (VLAN 29 and subnet 10.1.163.64/26). Tomorrow the vendor will come in to install his systems and the network team has been asked to ensure that the new VLANs have been implemented and that there is IP connectivity between the headquarters CCTV VLAN and the branch CCTV VLAN. Your team has been very busy lately, so this step was not done until yesterday. Yesterday afternoon, one of your colleagues implemented the VLANs while handling various other tasks, but did not have time to test the implementation. You were asked to verify his implementation.

Your task is to verify the implementation and ensure that there is IP connectivity between the two CCTV VLANs when the vendor comes in to implement the CCTV solution tomorrow.

Note	You are allowed to assign PCs CLT1 and CLT3 to the CCTV VLANs at the headquarters and branch sites for testing purposes. Make sure that you reassign the PCs to their original VLANs and verify proper operation after you have finished your tests.

Trouble Ticket I: Fire in the Server Room

Before starting to work on your assigned tasks, you first have to drive to one of the nearby branch offices to pick up some equipment that was delivered to the wrong office and is needed this afternoon. After fifteen minutes, you get an urgent phone call: You should return to the office immediately. A short circuit has caused a small fire in the server room and both routers CRO1 and CRO2, which were mounted in the same rack, were damaged. Luckily, you had two cold spares in storage. When you arrive at the office, two of your colleagues have already installed the two replacement routers, cabled them, and tried to configure the routers. However, the routers are not operational yet when you come in.

You receive a number of phone calls from network administrators who work in the branch offices asking about the loss of the WAN. Some of them have started to troubleshoot by themselves. You tell them what happened and ask them not to do anything until you have resolved the problem at the central site.

Your task is to work together with your colleagues on restoring routers CRO1 and CRO2 and regaining connectivity across the WAN.

Note	Because of the fire, you have also lost the OOB management connection to the consoles at the branch office. Therefore, the consoles of BRO1, BRO2, and BSW1 cannot be used during this exercise. This issue is not a problem that needs to be solved, but a condition that you will have to work around.

Trouble Ticket J: User in Branch Cannot Access the Internet

While you were on the road, just before the fire started, a user in the office LAN in Branch 1 (who uses client PC CLT2) complained that he did not have Internet access. When he tried to open the website <http://www.isp3.local> (which corresponds to IP address 172.34.224.1), he received an error message from his browser saying that it cannot display the web page. He can reach the internal server SRV1 without any problems. You know that there were some problems with Internet access yesterday evening, but your colleague who worked on the problem has called in sick today and the logs do not show any useful information.

Your task is to diagnose and solve this problem and make sure that the user regains connectivity to the Internet.

Instructions

Together with your team members, create a troubleshooting plan to divide the work, assign each team member appropriate roles, and coordinate device access between team members. Together, work on Trouble Tickets H, I, and J to resolve the issues. Document your progress in the following Troubleshooting Logs in order to help facilitate efficient communication within the team and to have an overview of your troubleshooting process for reference during the Lab Debrief discussions.

You are allowed a total of one and a half hours to complete as many of the trouble tickets as you can. After this amount of time has passed, the instructor will debrief the lab and review all trouble tickets and their solutions. The main objective for the troubleshooting labs in this course is to give you an opportunity to practice structured troubleshooting. Fixing the problems is secondary to practicing proper processes and procedures.

Lab Setup

The instructor will provide you with directions to prepare the lab equipment for this lab. After the instructor indicates that the lab is fully prepared, you are ready to start troubleshooting.

Troubleshooting Log

Use this log to document your actions and results during the troubleshooting process.

Trouble Ticket H

Your task is to verify your colleague's implementation and ensure that there is IP connectivity between the two CCTV VLANs when the vendor comes in to implement the CCTV solution tomorrow.

Note Refer to the Activity Verification items at the end of this log to verify that you have successfully completed this task.

[illegible]

Troubleshooting Log

Use this log to document your actions and results during the troubleshooting process.

Trouble Ticket I

Your task is to work together with your colleagues on restoring routers CRO1 and CRO2 and regaining connectivity across the WAN.

Note Refer to the Activity Verification items at the end of this log to verify that you have successfully completed this task.

[illegible]

Troubleshooting Log

Use this log to document your actions and results during the troubleshooting process.

Trouble Ticket J

Your task is to diagnose and solve the connectivity problem experienced on client PC CLT2 and make sure that the user regains connectivity to the Internet.


Note Refer to the Activity Verification items at the end of this log to verify that you have successfully completed this task.

[illegible]

Lab 5-1: Sample Troubleshooting Flows

Troubleshooting IP connectivity

The figure illustrates an example of a method that you could follow to diagnose and resolve problems related to IP connectivity.



Layer 3 is a common starting point for many troubleshooting procedures. An often-applied method is the divide-and-conquer approach: When a user reports a problem concerning connectivity to a certain service or application running on a server, a good first step is to determine if there is end-to-end IP connectivity between the client and the server. If this connectivity does exist, you can focus on the higher layers of the OSI reference model.

You can confirm end-to-end IP connectivity by using the **ping** or **traceroute** commands. The exact syntax of these commands may be slightly different for different operating systems, but almost every operating system supports these commands in some form.

A prerequisite to using this method is that the appropriate ICMP messages are allowed on the network and not blocked by any firewalls, including host-based firewalls on the destination host. If you cannot use ping and traceroute effectively, you may have to resort to analyzing traffic captures of the actual traffic flows to determine if packets can be sent at the network layer between the affected hosts.

Use the Correct Source Address

When using a **ping** or **traceroute** command from the first-hop router instead of a client, specify the correct source for the command.

```
BR01#ping 10.1.152.1 source FastEthernet 0/1.19
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.152.1, timeout is 2 seconds:
Packet sent with a source address of 10.1.160.252
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 12/15/16 ms
```

```
BR01#traceroute 10.1.152.1 source FastEthernet 0/1.19
Type escape sequence to abort.
Tracing the route to srv1.tshoot.local (10.1.152.1)


 0 10.1.193.1 4 msec 4 msec 8 msec
 1 10.1.192.1 8 msec 8 msec 8 msec
 2 srv1.tshoot.local (10.1.152.1) 8 msec 8 msec 4 msec
```

© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0-47

It is important for you to realize that a successful **ping** or **traceroute** response is dependent on two things: The availability of a route to the destination and a route back to the source. You have to make sure that you specify the source address of the **ping** or **traceroute**, particularly when you run tests from the first-hop router in the path. If you do not specify the source address, the router will use the IP address of the egress interface as the source for the packets. Using an address from a different source subnet than the client may lead you to reach wrong conclusions if the problem concerns the return path for the packets.

Sample Layer 3 Troubleshooting Flow



© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0-48

When you have determined that there is a problem with the end-to-end IP connectivity between the affected hosts, you need to reduce the potential scope of the problem and isolate the point or points in the path between the hosts where the connectivity is lost.

A commonly used method is to track the path of the packets. You can use this method to diagnose end-to-end IP connectivity problems:

- Determine the Layer 3 path. Based on documentation, baselines, and knowledge of your network in general, the first step you should take is to determine the path that you would expect packets to follow between the affected hosts. Determining the expected traffic path beforehand will help you in two ways: It will give you a starting point for gathering information about what is actually happening on the network and it will make it easier to find abnormal behavior. The second step in determining the Layer 3 path is to follow the expected path and verify that the links on the expected path are actually up and forwarding traffic. If the actual traffic path is different from your expected path, this step may give you clues about the particular links or protocols that are failing and the help you determine the cause of these failures.
- To track the path of the packets between the hosts, you should first track the path that is being used according to the control plane information: Start at the client and verify the IP address, subnet mask, and default gateway. Then go to the router that is listed as the default gateway and see which route is used for the destination IP address. Determine the next-hop router based on the information in the routing table. Connect to the next hop router and repeat this procedure until you arrive at the router that is directly connected to the destination host. Then repeat the process for the route back from the destination to the source.
- If at any point during this procedure you find that the router has no route in the table for the destination network, you need to diagnose the process that is the source of the routing information on this router, such as the routing protocol or static routes.
- If you have verified that routing information is present on the complete path from the source to the destination and from the destination back to the source, but connectivity is failing, then you will again have to track the path, but this time determine at which point packets are being dropped. The likely causes for the packets to be dropped are Layer 1 problems, Layer 2 problems, or Layer 3 to Layer 2 mapping problems. When you have determined the point at which the packets are dropped, you need to use the specific troubleshooting methods appropriate for the Layer 2 technology that is used on the egress interface.

These steps do not necessarily have to be taken in the order presented here. Often, different aspects of this generic procedure are combined and shortcuts may be taken based on the result. For instance, determining proper packet forwarding will often be done in parallel with the determination of the routes by using **ping** to verify the reachability of the next-hop derived from the route or using **ping** and **tracert** to the final destination from intermediate routers in the path. If you find that a ping is successful from a particular point in the path, you know that routes to the destination must be available on all the downstream routers and you can use **tracert** to determine the path to the destination instead of connecting to each router in the path. However, be aware that this method has a hidden assumption, which is that packets traveling to the same destination use the same path, regardless of their source. This assumption is not necessarily the case in a redundant network with equal cost paths to a certain destination. The source address is typically used as part of the load-balancing algorithm that determines the path used when equal cost paths are available. It is important to determine the exact path for the actual source and destination IP address pair that is affected, especially in those cases where control plane information is available in both directions but packets are dropped.

Verify the Routing Table

- The **show ip route *ip-address*** command displays the best match in the routing table for a specific destination IP address.

```
CRO1#show ip route 10.1.152.1
Routing entry for 10.1.152.0/24
  Known via "eigrp 1", distance 90, metric 28416, type internal
  Redistributing via eigrp 1
  Last update from 10.1.192.9 on FastEthernet0/1, 22:18:43 ago
  Routing Descriptor Blocks:
    10.1.192.9, from 10.1.192.9, 22:18:43 ago, via FastEthernet0/1
      Route metric is 28416, traffic share count is 1
      Total delay is 110 microseconds, minimum bandwidth is 100000 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 1
    * 10.1.192.1, from 10.1.192.1, 22:18:43 ago, via FastEthernet0/0
      Route metric is 28416, traffic share count is 1
      Total delay is 110 microseconds, minimum bandwidth is 100000 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 1
```

- **Note:** If the default route is the best match in the routing table for the destination IP address, the message “% Network not in table” will be displayed, even if the packets are forwarded.

© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0-49

When you are troubleshooting IP connectivity to a specific destination IP address, you can use the **show ip route *ip-address*** command to determine the best prefix match for the IP address, the egress interface and, for multipoint interfaces, the next-hop IP address. If multiple equal cost paths are present, as can be seen in the example in the figure, each of those entries will be listed.

In addition, the routing source will be listed, such as directly connected, static, or the routing protocol. Additional control plane parameters that are associated with the route source, such as the administrative distance, routing protocol metrics, source router, and route age are also displayed. To interpret these additional parameters, you need more detailed knowledge of the specific routing protocol, and often, information that is more detailed can be gathered from the data structures of that specific protocol.

This command will never display the default route, 0.0.0.0/0 as a match, even if it is the longest prefix match for a packet. Therefore, if this command displays the message “% Network not in table” you cannot conclude that packets will be dropped, but you need to verify whether a default route is present by using the command **show ip route 0.0.0.0 0.0.0.0**.

Verify the Cisco Express Forwarding FIB

- The **show ip cef *ip-address*** command displays the best match in the FIB for a specific destination IP address.

```
CRO1#show ip cef 10.1.152.1
10.1.152.0/24, version 344, epoch 0, per-destination sharing
0 packets, 0 bytes
Flow: AS 0, mask 24
via 10.1.192.9, FastEthernet0/1, 0 dependencies
traffic share 1
next hop 10.1.192.9, FastEthernet0/1
valid adjacency
via 10.1.192.1, FastEthernet0/0, 0 dependencies
traffic share 1
next hop 10.1.192.1, FastEthernet0/0
valid adjacency
0 packets, 0 bytes switched through the prefix
tmstats: external 0 packets, 0 bytes
internal 0 packets, 0 bytes
```

- **Note:** This command displays the default route 0.0.0.0/0 if it is the best match.

© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0-50

To see the best match for a specific IP address in the Cisco Express Forwarding FIB, use the command **show ip cef *ip-address***. This command lists the same forwarding information as the **show ip route** command but without the associated control plane information, such as routing protocol metrics, administrative distance, and so on. This command displays the default route 0.0.0.0/0 if it is the best match for the destination IP address.

If the routing table for a route contains multiple entries, these same entries will also be present in the FIB.

Verify the Cisco Express Forwarding FIB (Cont.)

- In cases where the routing table and FIB list multiple entries for the same destination prefix, the **show ip cef exact-route *source-ip-address destination-ip-address*** command displays which of those entries is used for a specific source and destination IP address pair.

```
CRO1#show ip cef exact-route 10.1.160.65 10.1.152.1
10.1.160.65 -> 10.1.152.1 : FastEthernet0/0 (next hop 10.1.192.1)
```

© 2009 Cisco Systems, Inc. All rights reserved.


TSHOOT v1.0-51

When you trace the packet flow between two specific hosts and the routing table and FIB list multiple entries (because there are multiple equal cost paths), you need to determine which of those entries is used to forward the packets associated with the specific source and destination IP address pair that you are troubleshooting. The **show ip cef exact-route** command can be used in these situations to determine the specific egress interface and next-hop IP address for the specific IP address pair.

On multilayer switches, instead of consulting the FIB that is stored in the main memory of the switch, you have to consult the forwarding information stored in the hardware TCAM, because packet forwarding is handled by the TCAM, not the Cisco Express Forwarding FIB.

Although the FIB is used to compile the information that is loaded into the TCAM, the load-balancing algorithms that are used are different and do not necessarily yield the same result.

To learn more about the commands that can be used to verify the Layer 3 forwarding information contained in the TCAM, consult the multilayer switching sections of the Student Guide for this course and this Lab Guide.




When you find a point in the network where no route is present in the routing table for the destination IP address of the session (or when you are analyzing the return path for the source IP address of the session and you find that no route is present for the source address), you need to investigate what caused that route not to be installed in the routing table.

To diagnose correctly why a particular route is missing from the routing table, you first need to consult your documentation and baselines to find out what the expected routing source for this route would be. Is static routing used on this router or a routing protocol?

If a static route has been configured, but it is not listed in the routing table, you need to verify the status of the associated egress interface. If the egress interface for a static route is down, the route will not be installed in the routing table. If the route is not configured with an egress interface, but with a next-hop IP address, the same rule applies. The router will execute a recursive routing table lookup on the next hop for the static route. If no matching route and

associated egress interface can be found for the configured next-hop IP address of the static route, the route will not be installed in the routing table. If a match is found for the next-hop IP address, the static route will be installed in the routing table.

For dynamic routing protocols, you need to initiate a troubleshooting process that is appropriate for that specific protocol and try to determine why the route was not learned on this router or, if it was learned, why it is not used.



When you have verified the presence of correct routing information along the paths in both directions, but you find that packets are dropped at a certain hop in the path, you need to diagnose the packet forwarding process.

If a route is present in the routing table (and FIB if Cisco Express Forwarding is used), but packets are not forwarded correctly, you should verify whether there is a correct mapping between the IP next hop and the Layer 2 protocol that is used on the egress interface. If the router cannot find all the Layer 2 information that is needed to construct a frame to encapsulate a packet, then the packet will be dropped even if the routing information is present in the routing table.

Which command you should use to verify the Layer 3 to Layer 2 protocol mapping depends on the Layer 2 technology used on the egress interface. Examples are the **show ip arp** command for Ethernet networks and the **show frame-relay map** command for Frame Relay.

For more information about the exact command syntax, research the Layer 2 technology used in the configuration guides and command references on <http://www.cisco.com>.

If you find incorrect mappings or if you find the mappings to be correct, but frames are not forwarded correctly, you should initiate a Layer 2 troubleshooting procedure for the Layer 2 technology that is being used.

Verify the Adjacency Table

- The **show adjacency interface detail** command displays the adjacent neighbors or hosts that are reachable via the specified interface and the Layer 2 frame header that will be used to encapsulate packets that are forwarded through that adjacency.

```
CRO1#show adjacency FastEthernet 0/0 detail
Protocol Interface Address
IP          FastEthernet0/0 10.1.192.1 (277)
                    5021 packets, 1062633 bytes
                    001EF7BBF7C20019562C8FB40800
                    ARP      04:00:42
                    Epoch: 0
```

© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0--54

Regardless of the Layer 2 technology, if Cisco Express Forwarding is used as the Layer 3 forwarding method, you can verify the availability of Layer 2 forwarding information using the **show adjacency detail** command.

As can be seen in the example in the figure, this command lists the Layer 2 frame header that is used to encapsulate packets transmitted via the listed adjacency. In this example, the frame header is 001EF7BBF7C20019562C8FB40800, which can be dissected as follows:


- **001EF7BBF7C2**: This address is the destination MAC address of the frame, which corresponds to the MAC address of the next hop 10.1.192.1.
- **0019562C8FB4**: This address is the source MAC address of the frame, which corresponds to the MAC address of interface FastEthernet 0/0.
- **0800**: This value is the Ethernet type field, which indicates that the frame contains an IP packet, because Ethernet type value 0x800 is registered as the value for IP.

If you are troubleshooting a Layer 3 forwarding problem and the IP next hop and interface that are listed in the routing table are not present in the adjacency table, you know that there is a problem with the Layer 3 to Layer 2 mapping mechanisms.

If a Layer 2 frame header is listed in the adjacency table, but the frames are not forwarded correctly across the Layer 2 medium, you will have to troubleshoot the underlying Layer 2 technology. The information contained in the header can be useful information when you start the Layer 2 troubleshooting process.

Troubleshooting EIGRP

The figure illustrates an example of a method that you could follow to diagnose and resolve problems related to the EIGRP.



A common circumstance that may require you to investigate the routing protocol operation is when you are troubleshooting IP connectivity to a particular destination and you find that the route to the destination network is missing from the routing table of one of the routers or that a different route than expected was selected to forward the packets to that destination.

In order to install a route into the routing table, each router that uses a routing protocol, goes through several stages:


- Discovers neighbors and establish a neighbor relationship
- Exchanges routing information with neighbors and store the received information in protocol-specific data structures
- Selects the best route from the available routes and install it in the routing table

Errors during any of these stages can lead to missing routing information or to the wrong routing information being installed in the routing table.

The exact processes that take place, the data structures that are used, and the commands that are used to gather information about these processes and data structures are protocol-specific, but the generic troubleshooting principles are similar for all routing protocols.

The order of verification of the different stages of this process is not important, as long as a structured approach is used.

Sample EIGRP Troubleshooting Flow



© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0—56

EIGRP uses hello packets to discover and maintain neighbor relationships. Neighbors that are discovered are registered in the EIGRP neighbor table remain in the neighbor table as long as hello packets are received. A neighbor will be removed from the table when its hold time expires or when the interface on which the neighbor is registered goes down. The default EIGRP hello timer is 5 seconds for these interfaces:

- High-speed multipoint interfaces, such as Ethernet interfaces.
- Point-to-point interfaces such as the following:
 - Serial interfaces running PPP or HDLC.
 - Point-to-point Frame Relay subinterfaces
 - Point-to-point ATM subinterfaces.

The default hold time for these interfaces is 15 seconds. Each router advertises hello and hold timers that it uses in its hellos. Although it is recommended that the timers be changed in a consistent manner on all routers if the timers need to be tuned, the timers do not need to match between two routers to allow them to become neighbors.

Verify the EIGRP Interface Table

- The **show ip eigrp interfaces** command displays all interfaces that have been enabled for EIGRP processing.
- If an interface is not listed in this command, then one of these items is true:
 - The IP address of the interface is not covered by any of the EIGRP **network** statements.
 - The interface was configured as a passive interface.

```
CR01#show ip eigrp interfaces
IP-EIGRP interfaces for process 1
```

Interface	Peers	Xmit Queue Un/Reliable	Mean SRTT	Pacing Time Un/Reliable	Multicast Flow Timer	Pending Routes
Fa0/0	1	0/0	85	0/1	832	0
Fa0/1	1	0/0	70	0/1	540	0
Se0/0/1	1	0/0	223	5/190	910	0
Se0/0/0.121	1	0/0	502	10/380	2492	0
Se0/0/0.122	1	0/0	306	10/380	2252	0
Lo0	0	0/0	0	0/1	0	0

© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0-57

Neighbors can only be discovered on an interface that is operational and has been activated for EIGRP processing. An interface will be activated for EIGRP packet processing if the IP address of the interface is covered by one of the **network** statements that is configured under the **router eigrp** process and the interface is not configured as a passive interface. You can use the **show ip eigrp interface** command to display the list of EIGRP interfaces. An interface does not need to be operational to be listed in the output of this command. You need to verify the operational status of the interface by using the **show interfaces**, **show interface status** or **show ip interfaces brief** commands.

If you find that an interface is not listed in the output of the **show ip eigrp interfaces** command as expected, you should verify the **network** and **passive-interface** commands under the **router eigrp** configuration.

Verify the EIGRP Neighbor Table

- The **show ip eigrp neighbors** command displays all neighbors that were discovered on active EIGRP interfaces.

```
CRO1#show ip eigrp neighbors
IP-EIGRP neighbors for process 1
```

H	Address	Interface	Hold (sec)	Uptime	SRTT (ms)	RTO	Q Cnt	Seq Num
0	10.1.192.1	Fa0/0	14	00:06:26	3	200	0	828
4	10.1.194.6	Se0/0/0.122	11	1d03h	104	2280	0	613
3	10.1.194.2	Se0/0/0.121	12	1d03h	117	2280	0	474
2	10.1.193.2	Se0/0/1	14	1d03h	49	1140	0	473
1	10.1.192.9	Fa0/1	14	1d03h	9	200	0	2069

Hold:
Based on a 15-second hold time and 5-second hellos, this value should be between 10 and 15 seconds.

Uptime:
This column indicates how long the neighbor has continuously been up.

© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0--58

You can use the **show ip eigrp neighbors** command to display the EIGRP table, from which you can verify that all expected neighbor relationships are operational.

The two most relevant columns in this output for troubleshooting purposes are the “Hold” column, which lists the number of seconds that will pass before a neighbor will expire from the table and the “Uptime” column, which lists how long this neighbor has been operational since it was last discovered. These two parameters can give you a good indication of the stability of the neighbor relationship. The uptime tells you how long the neighbor relationship has been successfully maintained, while displaying the hold time several times in a row can tell you if hellos are being received in a timely fashion. Based on the default 5-second hello and 15-second hold time, the value in this column should be between 15 and 10 seconds, because it counts down and will be reset to the hold time any time a hello is received from the neighbor.

If you find that the uptime of a neighbor is shorter than expected, you should verify the logs for interface-related events or EIGRP neighbor-related events, such as the following:

```
Apr 13 06:25:01 PDT: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/11, changed state to down
```

```
Apr 13 06:25:02 PDT: %LINK-3-UPDOWN: Interface FastEthernet0/11, changed state to down
```

```
Apr 13 06:25:02 PDT: %DUAL-5-NBRCHANGE: EIGRP-IPv4:(1) 1: Neighbor 10.1.192.2 (FastEthernet0/11) is down: interface down
```

```
Apr 13 06:25:14 PDT: %DUAL-5-NBRCHANGE: EIGRP-IPv4:(1) 1: Neighbor 10.1.192.2 (FastEthernet0/11) is up: new adjacency
```

```
Apr 13 06:25:16 PDT: %LINK-3-UPDOWN: Interface FastEthernet0/11, changed state to up
```

```
Apr 13 06:25:17 PDT: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/11, changed state to up
```

Specifically the %DUAL-5-NBRCHANGE messages are very useful in troubleshooting, because they give you an indication of why the neighbor was lost. (In this case, it was caused by the interface going down.)

Debug EIGRP Packet Exchange

- The **debug eigrp packets** command can be used to verify the reception and transmission of EIGRP packets.
- The output of the debug can be limited by specifying the packet type or limiting output to a specific neighbor.

```
CRO1#debug ip eigrp 1 neighbor 10.1.192.1
First enable IP-EIGRP Route Events or EIGRP packet debug
CRO1#debug eigrp packets
EIGRP Packets debugging is on
  (UPDATE, REQUEST, QUERY, REPLY, HELLO, IPXSAP, PROBE, ACK, STUB, SIAQUERY,
  SIAREPLY)
CRO1#debug ip eigrp 1 neighbor 10.1.192.1
IP Neighbor target enabled on AS 1 for 10.1.192.1
IP-EIGRP Neighbor Target Events debugging is on
CRO1#
Apr 13 06:57:26.964 PDT: EIGRP: Received HELLO on FastEthernet0/0 nbr 10.1.192.1
Apr 13 06:57:26.964 PDT:   AS 1, Flags 0x0, Seq 0/0 idbQ 0/0 iidbQ un/rely 0/0
peerQ un/rely 0/0
Apr 13 06:57:27.528 PDT: EIGRP: Sending HELLO on FastEthernet0/0
Apr 13 06:57:27.528 PDT:   AS 1, Flags 0x0, Seq 0/0 idbQ 0/0 iidbQ un/rely 0/0
Apr 13 06:57:31.764 PDT: EIGRP: Received HELLO on FastEthernet0/0 nbr 10.1.192.1
Apr 13 06:57:31.764 PDT:   AS 1, Flags 0x0, Seq 0/0 idbQ 0/0 iidbQ un/rely 0/0
peerQ un/rely 0/0
```

© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0-59

If you do not see an expected neighbor listed in the neighbor table on a specific interface and you have confirmed that the interface is operational and is listed in the interface table, you can use the **debug** command **debug eigrp packets** to display the transmission and reception of EIGRP protocol packets in real time. This command can potentially generate a large amount of output and should be enabled with care.

You can limit the output of this command by specifying the packet type (update, request, query, reply, hello, ipxsap, probe, ack, stub, siaquery, or siareply). Additional conditions can be imposed using the **debug ip eigrp as-number** command, such as limiting the output to a specific neighbor or network.

To reduce the impact of the command further, it may be good to disable logging to the console and log to buffers in the router instead. You can then display the content of the log buffer using the **show logging** command. The following example shows you how to use this technique:

```
CRO1#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
CRO1(config)#no logging console
CRO1(config)#logging buffered 16384
CRO1(config)#^Z
CRO1#debug eigrp packets
EIGRP Packets debugging is on
  (UPDATE, REQUEST, QUERY, REPLY, HELLO, IPXSAP, PROBE, ACK, STUB, SIAQUERY,
  SIAREPLY)
CRO1#debug ip eigrp 1 neighbor 10.1.192.1
IP Neighbor target enabled on AS 1 for 10.1.192.1
IP-EIGRP Neighbor Target Events debugging is on
CRO1#clear logging
Clear logging buffer [confirm]
CRO1#show logging
Syslog logging: enabled (1 messages dropped, 108 messages rate-limited,
                  0 flushes, 0 overruns, xml disabled, filtering disabled)
```


Console logging: disabled
Monitor logging: level debugging, 0 messages logged, xml disabled,
filtering disabled
Buffer logging: level debugging, 13924 messages logged, xml disabled,
filtering disabled
Logging Exception size (4096 bytes)
Count and timestamp logging messages: disabled

No active filter modules.

Trap logging: level informational, 242 message lines logged
Logging to 10.1.152.1(global) (udp port 514, audit disabled, link
up), 242 message lines logged, xml disabled,
filtering disabled

Log Buffer (16384 bytes):

```
Apr 13 07:40:38.177 PDT: EIGRP: Received HELLO on FastEthernet0/0 nbr
10.1.192.1
Apr 13 07:40:38.177 PDT:   AS 1, Flags 0x0, Seq 0/0 idbQ 0/0 iidbQ un/rely 0/0
peerQ un/rely 0/0
Apr 13 07:40:42.517 PDT: EIGRP: Received HELLO on FastEthernet0/0 nbr
10.1.192.1
Apr 13 07:40:42.517 PDT:   AS 1, Flags 0x0, Seq 0/0 idbQ 0/0 iidbQ un/rely 0/0
peerQ un/rely 0/0
Apr 13 07:40:47.237 PDT: EIGRP: Received HELLO on FastEthernet0/0 nbr
10.1.192.1
Apr 13 07:40:47.237 PDT:   AS 1, Flags 0x0, Seq 0/0 idbQ 0/0 iidbQ un/rely 0/0
peerQ un/rely 0/0
CRO1#
```



After you have verified that neighbor relationships have been established as expected, you should verify that the route for the destination network that you are troubleshooting has been received correctly from all appropriate neighbors. EIGRP stores all routes that it receives from its neighbors in its topology table and then selects the best route from these routes to be installed in the routing table.

By investigating the available routes to the destination network in the topology table, you can see if all options that you expected were learned and if they have the correct associated metrics.

If routes are missing from the topology table, you may need to debug the EIGRP route exchange process to see if the problem is that they were not received or that they were not entered into the topology table.

Verify the EIGRP Topology Table

- The **show ip eigrp topology network mask** command displays all available routes present in the EIGRP topology table for a specific prefix. These routes include successors (best routes), feasible successors (backup routes), and possible successors (potential backup routes).

```
CRO1#show ip eigrp topology 10.1.152.0 255.255.255.0
IP-EIGRP (AS 1): Topology entry for 10.1.152.0/24
  State is Passive, Query origin flag is 1, 2 Successor(s), FD is 28416
  Routing Descriptor Blocks:
    10.1.192.1 (FastEthernet0/0), from 10.1.192.1, Send flag is 0x0
      Composite metric is (28416/2816), Route is Internal
      Vector metric:
        Minimum bandwidth is 100000 Kbit
        Total delay is 110 microseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 1
    10.1.192.9 (FastEthernet0/1), from 10.1.192.9, Send flag is 0x0
      Composite metric is (28416/2816), Route is Internal
      Vector metric:
        Minimum bandwidth is 100000 Kbit
<... further output omitted ...>
```

© 2009 Cisco Systems, Inc. All rights reserved.TSHOOT v1.0-61

The EIGRP topology table contains all routes that were received from all neighbors.

For each particular prefix, the following three types of entries may appear:

- **Successors:** These routes are the entries that were selected from the topology table as the best routes and then were installed in the routing table. For a route to be a successor route, it needs to meet the following criteria: Its metric needs to be the best possible routing metric from all the routes in the topology table for that prefix (this metric is also called the feasible distance). Secondly, it will only be marked as a successor if it was actually installed in the topology table. If a competing route for that prefix, such as a static route, was installed in the routing table instead (because it had a better administrative distance), the EIGRP route will not be marked as a successor in the topology table.
- **Feasible successors:** These routes have a metric that is higher than the feasible distance for the prefix, but meet the “feasibility condition.” The feasibility condition is met if the advertised distance of the route is lower than the feasible distance. This result means that the route is considered a backup route and, if the best route is lost, the feasible successor route can be used immediately without confirming the feasibility of the backup route through a query and reply process.
- **Possible successors:** These routes do not meet the feasibility condition. They are potential backup routes, but if the best route is lost, you must perform a query and reply process to confirm that they are valid backup routes.

As an example, the content of the EIGRP topology table for network 10.1.152.0/24 is listed here and comments are interspersed with the output to help interpret the entries.

```
CRO1#show ip eigrp topology 10.1.152.0 255.255.255.0
IP-EIGRP (AS 1): Topology entry for 10.1.152.0/24
  State is Passive, Query origin flag is 1, 2 Successor(s), FD is 28416
```

There are two successors for this prefix and the feasible distance is 28416.

```
Routing Descriptor Blocks:
 10.1.192.1 (FastEthernet0/0), from 10.1.192.1, Send flag is 0x0
   Composite metric is (28416/2816), Route is Internal
```

This entry is one of the two successors, because its distance of 28416 (the first number between the parentheses) is equal to the feasible distance of 28416.

```
Vector metric:
  Minimum bandwidth is 100000 Kbit
  Total delay is 110 microseconds
  Reliability is 255/255
  Load is 1/255
  Minimum MTU is 1500
  Hop count is 1
 10.1.192.9 (FastEthernet0/1), from 10.1.192.9, Send flag is 0x0
   Composite metric is (28416/2816), Route is Internal
```

This entry is the second successor, because its distance of 28416 is also equal to the feasible distance of 28416.

```
Vector metric:
  Minimum bandwidth is 100000 Kbit
  Total delay is 110 microseconds
  Reliability is 255/255
  Load is 1/255
  Minimum MTU is 1500
  Hop count is 1
 10.1.194.2 (Serial0/0/0.121), from 10.1.194.2, Send flag is 0x0
   Composite metric is (41026816/20514816), Route is Internal
```

This entry is not a successor, because its distance of 41026816 is higher than the feasible distance of 28416. It is not a feasible successor either, because its advertised distance (the second number between the parentheses) is not lower than the feasible distance of 28416. Therefore, it is a possible successor.

```
Vector metric:
  Minimum bandwidth is 64 Kbit
  Total delay is 40110 microseconds
  Reliability is 255/255
  Load is 1/255
  Minimum MTU is 1500
  Hop count is 3
 10.1.194.6 (Serial0/0/0.122), from 10.1.194.6, Send flag is 0x0
   Composite metric is (41026816/20514816), Route is Internal
```

This entry is not a successor, because its distance of 41026816 is higher than the feasible distance of 28416. It is not a feasible successor either, because its advertised distance is not lower than the feasible distance of 28416. Therefore, it is a possible successor.

```
Vector metric:
  Minimum bandwidth is 64 Kbit
  Total delay is 40110 microseconds
  Reliability is 255/255
  Load is 1/255
  Minimum MTU is 1500
  Hop count is 3
```

Debug EIGRP Route Exchange

- The **debug ip eigrp** command can be used to verify the processing of routing events such as updates, queries, and replies.
- **Use with extreme care:** This debug produces very large amounts of output even when restricted to specific neighbors or routes.

```
CRO1#debug ip eigrp
IP-EIGRP Route Events debugging is on
CRO1#debug ip eigrp 1 neighbor 10.1.192.1
IP Neighbor target enabled on AS 1 for 10.1.192.1
IP-EIGRP Neighbor Target Events debugging is on
CRO1#clear ip eigrp neighbors 10.1.192.1
CRO1#
Apr 13 08:44:39 PDT: %DUAL-5-NBRCHANGE: IP-EIGRP(0) 1: Neighbor 10.1.192.1
(FastEthernet0/0) i
Apr 13 08:44:39.076 PDT: IP-EIGRP(Default-IP-Routing-Table:1): 0.0.0.0/0 routing
table not updated thru 10.1.194.2
Apr 13 08:44:39.076 PDT: IP-EIGRP(Default-IP-Routing-Table:1): 0.0.0.0/0 routing
table not updated thru 10.1.194.6
Apr 13 08:44:39.076 PDT: IP-EIGRP(Default-IP-Routing-Table:1): route installed
for 10.1.130.128 ()
Apr 13 08:44:39.076 PDT: IP-EIGRP(Default-IP-Routing-Table:1): 10.1.130.128/26
routing table not updated thru 10.1.194.2
<... further output omitted ...>
```


© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0-62

If you find that expected route entries are missing from the topology table, you can consider using the **debug ip eigrp** command to display the processing of routing events by the router. However, this command can produce very large numbers of messages and as a result, has a high risk of disrupting the operation of the router. This debug command should not be used unless Cisco TAC tells you to use it or unless you are in a nonoperational network, such as a lab network that you built to reproduce a problem.

Like the **debug eigrp packets** command, you can limit the impact of this command by logging to buffers instead of the console and by limiting the output to specific neighbors or routes. Even then, extreme care should be taken when using this **debug** command.

Sample EIGRP Troubleshooting Flow



© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0-63

If you find that an EIGRP route for a specific destination network is available in the topology table, but a different route is present in the routing table, you should compare the value of the administrative distance of the route in the routing table to the value of the EIGRP route (which is 90 for internal routes and 170 for external routes by default). If the distance of the EIGRP route is higher than the distance of the competing route, the EIGRP route will not be installed in the routing table.

Lab Debrief Notes

Use these notes sections to write down the primary learning points that are discussed during the Lab Debrief.

Lab 5-1: Alternate Solutions

[illegible]

Lab 5-1: Alternate Methods and Processes

[illegible]

[illegible]

Lab 5-1: Important Commands and Tools

[illegible]

Lab 5-1: References

If you need more information on the commands and their options, you can go to the following sections of <http://www.cisco.com>.

- Cisco Systems, Inc. *Cisco IOS IP Routing Protocols Command Reference*. San Jose, California, November 2008:
http://www.cisco.com/en/US/docs/ios/iproute/command/reference/irp_book.html
- Cisco Systems, Inc. *Cisco IOS IP Switching Command Reference*. San Jose, California, November 2008:
http://www.cisco.com/en/US/docs/ios/ipswitch/command/reference/isw_book.html
- Cisco Systems, Inc. Enhanced Interior Gateway Routing Protocol Troubleshooting TechNotes:
http://www.cisco.com/en/US/tech/tk365/tsd_technology_support_troubleshooting_technotes_list.html#anchor3

Lab 5-2: OSPF and Route Redistribution

Complete this lab activity to practice what you learned in the related module.

Activity Objective


In this activity, you will troubleshoot various problems related to OSPF routing protocol and route redistribution between routing protocols. After completing this activity, you will be able to meet these objectives:

- Diagnose and resolve problems related to the OSPF routing protocol
- Diagnose and resolve problems related to route redistribution
- Document troubleshooting progress, configuration changes, and problem resolution

Information Packet

The figure illustrates what you will accomplish in this activity:

Visual Objective



Required Resources

These are the resources and equipment that are required to complete this activity:

- One PC with Internet access per team member
- Four Cisco Catalyst 3560 Series Switches per team
- Six Cisco 1841 Integrated Services Routers per team
- Three client PCs per team

- One server per team

Job Aids

These job aids are available to help you complete the lab activity.

- Trouble tickets
- Troubleshooting log


Introduction: Migration to OSPF

Your company has decided to migrate from using the EIGRP to using OSPF as the routing protocol. This migration is going to be executed in two phases. During the first phase, the headquarters campus will be migrated to OSPF. During this phase, EIGRP will still be used on the WAN towards the branch offices. On routers CRO1 and CRO2, redistribution will be configured between OSPF and EIGRP to ensure connectivity between headquarters and the branches. During the second phase, the branch offices will be migrated one by one until OSPF is used in the entire network.

The migration has been planned and designed by the engineering team, but the support team will have to support the new network and will be involved in migrating most of the branches during the second phase.

Today is Saturday and the engineering team has been busy implementing OSPF and removing EIGRP at the headquarters site. Although you have not taken part in the actual implementation, some of the senior engineers in the support team are on standby to assist during the verification and troubleshooting phase. Together with the engineering team, you will have to make the decision on Sunday to either accept the implementation or if major issues are uncovered that would threaten the stability of the network, roll back to the original configurations.


The OSPF design is outlined in the following figure:



The branch sites will be in separate areas that are configured as totally stubby areas.

To test both the branch connectivity using redistribution between EIGRP and OSPF and the eventual situation of only using OSPF, branch routers BRO1 and BRO2 have been specifically prepared for both these scenarios. The access VLANs at the Branch 1 site have been divided between routers BRO1 and BRO2 by disabling the corresponding interface on the other router. As a result, router BRO1 will function as the default gateway for VLANs 16 and 17, while router BRO2 will be the default gateway for VLANs 18, 19, and 128. Router BRO1 will run EIGRP as usual, while router BRO2 has been converted to run OSPF in area 11. This setup allows testing of the EIGRP redistribution from client PC CLT2 and testing of OSPF from client PC CLT3.

The branch office migration design is outlined in the following figure:



All of the following trouble tickets are related to the verification and acceptance of the first phase of the OSPF migration.

Note	Any interfaces that have been shut down on routers BRO1 and BRO2 should remain shut down for the duration of this lab exercise.

Trouble Ticket K: No Connectivity from Client PC CLT2

After the implementation of OSPF and the implementation of redistribution between EIGRP and OSPF, the connectivity from client PC CLT2, which uses router BRO1 as its default gateway to server SRV1 at headquarters, is tested. A ping from client PC CLT2 to server SRV1 fails. The connectivity problem is not limited to SRV1. An attempt to browse to <http://www.isp3.local> also fails.

Your task is to diagnose this problem and if possible, resolve it. Connectivity from PC CLT2 to server SRV1 and to server <http://www.isp3.local> on the Internet is mandatory for this phase of the migration to be considered successful.

Trouble Ticket L: No Connectivity from Client PC CLT3

After the implementation of OSPF, the connectivity from client PC CLT3, which uses router BRO2 as its default gateway, to server SRV1 at headquarters is tested. A ping from client PC CLT3 to server SRV1 fails. The connectivity problem is not limited to SRV1. An attempt to browse to <http://www.isp3.local> also fails.

Your task is to diagnose this problem and, if possible, resolve it. Connectivity from PC CLT3 to server SRV1 and to server <http://www.isp3.local> on the Internet is mandatory for this phase of the migration to be considered successful.

Trouble Ticket M: Internet not Reachable from Client PC CLT1

After the implementation of OSPF, the connectivity from client PC CLT1 to the Internet does not seem to be working. A ping from client PC CLT1 to server SRV1 succeeds, but an attempt to browse to <http://www.isp3.local> fails.

Your task is to diagnose this problem and, if possible, resolve it. Connectivity from PC CLT1 to server <http://www.isp3.local> on the Internet is mandatory for this phase of the migration to be considered successful.

Trouble Ticket N: OSPF Authentication Not Working

Yesterday one of the engineers suggested that it might be a good idea to secure the OSPF implementation by using MD5 authentication between the routers. Because action this could complicate the implementation, it was decided that it was too late to include this in the implementation for all areas now. However, to test the concept it was decided to enable the authentication for area 1 only to test the concept. If this test is successful, the authentication will be added to other areas during the second phase of the implementation. If the test is not successful, a separate project will be initiated to implement the authentication.

One of your colleagues has enabled MD5 authentication for area 1 on VLAN 111, which is used as a transit link between the core switches CSW1 and CSW2 in area 1. Unfortunately, the neighbor relationship between CSW1 and CSW2 on VLAN 111 is not established.

Your task is to diagnose this problem and, if possible, resolve it. The inability to resolve the authentication problem is not considered a reason to roll back the OSPF migration. You are allowed to remove the authentication for area 1 if necessary. However, in this case, you still need to make sure that the neighbor relationship between switches CSW1 and CSW2 on VLAN 111 in area 1 is established correctly.

Instructions

Together with your team members, create a troubleshooting plan to divide the work, assign each team member appropriate roles, and coordinate device access between team members. Together, work on Trouble Tickets K, L, M, and N to resolve the issues. Document your progress in the following Troubleshooting Logs in order to help facilitate efficient communication within the team and to have an overview of your troubleshooting process for reference during the Lab Debrief discussions.

You are allowed a total of two and a half hours to complete as many of the trouble tickets as you can. After two and a half hours, the instructor will debrief the lab and review all trouble tickets and their solutions. The main objective for the troubleshooting labs in this course is to give you an opportunity to practice structured troubleshooting. Your most important goal is to fix the problems that are introduced in the lab; practicing proper processes and procedures is a secondary goal.

Lab Setup

The instructor will provide you with directions to prepare the lab equipment for this lab. After the instructor indicates that the lab is fully prepared, you are ready to start troubleshooting.

Troubleshooting Log

Use this log to document your actions and results during the troubleshooting process.

Trouble Ticket K

Your task is to diagnose the connectivity problem from client PC CLT2 and if possible, resolve it. Connectivity from PC CLT2 to server SRV1 and to server <http://www.isp3.local> on the Internet is mandatory for this phase of the migration to be considered successful.

Note Refer to the Activity Verification items at the end of this log to verify that you have successfully completed this task.

[illegible]

Troubleshooting Log

Use this log to document your actions and results during the troubleshooting process.

Trouble Ticket L

Your task is to diagnose the connectivity problem from client PC CLT3 and, if possible, resolve it. Connectivity from PC CLT3 to server SRV1 and to server <http://www.isp3.local> on the Internet is mandatory for this phase of the migration to be considered successful.

Note Refer to the Activity Verification items at the end of this log to verify that you have successfully completed this task.

[illegible]

Troubleshooting Log

Use this log to document your actions and results during the troubleshooting process.

Trouble Ticket M

Your task is to diagnose the connectivity problem from client PC CLT1 and, if possible, resolve it. Connectivity from PC CLT1 to server <http://www.isp3.local> on the Internet is mandatory for this phase of the migration to be considered successful.

Note Refer to the Activity Verification items at the end of this log to verify that you have successfully completed this task.

[illegible]

Troubleshooting Log

Use this log to document your actions and results during the troubleshooting process.

Trouble Ticket N

Your task is to diagnose the OSPF authentication problem and, if possible, resolve it. The inability to resolve the authentication problem is not considered a reason to roll back the OSPF migration. You are allowed to remove the authentication for area 1 if necessary. However, in this case, you still need to make sure that the neighbor relationship between switches CSW1 and CSW2 on VLAN 111 in area 1 is established correctly.


Note Refer to the Activity Verification items at the end of this log to verify that you have successfully completed this task.

[illegible]

Lab 5-2: Sample Troubleshooting Flows

Troubleshooting the OSPF Routing Protocol

The figure illustrates an example of a method that you could follow to diagnose and resolve problems related to the OSPF routing protocol.



A common circumstance that may require you to investigate the routing protocol operation is when you are troubleshooting IP connectivity to a particular destination and you find that that the route to the destination network is missing from the routing table of one of the routers, or that a different route than expected was selected to forward the packets to that destination.

In order to install a route into the routing table, each router that uses a routing protocol, goes through several stages:


- Discovers neighbors and establish a neighbor relationship
- Exchanges routing information with neighbors and store the received information in protocol-specific data structures
- Selects the best route from the available routes and install it in the routing table

Errors during any of these stages can lead to missing routing information or to the wrong information being installed in the routing table.

The exact processes that take place, the data structures that are used, and the commands that are used to gather information about these processes and data structures are protocol specific, but the generic troubleshooting principles are similar for all routing protocols.

The order of verification of the different process stages is not important as long as a structured approach is used.

Sample OSPF Troubleshooting Flow



© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0-68

OSPF uses hello packets to establish and maintain neighbor relationships. Neighbors, from which a hello packet is received, are entered in the neighbor table. Subsequently OSPF establishes an adjacency by transitioning through several stages in which the link-state databases of the router are synchronized with its neighbor. After the completion of the database synchronization, the neighbors are considered fully adjacent and both link-state updates and user traffic can be passed between the neighbors. The neighbor remains registered in the neighbor table as long as hello packets are received regularly. A neighbor will be removed from the neighbor table when its dead time expires or when the interface on which the neighbor is registered goes down. The default OSPF hello timer is 10 seconds for point-to-point interfaces (such as serial interfaces running PPP or HDLC), point-to-point Frame Relay subinterfaces, and point-to-point ATM subinterfaces and for broadcast type interfaces such as Ethernet. The default dead time for these interfaces is 40 seconds. Each router advertises its hello and hold times in its hello packets and these values need to match for two routers to become neighbors.

Verify the OSPF Interface Table

- The **show ip ospf interface brief** command lists all interfaces that have been enabled for OSPF.
- If an interface is not listed in this command, this means that:
 - The IP address of the interface is not covered by any of the OSPF **network** statements.
 - OSPF has not been enabled directly on the interface by use of the **ip ospf process-id area area-id** command.

```
CR01#show ip ospf interface brief
```

Interface	PID	Area	IP Address/Mask	Cost	State	Nbrs	F/C
Lo0	100	0	10.1.220.1/32	1	LOOP	0/0	
Fa0/1	100	0	10.1.192.10/30	1	BDR	1/1	
Fa0/0	100	0	10.1.192.2/30	1	DOWN	0/0	

- Passive interfaces are listed in the output of this command.
- Interfaces that are down are listed in the output of this command with their state marked as “DOWN.”

© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0-68

Neighbors can only be discovered on an interface that has been enabled for OSPF and has not been configured as a passive interface. An interface can be enabled for OSPF in two different ways. An interface is enabled for OSPF if the IP address of the interface is covered by one of the **network** statements configured under the **router ospf** process, which assigns it to an area. Alternatively, an interface can be enabled for OSPF by an explicit **ip ospf process-id area area-id** command configured on the interface, which assigns the interface to an area. You can use the **show ip ospf interface brief** command to display the list of OSPF-enabled interfaces. This list includes interfaces that are down or that have been configured as passive interfaces. You can recognize interfaces that are down by the fact that their state is marked as “DOWN.” However, passive interfaces are not easily recognizable in the output of this command.

Verify the OSPF Interface Table (Cont.)

- The **show ip ospf interface *interface-id*** command displays the configured parameters and detailed information about the operational state of the specified interface.

```
CRO1#show ip ospf interface FastEthernet 0/0
FastEthernet0/0 is up, line protocol is up
  Internet Address 10.1.192.2/30, Area 0
  Process ID 100, Router ID 10.1.220.1, Network Type BROADCAST, Cost: 1
  Transmit Delay is 1 sec, State DR, Priority 1
  Designated Router (ID) 10.1.220.1, Interface address 10.1.192.2
  No backup designated router on this network
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    oob-resync timeout 40
  No Hellos (Passive interface)
  Supports Link-local Signaling (LLS)
  Index 1/1, flood queue length 0
  Next 0x0(0)/0x0(0)
  Last flood scan length is 0, maximum is 0
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 0, Adjacent neighbor count is 0
  Suppress hello for 0 neighbor(s)
```

© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0-70

Use the **show ip ospf interface *interface-id*** command to verify whether an interface is marked as a passive interface. Instead of a short list, this command displays comprehensive details of the OSPF parameters and the operational state for the specified interface. This command can also be used to verify timer values, such as the hello and dead timers, which could prevent a neighbor relationship from being established.

What does this situation mean from a troubleshooting standpoint?

If you find that an interface is not listed in the output of the **show ip ospf interface brief** command as you expected, you should verify the **network** commands under the **router ospf** configuration.

If you find that an interface is listed, but no neighbors are registered on the interface, you should issue the **show ip ospf interface *interface-id*** command for that interface to verify that the interface was not marked as passive.

Verify the OSPF Neighbor Table

- The **show ip ospf neighbor** command displays all neighbors that were discovered on active OSPF interfaces.

```
CSW1#show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
10.1.220.2	1	FULL/BDR	00:00:36	10.1.192.6	FastEthernet0/12
10.1.220.1	1	FULL/BDR	00:00:36	10.1.192.2	FastEthernet0/11
10.1.220.253	1	FULL/DR	00:00:37	10.1.151.194	Vlan111
10.1.220.3	1	FULL/BDR	00:00:30	10.1.192.19	Vlan129
10.1.220.4	1	2WAY/DROTHER	00:00:37	10.1.192.20	Vlan129
10.1.220.253	1	FULL/DR	00:00:35	10.1.192.18	Vlan129

State:

"FULL" indicates that the link-state database has been synchronized with the neighbor.
"2WAY/DROTHER" is acceptable on broadcast or nonbroadcast networks with more than three routers.

Dead Time:

Based on a 40-second dead time and 10-second hellos, this value should be between 30 and 40 seconds.

© 2009 Cisco Systems, Inc. All rights reserved.TSHOOT v1.9-71

To verify that all expected neighbor relationships are operational, you can use the **show ip ospf neighbor** command to display the OSPF neighbor table.

While two routers establish an adjacency and synchronize their link-state databases, they go through the following phases: Attempt (optional), Init, 2-Way, Exstart, Exchange, Loading, and Full. Therefore, the expected state for a neighbor relationship is "Full." The other states are transitory states and a neighbor should not be stuck in any of those states for an extended period.

The only exception to this rule is a broadcast or nonbroadcast network with more than three routers. On these types of networks, a designated router (DR) and backup designated router (BDR) are elected and all routers establish a full adjacency with the DR and BDR. Therefore, any two routers that are both not a DR and not a BDR (marked "DROTHER" in the **show** commands) will not transition any further than the two-way state.

In the example in the output, you can see that on interface Vlan129, this router has three neighbors: Neighbor 10.1.220.253, which is the DR; neighbor 10.1.220.3, which is the BDR; and neighbor 10.1.220.4, which is neither the DR nor the BDR. This router has transitioned to the full state with neighbor 10.1.220.253 and 10.1.220.3 (DR and BDR) and to the 2WAY state with neighbor 10.1.220.4 (DROTHER). This behavior is the expected behavior for these interfaces.

Debug OSPF Packet Exchange

- The **debug ip ospf packet** command can be used to verify the reception of OSPF packets.
- Transmitted packets are not listed.
- Only the OSPF packet headers are displayed.

```
CSW1#debug ip ospf packet
OSPF packet debugging is on
CSW1#
Apr 19 06:59:10.786 PDT: OSPF: rcv. v:2 t:1 l:56 rid:10.1.220.253
aid:0.0.0.100 chk:BD06 aut:0 auk: from Vlan129
Apr 19 06:59:10.937 PDT: OSPF: rcv. v:2 t:1 l:48 rid:10.1.220.1
aid:0.0.0.0 chk:8A97 aut:0 auk: from FastEthernet0/11
Apr 19 06:59:11.205 PDT: OSPF: rcv. v:2 t:1 l:48 rid:10.1.220.2
aid:0.0.0.0 chk:8A8E aut:0 auk: from FastEthernet0/12
Apr 19 06:59:11.583 PDT: OSPF: rcv. v:2 t:1 l:56 rid:10.1.220.4
aid:0.0.0.100 chk:BD06 aut:0 auk: from Vlan129
```

Area in IP Address Notation:
area 0.0.0.100 = area 100

Receiving Interface

Neighbor Router ID

© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0-72

When an OSPF neighbor relationship is not properly established, you can use several **debug** commands to display events related to the establishment of neighbor relationships. The most elementary command is the **debug ip ospf packet** command, which displays the headers of OSPF packets as they are received by the router.

This command lists only received packets. Transmitted packets are not displayed. Secondly, because interfaces that are not enabled for OSPF do not listen to the OSPF multicast addresses, packets are only shown for interfaces that are enabled for OSPF.

The most relevant fields in the header description of these packets are the following:

- **Type “t”**: The type field lists the type of the packet. The possible packet types are:
 - **Type 1**: Hello packets
 - **Type 2**: Database description packets
 - **Type 3**: Link-state request packets
 - **Type 4**: Link-state update packets
 - **Type 5**: link-state acknowledgement packets
- **Router ID (rid)**: The router ID field lists the router ID of the sending router. Note that the router ID is usually not the same as the source address of the packet.
- **Area ID (aid)**: The 32-bit area ID of the sending router is represented in dotted-decimal IP address format.
- **Authentication (aut)**: This field lists the authentication type. The possible types are:
 - **Type 0**: No (null) authentication
 - **Type 1**: Clear-text authentication
 - **Type 2**: Message Digest 5 (MD5) authentication
- **Interface (from)**: The interface on which the packet was received is listed here.

Note Only successfully received and accepted packets are listed in the output of the **debug ip ospf packet** command. If a mismatch exists between essential parameters in the header, such as the area id, the authentication type, or authentication data between this router and the neighbor, the packets from that neighbor will be silently discarded and not listed in the output of the debug.

The usefulness of this command for troubleshooting is limited because it does not display sent packets, packets received on an interface that is not enabled for OSPF, or packets that carry mismatched header information. However, because of the relatively limited amount of generated output, the command can be used to confirm the reception of correct hellos from a neighbor.

Debug OSPF Adjacencies

- The **debug ip ospf adj** command displays the different stages of the OSPF adjacency building process and other neighbor-related events such as the reception of LSAs.

```
CSW1#debug ip ospf adj
OSPF adjacency events debugging is on
*Apr 24 03:38:39.344 PDT: OSPF: Interface Serial0/0/0.122 going Up
*Apr 24 03:38:39.480 PDT: OSPF: 2 Way Communication to 10.1.221.2
on Serial0/0/0.122, state 2WAY
*Apr 24 03:38:39.480 PDT: OSPF: Send DBD to 10.1.221.2 on
Serial0/0/0.122 seq 0x761 opt 0x50 flag 0x7 len 32
*Apr 24 03:38:39.520 PDT: OSPF: Rcv LS UPD from 10.1.220.252 on
FastEthernet0/0 length 64 LSA count 1
*Apr 24 03:38:39.524 PDT: OSPF: Rcv LS UPD from 10.1.220.253 on
FastEthernet0/1 length 64 LSA count 1
*Apr 24 03:38:39.676 PDT: OSPF: Rcv DBD from 10.1.221.2 on
Serial0/0/0.122 seq 0x1921 opt 0x50 flag 0x7 len 32 mtu 1500 state
EXSTART
*Apr 24 03:38:39.676 PDT: OSPF: NBR Negotiation Done. We are the
SLAVE
*Apr 24 03:38:39.676 PDT: OSPF: Send DBD to 10.1.221.2 on
Serial0/0/0.122 seq 0x1921 opt 0x50 flag 0x2 len 132
<... further output omitted ...>
```

© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0-73

A very useful command for troubleshooting OSPF neighbor-related events is the **debug ip ospf adj** command, which displays all the different stages of the OSPF adjacency building process, as two neighbors transition from the init state to the full state. This command can be very helpful in diagnosing problems where a neighbor relationship is stuck in a particular stage of the adjacency building process.

This command also reveals mismatches in the basic parameters contained in the OSPF packet header, such as area ID mismatches, the source being on the wrong subnet, or authentication mismatches. However, it does not reveal other mismatches in hello parameters, such as hello timers, subnet masks, or flags.

Debug OSPF Events

- The **debug ip ospf events** command combines information that is displayed by the **debug ip ospf adj** command with information about reception and transmission of hellos.
- Mismatches in hello parameters are indicated:
 - Mismatched area number
 - Mismatched hello and dead time
 - Mismatched authentication type or data
 - Mismatched subnet and mask
 - Mismatched flags (stub area, NSSA)

```
CSW1#debug ip ospf events
OSPF events debugging is on
*Apr 24 03:41:17.168 PDT: OSPF: Rcv hello from 10.1.220.1 area 11 from
Serial0/0/0.111 10.1.194.5
*Apr 24 03:41:17.168 PDT: OSPF: Mismatched hello parameters from 10.1.194.5
*Apr 24 03:41:17.168 PDT: OSPF: Dead R 20 C 40, Hello R 5 C 10
*Apr 24 03:41:17.820 PDT: OSPF: Send hello to 224.0.0.5 area 11 on
Serial0/0/1 from 10.1.193.6
```

© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0-74

A third **debug** command that can be useful in troubleshooting the establishment of OSPF neighbor relationships is the **debug ip ospf events** command. This command displays the same information that is displayed by the **debug ip ospf adj** command. In addition, it displays the transmission and reception of hello packets and reports mismatches in the hello parameters.


Confirming the transmission of hello packets by using this command can be useful to you, because the **debug ip ospf packet** or **debug ip ospf adj** commands do not display the transmission of hello packets.

Secondly, this command can be used to display the reception of invalid hello packets. If a mismatch exists between the neighbors in the hello parameters that prevents the neighbor relationship from forming, this command will display the type of parameter mismatch and the value of the mismatched parameters. This command displays mismatches for the following parameters:

- Hello and dead timers
- Area ID
- Subnet and subnet mask
- Authentication type and authentication data
- Flags that signify the area type, such as stub or not-so-stubby area (NSSA)

Because this command displays more events than the **debug ip ospf adj** command, it is often better if you first enable the **debug ip ospf adj** command and only add the **debug ip ospf event** command if the **debug ip ospf adj** command does not yield the information you are interested in.

Sample OSPF Troubleshooting Flow



After you have verified that neighbor relationships have been established as expected, you should verify that the network topology information for the destination network that you are troubleshooting has been received correctly and entered into the OSPF link-state database.

The presence or absence of specific topology information in the OSPF link-state database can help in isolating the source of the problem.

Verify the OSPF Link-State Database for Intra-Area Routes

- OSPF intra-area routes are calculated based on the type 1 and type 2 LSAs for the area.

```
CSW1#show ip ospf database router 10.1.220.3

      OSPF Router with ID (10.1.220.252) (Process ID 100)

      Router Link States (Area 100)

<... output omitted ...>

Link connected to: a Stub Network
(Link ID) Network/subnet number: 192.168.224.240
(Link Data) Network Mask: 255.255.255.240
Number of TOS metrics: 0
TOS 0 Metrics: 1

Link connected to: a Transit Network
(Link ID) Designated Router address: 10.1.192.18
(Link Data) Router Interface address: 10.1.192.19
Number of TOS metrics: 0
TOS 0 Metrics: 1
```

In order to decide what information to look for in the link-state database, you first need to discern what type of route you are interested in. If the destination network that you are troubleshooting is in the same area as the router that you are troubleshooting from, you know that the path to this destination network was derived from the type 1 and type 2 LSAs in the database of that area. To begin with, you can verify whether the directly connected routers properly advertise the destination network. To do this verification, you should display the router (type 1) for the connected routers by issuing the command **show ip ospf database router-id** for these routers. To troubleshoot OSPF effectively, you need to know the router IDs of all routers in your network because these IDs are used to identify a router in many of the OSPF show commands.

As part of the type 1 router LSA for a specific router, all subnets corresponding to a point-to-point link, loopback interface, or nontransit broadcast network (Ethernet) are listed as stub networks. If the target network is missing in this list, this absence indicates that the interface on the advertising router has not been enabled for OSPF.

In the example in the figure, you can see that subnet 192.168.224.240/28 is advertised by router 10.1.220.3 in area 1.

For transit networks (such as an Ethernet LAN with multiple routers attached), a link to the DR for the segment is listed. This listing points to the type 2 network LSA that contains the full topology information for the segment.

In the example in the figure, you can see that this router is connected to a transit network with router 10.1.192.18 as the DR. Note that this IP address is the interface IP address of the DR, not the router ID.

Verify the OSPF Link-State Database for Intra-Area Routes

- OSPF intra-area routes are calculated based on the type 1 and type 2 LSAs for the area.

```
CSW1#show ip ospf database network 10.1.192.18

      OSPF Router with ID (10.1.220.252) (Process ID 100)

      Net Link States (Area 100)

Routing Bit Set on this LSA
LS age: 543
Options: (No TOS-capability, DC)
LS Type: Network Links
Link State ID: 10.1.192.18 (address of Designated Router)
Advertising Router: 10.1.220.253
LS Seq Number: 80000029
Checksum: 0x6B40
Length: 40
Network Mask: /29
  Attached Router: 10.1.220.253
  Attached Router: 10.1.220.3
  Attached Router: 10.1.220.4
  Attached Router: 10.1.220.252
```

© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0-77

Full information about a transit LAN can be displayed by issuing the **show ip ospf database designated-router** command, using the IP address of the DR that was listed in the type 1 router LSA for one of the routers connected to the transit LAN. In the type 2 LSA, the DR advertises the subnet mask and connected routers for the segment. The connected routers are listed by their router ID values.

In the example in the figure, a subnet mask of /29 is advertised for the transit LAN and four connected routers are listed.

Verify the OSPF Link-State Database for Interarea Routes

- OSPF interarea routes are calculated based on type 3 LSAs generated by the ABR for the area.

```
IR01#show ip ospf database summary 10.1.152.0

      OSPF Router with ID (10.1.220.3) (Process ID 100)

      Summary Net Link States (Area 100)

Routing Bit Set on this LSA
LS age: 41
Options: (No TOS-capability, DC, Upward)
LS Type: Summary Links(Network)
Link State ID: 10.1.152.0 (summary Network Number)
Advertising Router: 10.1.220.252
LS Seq Number: 8000001D
Checksum: 0xE3B3
Length: 28
Network Mask: /24
      TOS: 0 Metric: 1
```

If the destination network that you are troubleshooting is in a different area than the area of the router that you are troubleshooting from, the router will not learn about this network through type 1 and type 2 LSAs, because these are only used for intra-area routes. OSPF interarea routes are calculated based on type 3 LSAs that are generated by the Area Border Routers (ABR) for the area.

To verify the availability of a specific target network in a different area, you can use the **show ip ospf database summary subnet** command, where *subnet* is the subnet IP address of the prefix that you are interested in.

The type 3 summary LSA contains the subnet, mask, and cost of the targeted subnet and lists the router ID of the advertising ABR. If multiple ABRs are advertising the same network, all entries are listed.

In the example in the figure, we see that subnet 10.1.152.0/24 is advertised with a cost of 1 by ABR 10.1.220.252. The cost advertised by the ABR is the cost from the advertising ABR to the target network. When the router executes the SPF algorithm, it calculates its own cost to reach the ABR within the area and add that to the cost advertised by the ABR.

If you do not find an entry for the target network, the next step is to connect to the ABR, which you expected to be advertising the route, and verify if the route is available there.

Verify the OSPF Link-State Database for External Routes

- OSPF external routes are calculated based on type 5 LSAs generated by the redistributing ASBR.

```
IR01#show ip ospf database external 10.1.160.64

OSPF Router with ID (10.1.220.3) (Process ID 100)

Type-5 AS External Link States

Routing Bit Set on this LSA
LS age: 1100
Options: (No TOS-capability, DC)
LS Type: AS External Link
Link State ID: 10.1.160.64 (External Network Number )
Advertising Router: 10.1.220.1
LS Seq Number: 80000023
Checksum: 0xBD2B
Length: 36
Network Mask: /26
Metric Type: 2 (Larger than any link state path)
TOS: 0
Metric: 20
Forward Address: 0.0.0.0
External Route Tag: 0
```

© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0-79

Finally, if the destination network that you are troubleshooting did not originate in the OSPF network, but was redistributed from a different source, the OSPF router will learn about this network through type 5 external routes that are injected into the OSPF database by an Autonomous System Boundary Router (ASBR).

To verify the availability of a specific type 5 external LSA in the OSPF database, issue the command **show ip ospf database external *subnet***, where *subnet* is the subnet IP address of the prefix that you are interested in.

The type 5 summary LSA contains the subnet, mask, metric type, and cost of the targeted subnet. In addition, it lists the router ID of the advertising ASBR. If multiple ASBRs are advertising the same network, all entries are listed.

In the example in the figure, you see that subnet 10.1.160.64/26 is advertised with a cost of 20 as a metric-type 2 external route by ASBR 10.1.220.1.

If you do not find an entry for the target network, the next step is to connect to the ASBR that you expected to be advertising the route, and verify whether the route is available. If the route is available, but not advertised by the ASBR, you should troubleshoot the route redistribution process on that router.

Verify the OSPF Link-State Database for External Routes (Cont.)

- ASBRs advertise their ASBR status in their type 1 LSA.

```
CSW1#show ip ospf database router 10.1.220.1

      OSPF Router with ID (10.1.220.252) (Process ID 100)

      Router Link States (Area 0)

Routing Bit Set on this LSA
LS age: 194
Options: (No TOS-capability, DC)
LS Type: Router Links
Link State ID: 10.1.220.1
Advertising Router: 10.1.220.1
LS Seq Number: 80000024
Checksum: 0x9557
Length: 60
AS Boundary Router
Number of Links: 3

<... further output omitted ...>
```

Instead of connecting to the ASBR, the OSPF database can also be used to verify whether any form of redistribution has been configured on the router that is supposed to be an ASBR. If that router is in the same area as the router that you are troubleshooting from, you can inspect the type 1 router LSA for the ASBR and verify that it advertises itself as an ASBR.

In the example in the figure, you can see that router 10.1.220.1 announces its ASBR status in its type 1 LSA.

If the router does not advertise its ASBR status in its type 1 LSA, redistribution has not been configured correctly on that router.

Verify the OSPF Link-State Database for External Routes (Cont.)

- Type 4 summary ASB LSAs are generated by ABRs to announce the availability of ASBRs to other areas.

```
IR01#show ip ospf database asbr-summary 10.1.220.1

      OSPF Router with ID (10.1.220.3) (Process ID 100)

      Summary ASB Link States (Area 100)

Routing Bit Set on this LSA
LS age: 1104
Options: (No TOS-capability, DC, Upward)
LS Type: Summary Links(AS Boundary Router)
Link State ID: 10.1.220.1 (AS Boundary Router address)
Advertising Router: 10.1.220.252
LS Seq Number: 80000020
Checksum: 0xD677
Length: 28
Network Mask: /0
      TOS: 0 Metric: 1

<... further output omitted ...>
```

© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0--81

If the ASBR is not in the same area as the router that you are troubleshooting from, you do not have the type 1 LSA of the ASBR in the database of the router, and as a result, you cannot verify the ASBR status of the redistributing router by displaying the type 1 LSA. However, if an ASBR is available in a different area, the area border routers for the area will generate a type 4 summary ASB entry to announce the availability of the ASBR. The presence or absence of a type 4 entry can also yield a clue about the operation of the redistribution.

The **show ip ospf asbr-summary** *router-id* command can be used to verify if a type 4 summary ASB LSA exists for the ASBR with the specified router ID.

In the figure in the example, ABR 10.1.220.252 announces the availability of ASBR 10.1.220.1.

Verify the Availability of ABRs and ASBRs

- ABRs and ASBRs that are available in an area and the path and metric to reach these routers can be displayed using the **show ip ospf border-routers** command.

```
IR01#show ip ospf border-routers

OSPF Process 100 internal Routing Table

Codes: i - Intra-area route, I - Inter-area route

i 10.1.220.252 [1] via 10.1.192.17, FastEthernet0/1, ABR, Area 100, SPF 16
i 10.1.220.253 [1] via 10.1.192.18, FastEthernet0/1, ABR, Area 100, SPF 16
I 10.1.220.1 [2] via 10.1.192.17, FastEthernet0/1, ASBR, Area 100, SPF 16
I 10.1.220.1 [2] via 10.1.192.18, FastEthernet0/1, ASBR, Area 100, SPF 16
I 10.1.220.2 [2] via 10.1.192.17, FastEthernet0/1, ASBR, Area 100, SPF 16
I 10.1.220.2 [2] via 10.1.192.18, FastEthernet0/1, ASBR, Area 100, SPF 16
I 10.1.220.4 [1] via 10.1.192.20, FastEthernet0/1, ASBR, Area 100, SPF 16
```

i = intra-area
I = interarea

Router ID of
ABR or ASBR

Cost to reach the
ABR or ASBR


© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0-82

During the execution of the SPF algorithm, a router combines the information from the various LSAs that contain information about ABR and ASBR status and calculates the shortest paths to each of the ABRs and ASBRs. You can use the **show ip ospf border-routers** command to view the result of this calculation.

In the example in the figure, you can see that area 100 has two ABRs: 10.1.220.252 and 10.1.220.253. The cost to reach each of those two routers is 1, as can be seen from the number in the square brackets. This cost is important to know, because it is added to the cost advertised by these routers in their type 3 LSAs to obtain the total cost to the destination network.

Sample OSPF Troubleshooting Flow



© 2009 Cisco Systems, Inc. All rights reserved.


TSHOOT v1.0-83

If all appropriate entries are available in the OSPF link-state database, the result should be correct routes in the IP routing table after calculation of the SPF algorithm. Unfortunately, the results of the SPF algorithm for each individual route cannot be directly verified.

Remember that OSPF competes with other routing sources to install routes in the routing table and, therefore, an OSPF route may not be installed in the routing table because a route with a better administrative distance from a different source is available.

Troubleshooting Route Redistribution

The figure illustrates an example of a method that you could follow to diagnose and resolve problems related to route redistribution.




When do you start troubleshooting route redistribution?

There are two major reasons to start troubleshooting the route redistribution. The first reason is that you are experiencing IP connectivity problems in an environment where information from a specific routing domain is redistributed into a different routing domain and the connectivity problem is caused by a route from the source routing domain that is not available on one or more of the routers participating in the destination routing domain. (Note that in this section the terms *source* and *destination* are used to indicate the source and destination of the routing information, not the source and destination of a traffic flow.) In this scenario, the cause of the problem is that the exchange of routing information between the source routing domain and the destination routing domain is not working correctly.

The second reason to start troubleshooting route redistribution is if you are experiencing IP connectivity problems caused by the use of incorrect routing information by some of the routers in a network that uses route redistribution. This behavior could be caused by routing information feedback or improper route selection.

Sample troubleshooting flows for each of these scenarios are provided in this section.

Sample Route Redistribution Troubleshooting Flow



© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0-85

The first scenario in which you start troubleshooting route redistribution is when redistribution is configured and you are troubleshooting connectivity problems to a network in the source routing domain from a router in the destination routing domain. This type of problem is usually discovered during a generic IP connectivity troubleshooting process. During that process, you discover that a route is missing from the routing table on one of the routers in the destination routing domain, while the route is present in the routing tables of the routers in the source routing domain.


Troubleshooting redistribution consists of four generic steps:

- Troubleshooting the source routing protocol
- Troubleshooting route selection and installation
- Troubleshooting redistribution
- Troubleshooting the destination protocol

In this scenario, the reason you would start troubleshooting the redistribution is that the route is available in the source routing domain, but not in the destination routing domain. Therefore, the first step has already been taken at this point. If the route is not available everywhere in the source routing domain to begin with, you do not have any reason to start troubleshooting redistribution, but you should initiate a troubleshooting process for the source routing protocol first.

Therefore, you should start at the second step: troubleshooting route selection and installation.

Sample Route Redistribution Troubleshooting Flow



© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0-86

Not many tools are specifically targeted at troubleshooting the redistribution process. The redistribution process takes routes from the routing table after they have been installed by the source routing protocol and then injects them into the data structures of the destination protocol. Therefore, the main tools that are available to track this flow of information are the commands that allow you to examine the routing table and the destination protocol data structures.

After you have verified that the routes are injected into the destination data structures of the protocol, you have finished troubleshooting the actual redistribution process. If the routes are not properly propagated by the destination protocol, you should initiate a troubleshooting process for the destination protocol.

Verify the Routing Table

- Routes that are being redistributed are marked in the routing table.
- If a route is redistributed and permitted by any possible route maps or distribute lists applied to the redistribution, it will be marked as “advertised by” the protocol that it is redistributed into.

```
CRO1#show ip route 10.1.152.0 255.255.255.0
Routing entry for 10.1.152.0/24
  Known via "ospf 100", distance 110, metric 2, type inter area
  Redistributing via eigrp 1
  Advertised by eigrp 1 metric 64 10000 255 1 1500
  Last update from 10.1.192.1 on FastEthernet0/0, 22:15:58 ago
  Routing Descriptor Blocks:
    * 10.1.192.9, from 10.1.220.253, 22:15:58 ago, via FastEthernet0/1
      Route metric is 2, traffic share count is 1
    10.1.192.1, from 10.1.220.252, 22:15:58 ago, via FastEthernet0/0
      Route metric is 2, traffic share count is 1
```


© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0--87

The best tool available in troubleshooting redistribution problems is the **show ip route network mask** command. Routes that are being redistributed and advertised to other routers by the destination protocol are marked with a line starting with “Advertised by” and then it lists the destination protocol and any parameters configured on the redistribution statement, such as configured metrics and metric type.

What makes this command very useful is that it takes into account any route maps or distribute lists that are applied to the redistribution.

Sample Route Redistribution Troubleshooting Flow




© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0--88

The second common scenario that may lead you to start troubleshooting route redistribution is a scenario in which you discover that traffic is using unexpected suboptimal routes to reach certain destinations or traffic enters a routing loop. This situation is often discovered when you are troubleshooting IP connectivity to a certain destination and using the **show ip route** and **traceroute** commands to track the flow of traffic. When you are redistributing routing information between routing protocols, you have to be aware that improper route selection or routing feedback may cause suboptimal paths to be used or may cause traffic to enter a routing loop. Therefore, whenever you spot unexpected routing behavior in a network that uses redistribution, you should consider routing feedback or improper route selection as a possible cause.

The following symptom is typical in the case of a redistribution problem: On the router that you are troubleshooting, the expected route is available, but it is not selected as the best route in the routing table. A route from a different protocol, or a route of the same protocol, but one that originated from a different source, is selected as the best route and installed in the routing table.



The first question you need to ask yourself at this point is if the route is only improperly selected. In other words, you expected this route to be present, but did not want it to be selected as the best route. If this scenario is the case, you can manipulate the route selection process by changing the administrative distance. This change can be done for all routes that were learned via a particular routing protocol or selectively, using an access list.

If the route was not only improperly selected, but also should not have been present in the routing protocol data structures in this router at all, you need to track the source of the route and use route filtering techniques at the source to stop it from being advertised.

Verify the Routing Table

- The source of a route is marked in the routing table.
- By tracking the source of a route, you can find the point where incorrect routing information is being injected.

```
CRO1#show ip route 10.1.152.0 255.255.255.0
Routing entry for 10.1.152.0/24
  Known via "ospf 100", distance 110, metric 2, type inter area
  Redistributing via eigrp 1
  Advertised by eigrp 1 metric 64 10000 255 1 1500
  Last update from 10.1.192.1 on FastEthernet0/0, 22:15:58 ago
  Routing Descriptor Blocks:
    * 10.1.192.9, from 10.1.220.253, 22:15:58 ago, via FastEthernet0/1
      Route metric is 2, traffic share count is 1
    10.1.192.1, from 10.1.220.252, 22:15:58 ago, via FastEthernet0/0
      Route metric is 2, traffic share count is 1
```

© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0--80

The source of a route in the routing table is marked by the “from” field that follows the next-hop IP address. For distance vector protocols, the source and next-hop address are typically the same, but for a link-state protocol such as OSPF, the source is the router that originated the LSA that the route is based on. By tracking the routing source from router to router, you can determine the point where the incorrect routing information is injected into the routing protocol’s data structures and you can apply filtering to stop it from being propagated.

Lab Debrief Notes

Use these notes sections to write down the primary learning points that are discussed during the Lab Debrief.

Lab 5-2: Alternate Solutions

[illegible]

Lab 5-2: Alternate Methods and Processes

[illegible]

[illegible]

Lab 5-2: Important Commands and Tools

[illegible]

Lab 5-2: References

If you need more information on the commands and their options, you can go to the following sections of <http://www.cisco.com>.

- Cisco Systems, Inc. *Cisco IOS IP Routing Protocols Command Reference*. San Jose, California, November 2008:
http://www.cisco.com/en/US/docs/ios/iproute/command/reference/irp_book.html
- Cisco Systems, Inc. IP Routing Troubleshooting TechNotes, Open Shortest Path First:
http://www.cisco.com/en/US/tech/tk365/tsd_technology_support_troubleshooting_technotes_list.html#anchor8

Lab 5-3: Border Gateway Protocol

Complete this lab activity to practice what you learned in the related module.

Activity Objective


In this activity, you will troubleshoot various problems related to the BGP. After completing this activity, you will be able to meet these objectives:

- Diagnose and resolve problems related to the BGP exterior gateway protocol
- Document troubleshooting progress, configuration changes, and problem resolution

Information Packet

The figure illustrates what you will accomplish in this activity:

Visual Objective



Required Resources

These are the resources and equipment that are required to complete this activity:

- One PC with Internet access per team member
- Four Cisco Catalyst 3560 Series Switches per team
- Six Cisco 1841 Integrated Services Routers per team
- Three client PCs per team
- One server per team

Job Aids

These job aids are available to help you complete the lab activity.

- Trouble tickets
- Troubleshooting log

Introduction: Implementation of BGP


Your company has decided to implement several new Internet-based services. All current web services that the company offers are hosted at an external data center. The company has decided to build a new in-house data center from which the new services will be hosted. The servers that are currently externally hosted will also be moved to this new data center.

Your company is already using two different ISPs for redundant outbound Internet access by use of NAT. From a recent acquisition, your company has obtained a registered Autonomous System (AS) number (64568) and address block (172.17.76.0/22), which will be used for the new services. After consulting with both ISPs, Internet Service Provider 1 and Internet Service Provider 2, your company has decided to use BGP with both of these service providers to provide redundant inbound connectivity to your company AS and IP address block.

Because BGP is new technology to your organization, they have decided to implement BGP on the existing Internet routers IRO1 and IRO2 in preparation of the new data center project. This decision will allow some services to be migrated to the new IP address block before moving them to the newly built datacenter.

Your support team has been working closely together with the engineering team to prepare the implementation. You have received confirmation from both ISPs that they have prepared their routers for the BGP implementation.

The high-level BGP design is outlined in the following figure:



An external BGP peering will be established between routers IRO1 and ISP1 and a second external peering will be established between routers IRO2 and ISP2. In addition, an internal BGP peering between routers IRO1 and IRO2 will be established.

Routers IRO1 and IRO2 will advertise the full 172.17.76.0/22 block to both ISP routers, ISP1 and ISP2. No other prefixes are allowed to be advertised to routers ISP1 and ISP2 in order to protect the company AS from accidentally becoming a transit AS.

Each of the ISP routers will send a default route and a limited set of additional prefixes to routers IRO1 and IRO2. The default route will be redistributed into EIGRP by both routers IRO1 and IRO2. No other routes will be redistributed.

It is Friday evening and the engineering team has just configured routers IRO1 and IRO2 for BGP. To facilitate testing, VLAN 145 DMZ and the corresponding subnet 172.17.76.0/24 have been created. The client PC CLT1 has been temporarily assigned to this VLAN for testing purposes. All other devices, which have IP addresses in the 10.1.0.0/16 range, are still using NAT and their Internet access is not affected by the BGP configuration.

You are on standby to assist in troubleshooting and testing the solution.

Trouble Ticket O: BGP Peering to Router ISP1 Not Established

When you ask your colleague from the engineering team for a status update, you hear that the peering to router ISP2 has been established, but that the peering to router ISP1 is not being established. He has contacted the support team of Internet Service Provider 1 to find out if they have a problem at their end, but they state that everything is correctly configured on router ISP1. You offer to help troubleshoot this problem.

Your task is to diagnose the problem and, if possible, establish the BGP peering session between router IRO1 and router ISP1. After the peering has been established, you need to verify that traffic from subnet 172.17.76.0/24 can be sent to the Internet via router ISP1 and that the return traffic can be received via router ISP1.

Trouble Ticket P: Client CLT1 Cannot Reach the Internet

Even without the BGP session between routers IRO1 and ISP1, Internet connectivity from test PC CLT1 should be available via router IRO2 and Internet Service Provider 2. When one of your colleagues from engineering tries to test the connectivity by browsing to <http://www.isp3.local>, it does not work. You offer to assist in troubleshooting this problem.

Your task is to diagnose this problem and, if possible, resolve it.

Instructions

Together with your team members, create a troubleshooting plan to divide the work, assign each team member appropriate roles and coordinate device access between team members. Together, work on Trouble Tickets O and P to resolve the issues. Document your progress in the following Troubleshooting Logs in order to help facilitate efficient communication within the team and to have an overview of your troubleshooting process for reference during the Lab Debrief discussions.

You are allowed a total of 45 minutes to complete as many of the trouble tickets as you can. After 45 minutes, the instructor will debrief the lab and review all trouble tickets and their solutions. The main objective for the troubleshooting labs in this course is to give you an opportunity to practice structured troubleshooting. Fixing the problems is secondary to practicing proper processes and procedures.

Lab Setup

The instructor will provide you with directions to prepare the lab equipment for this lab. After the instructor indicates that the lab is fully prepared, you are ready to start troubleshooting.

Troubleshooting Log

Use this log to document your actions and results during the troubleshooting process.

Trouble Ticket O

Your task is to diagnose the BGP problem and, if possible, establish a peering session between router IRO1 and router ISP1. After the peering has been established, you need to verify that traffic from subnet 172.17.76.0/24 can be sent to the Internet via router ISP1 and that the return traffic can be received via router ISP1.

Note Refer to the Activity Verification items at the end of this log to verify that you have successfully completed this task.

[illegible]

Troubleshooting Log

Use this log to document your actions and results during the troubleshooting process.

Trouble Ticket P


Your task is to diagnose the problem of failing Internet connectivity via routers IRO2 and ISP2 and, if possible, resolve it.

Note Refer to the Activity Verification items at the end of this log to verify that you have successfully completed this task.

[illegible]

Lab 5-3: Sample Troubleshooting Flows

The figure illustrates an example of a method that you could follow to diagnose and resolve problems related to the BGP.



Encountering the following issues is typically cause to start investigating BGP operations when you are using BGP as an exterior gateway protocol to connect to other autonomous systems and you are troubleshooting IP connectivity to a destination in a different AS: You find that that a route to the destination network is missing from the routing table of one of the routers, a different route than expected was selected to forward the packets to that destination, or return traffic from the other AS is not making it back to the source.

Troubleshooting problems with missing return traffic usually requires coordination with those who are responsible for the routing in the destination AS and possibly even intermediate autonomous systems. The only thing you can verify from within your own AS is if your routing information is correctly passed to the neighbor AS. Propagation of your routes beyond your direct peers cannot be verified without access to routers in other autonomous systems.


Therefore, this flow will focus mainly on troubleshooting traffic to a destination network in a different AS than your own. However, commands that are helpful in troubleshooting route advertisement to a different AS are also highlighted where appropriate.

In order to install a route into the routing table, each router that uses BGP goes through several stages:

- Establishes neighbor relationships with its configured neighbors
- Exchanges routing information with neighbors and stores the received information in the BGP table
- Selects the best route from the available routes and installs it in the routing table

Errors during any of these stages can lead to missing routing information or to the wrong routing information being installed in the routing table.

The order of verification of the different stages of this process is not important, as long as a structured approach is used.



BGP does not discover neighbors. Neighbor relationships are established based on an explicit configuration on both routers that participate in the peering session.

BGP uses TCP as a transport protocol and, therefore, establishing a peering relationship always starts with the establishment of a TCP session on port 179 between the configured neighbor IP addresses. By default, both neighbors will attempt to initiate the TCP session to the configured IP address of the neighbor. When a router receives an incoming session request, it will compare the source IP address of the session to its list of configured neighbors. It will only accept the session if the source IP address matches one of the IP addresses of its configured neighbors. Therefore, it is important that a router always source the BGP packets that it sends to a specific neighbor from the IP address that has been configured as the neighbor IP address on the peer router. For neighbors that are directly connected on an interface, the correct source address is automatically used. For neighbors that are not directly connected, the appropriate source IP address for the session to a neighbor may need to be selected by using the **neighbor ip-address update-source interface-id** command.

Verify the BGP Neighbor Table

- The **show ip bgp summary** command presents an overview of all configured BGP neighbors in addition to general statistics and parameters.

```
IR01#show ip bgp summary
BGP router identifier 10.1.220.3, local AS number 64568
BGP table version is 44, main routing table version 44
6 network entries using 702 bytes of memory
11 path entries using 572 bytes of memory
8/4 BGP path/bestpath attribute entries using 992 bytes of memory
5 BGP AS-PATH entries using 120 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 2386 total bytes of memory
BGP activity 12/6 prefixes, 22/11 paths, scan interval 60 secs
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
10.1.220.4	4	64568	180	175	44	0	0	01:36:15	5
192.168.224.254	4	65525	246	243	44	0	0	02:01:36	5

© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0-95

To verify that all expected neighbor relationships are operational, you can use the **show ip bgp summary** command to display a summary of the BGP neighbor table. This command lists important BGP parameters, such as the AS number and router ID, statistics about the memory consumption of the various BGP data structures, and a brief overview of the configured neighbors and their state.

For each neighbor, the configured IP address and AS of the neighbor are listed. The “Up/Down” column lists the time that has elapsed since the last state change. For a neighbor that is currently up, it lists the time that has elapsed since the session was established. For a neighbor that is down, it lists the time that has elapsed since the session was lost.

The most important column that is used to verify the operational state of the neighbor is the “State/PfxRcd” column. This column can display the following values:

- **Idle:** This state indicates that there is no session with the peer and that the router is not currently attempting to establish a session with the peer. The router is ready to accept incoming sessions.
- **Idle (Admin):** This state indicates that the session has been administratively shut down by someone using the **neighbor ip-address shutdown** command.
- **Active:** The router is actively trying to open a TCP session with the neighbor. If it does not succeed in establishing the session, the router will toggle between the Idle and Active states.
- **Open Sent:** An Open message has been sent to the neighboring router containing the router ID, autonomous system number, BGP version, hold timer, and capabilities.
- **Open Confirm:** An Open message from the neighbor has been received, the parameters in the message have been processed and accepted, and a hello message has been sent to acknowledge the acceptance of the neighbor’s Open message.

- **Number of received prefixes:** After an acknowledgment from the neighbor, confirming the reception of this router's Open message, the state of the session moves to the Established state. At this point, the "State/PfxRcd" column will not list the state, but the number of prefixes that have been received from that neighbor and installed in the BGP table.

The desired result is to see a number listed in this column, because that indicates that the session with the peer has been successfully established. The Open Sent and Open Confirm states are transitory states. When the state for a neighbor toggles between Active and Idle, the toggling is an indication that the router is not being successful in establishing a session with the neighbor.

The **show ip bgp neighbor *ip-address*** command can be used to display additional parameters and extensive statistics about the peering session. For more detail about these parameters and statistics, consult the BGP command references on <http://www.cisco.com>.

Verify IP and TCP Connectivity

- The **debug ip tcp transactions** command can be used to gather clues about the TCP connections that is used by BGP.

```

IRO1#debug ip tcp transactions
TCP special event debugging is on
IRO1#
<... output omitted ...>
Apr 29 01:13:27.024 PDT: TCB63CC9170 bound to 10.1.220.3.50886
Apr 29 01:13:27.024 PDT: Reserved port 50886 in Transport Port Agent for
TCP IP type 1
Apr 29 01:13:27.024 PDT: TCP: sending SYN, seq 3840777289, ack 0
Apr 29 01:13:27.024 PDT: TCP0: Connection to 10.1.220.4:179, advertising
MSS 536
Apr 29 01:13:27.024 PDT: TCP0: state was CLOSED -> SYNSENT [50886 ->
10.1.220.4(179)]
Apr 29 01:13:27.024 PDT: TCP0: state was SYNSENT -> CLOSED [50886 ->
10.1.220.4(179)]
Apr 29 01:13:27.028 PDT: Released port 50886 in Transport Port Agent for
TCP IP type 1 delay 240000
Apr 29 01:13:27.028 PDT: TCP0: bad seg from 10.1.220.4 -- closing
connection: port 50886 seq 0 ack 3840777290 rcvnx 0 rcvwnd 0 len 0
Apr 29 01:13:27.028 PDT: TCP0: connection closed - remote sent RST
Apr 29 01:13:27.028 PDT: TCB 0x63CC9170 destroyed

```

If a session to one of the neighbors is not established correctly, you can take several steps to diagnose the issue. The first step is to test IP connectivity to the IP address of the neighbor by using the **ping** command. Make sure that you specify the same source interface for the **ping** command that is also used as the source interface for the BGP session. If this ping fails, you should initiate a troubleshooting process to restore IP connectivity to the neighbor first.

If the ping is successful, the next step is for you to determine whether the TCP session with the neighbor is established and successively torn down again, or if the TCP session is never established.

The **debug ip tcp transactions** command can be used to investigate whether the TCP session is refused (indicated by the reception of a TCP RST), established, and subsequently torn down again (indicated by the normal TCP initiation and termination handshakes), or if no response is received at all from the neighbor.

In the example output in the figure, you can see that the TCP session to IP address 10.1.220.4 and TCP port 179 is refused by the peer, as indicated by the reception of the TCP RST from the peer. Clues like these can help you eliminate possible problem causes. For instance, in this particular example, the output eliminates an access list as the cause of the problem because a TCP RST has been successfully received from the neighbor in response to the transmitted TCP SYN. In general, the fact that the peer refuses the session indicates that it does not recognize the session as coming from one of its configured neighbors. Possible causes are a missing neighbor statement or a mismatch between the configured IP address on the neighbor and the source IP address used by this router. Note that the source IP address and TCP port of the session are also displayed in the output of the debug as “bound to 10.1.220.3.50886.” You will have to work together with the party that manages the peer router to determine the exact cause of the problem.

Debug BGP Neighbor Establishment

- The **debug ip bgp** command displays the subsequent stages of BGP session establishment.

```
IR01#debug ip bgp
BGP debugging is on for address family: IPv4 Unicast
IR01#
BGP: 192.168.224.254 passive open to 192.168.224.241
BGP: 192.168.224.254 went from Active to Idle
BGP: 192.168.224.254 went from Idle to Connect
BGP: 192.168.224.254 rcv message type 1, length (excl. header) 26
BGP: 192.168.224.254 rcv OPEN, version 4, holdtime 180 seconds
BGP: 192.168.224.254 went from Connect to OpenSent
BGP: 192.168.224.254 sending OPEN, version 4, my as: 64568, holdtime 180
seconds
BGP: 192.168.224.254 rcv OPEN w/ OPTION parameter len: 16
BGP: 192.168.224.254 rcvd OPEN w/ optional parameter type 2 (Capability) len 6
BGP: 192.168.224.254 OPEN has CAPABILITY code: 1, length 4

<... output omitted ...>


BGP: 192.168.224.254 rcvd OPEN w/ remote AS 65525
BGP: 192.168.224.254 went from OpenSent to OpenConfirm
BGP: 192.168.224.254 send message type 1, length (incl. header) 45
BGP: 192.168.224.254 went from OpenConfirm to Established
```

© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0-97

If the TCP session is successfully established, but consecutively torn down again, the typical cause is one of the BGP peers rejecting one of the parameters in the received Open message from the peer. The **debug ip bgp** command displays the successive state transitions during the establishment of the BGP peering. If one of the peers decides to close the session because of a parameter problem, such as a mismatched AS or an invalid router ID, the debug will also display information about the exact cause.

Sample BGP Troubleshooting Flow



© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0-98


After you have verified that neighbor relationships have been established as expected, you should verify that the route for the destination network that you are troubleshooting has been received correctly from all appropriate neighbors. BGP stores all routes that it receives from its neighbors in the BGP table and then selects the best route for each prefix to be installed in the routing table and advertised to other neighbors.

By investigating all available paths to the destination network in the BGP table, you can see if all the paths you expected to find are available and if multiple paths to the same prefix are listed, which one was selected. In addition, you can see all the associated BGP attributes for the route, which can be useful for verifying the path selection process and the results of possible attribute manipulation by route maps that are used.

If routes are missing from the BGP table, you may need to debug the BGP route exchange process to see if they were not received, or if they were not entered into the BGP table.

Verify the BGP Table

- The **show ip bgp network mask** command displays all available paths and associated BGP attributes that are present in the BGP table for a specific prefix.



The BGP table contains all routes that were received from all neighbors and were not denied by an incoming access list, prefix list, or route map.

When you issue the command **show ip bgp network mask** to display the content of the BGP table for a specific prefix, the information is organized in the following manner. The entry for each available path in the table starts with the AS path attribute of the path (using the word “Local” to represent the empty AS path string). On the following lines, the other BGP attributes of the route, such as the next hop, origin code, and local preference are listed. In addition, other information associated with the route is displayed. For example, the route is marked as internal if it was received from a BGP neighbor in the same AS or external if it was received from a neighbor in a different AS. The path that was selected as the best path by the BGP path selection algorithm is marked with the word “best.”

The following section uses the output that is also displayed in the figure as an example to demonstrate how to interpret the output of this command. This output is interspersed with comments that explain the important fields and their interpretation.

```
IR01#show ip bgp 172.34.224.0 255.255.224.0
BGP routing table entry for 172.34.224.0/19, version 98
Paths: (2 available, best #1, table Default-IP-Routing-Table)
```

Two paths are available to reach prefix 172.34.224.0/19. The first path listed has been selected as the best path.

```
  Advertised to update-groups:
    2
```

The best path is advertised to all neighbors in update-group 2. (Use the **show ip bgp update-group** command to view the neighbors that are member of a specific update-group).

```
  65525 65486
```


The first path has “65525 65486” as its AS path attribute, which indicates that the route has originated in AS 65486, and then passed to AS 65525, which subsequently passed it to this AS.

```
192.168.224.254 from 192.168.224.254 (192.168.100.1)
```

The BGP next hop for this route is 192.168.224.1. The route was received from neighbor 192.168.224.1 and the router ID of that neighbor is 192.168.100.1.

```
Origin IGP, localpref 100, valid, external, best
```

The origin attribute for this route is “IGP” and the local preference attribute has a value of 100. This route is a valid route that is received from an external BGP peer, and it has been selected as the best path.

```
64566 65486
```

The second path has “64566 65486” as its AS path attribute, which indicates that the route has originated in AS 65486, and then passed to AS 64566, which subsequently passed it to this AS.

```
172.24.244.86 (metric 30720) from 10.1.220.4 (10.1.220.4)
```

The BGP next hop for this route is 172.24.244.86 and the IGP metric to reach this next hop IP address is 30720 (which is the EIGRP metric listed in the routing table to reach 172.24.244.86). The route was received from neighbor 10.1.220.4, and the router ID of that neighbor is also 10.1.220.4.

```
Origin IGP, metric 0, localpref 100, valid, internal
```

The origin attribute for this route is “IGP,” the multi-exit discriminator (MED) attribute has a value of 0, and the local preference attribute has a value of 100. The route is a valid route received from an internal BGP peer.

For troubleshooting purposes the AS path, next hop, and best path indicator are the most important fields in the output of this command. For a full description of all possible fields in the output of this command, refer to the BGP command references on <http://www.cisco.com>.

Instead of viewing a specific entry in the BGP table, you may also find it useful to select a set of routes from the BGP table based on certain criteria. The Cisco IOS BGP command toolkit includes the following options to select specific routes from the BGP table:

- **show ip bgp *network mask longer-prefixes***: This command lists all the more specific prefixes present in the BGP table (including the prefix itself) that are contained in the prefix specified by the *network* and *mask* options.
- **show ip bgp neighbor *ip-address* routes**: This command lists all routes in the BGP table that were received from the neighbor specified by the *ip-address* option.
- **show ip bgp neighbor *ip-address* advertised-routes**: This command lists all routes in the BGP table that will be advertised to the neighbor specified by the *ip-address* option.
- **show ip bgp regexp *regular-expression***: This command selects all routes from the BGP table that have an AS path string that is matched by the specified regular expression.

For more information about regular expressions and how to match specific AS paths using regular expressions, consult the “Understanding Regular Expressions” section in the *Cisco IOS Configuration Fundamentals Configuration Guide* at:

http://www.cisco.com/en/US/docs/ios/fundamentals/configuration/guide/cf_cli-basics_ps6350_TSD_Products_Configuration_Guide_Chapter.html#wp1002051

Debug BGP Route Exchange

- The **debug ip bgp update** command can be used to display the transmission and reception of BGP routing updates.
- **Use with extreme care:** Restrict use of this command to only displaying output for the specific routes and neighbors that you are interested in.

```
IR01#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
IR01(config)#access-list 37 permit 172.17.76.0 0.0.3.255
IR01(config)#^Z
IR01#
IR01#debug ip bgp 192.168.224.254 updates 37
BGP updates debugging is on for access list 37 for neighbor 192.168.224.254 for
address family: IPv4 Unicast
IR01#clear ip bgp 192.168.224.254 soft
IR01#
Apr 29 06:36:57.549 PDT: BGP(0): 192.168.224.254 send UPDATE (format)
172.17.76.0/22, next 192.168.224.241, metric 0, path Local
Apr 29 06:36:57.553 PDT: BGP(0): 192.168.224.254 rcv UPDATE w/ attr: nexthop
192.168.224.254, origin i, originator 0.0.0.0, path 65525 64568, community ,
extended community
Apr 29 06:36:57.553 PDT: BGP(0): 192.168.224.254 rcv UPDATE about 172.17.76.0/22
-- DENIED due to: AS-PATH contains our own AS;
```

© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0-100

If you find expected route entries to be missing from the BGP table, or you doubt whether the router is sending specific routes to a neighbor, you can consider using the **debug ip bgp updates** command to display the processing of BGP updates by the router. However, this command can generate a large number of messages, especially if your BGP table carries many routes. Consequently, it has a high risk of disrupting the operation of the router. In production networks, utmost care should be taken when using this command and additional command options should be used to limit the command to the prefixes and neighbor that you are troubleshooting.

The example in the figure shows how to limit the output of the **debug ip bgp updates** command by specifying a neighbor and using an access list to select only certain prefixes.

To illustrate the procedure, the commands are listed interspersed with comments that explain the procedure and output.

```
IR01#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
IR01(config)#access-list 37 permit 172.17.76.0 0.0.3.255
IR01(config)#^Z
IR01#
```

An access list with number 37 is created. When used to filter BGP routes, this access list matches any prefix in the 172.17.76.0 – 172.17.79.0 IP range.

```
IR01#debug ip bgp 192.168.224.254 updates 37
BGP updates debugging is on for access list 37 for neighbor
192.168.224.254 for address family: IPv4 Unicast
```

The debug is enabled for neighbor 192.168.224.254 and access list 37. Only update messages transmitted to or received from neighbor 192.168.224.254 (that are permitted by access list 37) will be displayed.

```
IR01#clear ip bgp 192.168.224.254 soft
```

A “soft” clear of BGP neighbor 192.168.224.254 is issued. As opposed to a “hard” clear, this clear will not tear down and restart the session completely, but merely forces the routes between this router and the neighbor to be retransmitted.

```
IR01#
```

```
Apr 29 06:36:57.549 PDT: BGP(0): 192.168.224.254 send UPDATE (format)
172.17.76.0/22, next 192.168.224.241, metric 0, path Local
```


An update about prefix 172.17.76.0/22 is transmitted to neighbor 192.168.224.254. Note that both the neighbor and the prefix match the imposed restrictions.

```
Apr 29 06:36:57.553 PDT: BGP(0): 192.168.224.254 rcv UPDATE w/ attr:
nexthop 192.168.224.254, origin i, originator 0.0.0.0, path 65525
64568, community , extended community
```

```
Apr 29 06:36:57.553 PDT: BGP(0): 192.168.224.254 rcv UPDATE about
172.17.76.0/22 -- DENIED due to: AS-PATH contains our own AS;
```

An update about prefix 172.17.76.0/22 is received, but denied because the AS path attribute contains the AS (AS 64568) of this router.

Many more updates were sent between this router and its neighbor, but only updates that match the imposed restrictions were displayed, limiting the impact of the command.



If you find that a route is available in the BGP table, but not in the routing table, there are two possible explanations. Either BGP has not been able to select any of the paths as the best path, or it has selected a best path, but a competing route from a different source with a better administrative distance is present and has been installed in the routing table.

If BGP has not selected any of the paths as the best path, this failure will be clearly visible in the BGP table and clues about the cause of the best path selection failure can be gathered from the BGP table. For example, if none of the paths has a next hop that can be resolved in the IP routing table, the text “Inaccessible” will be displayed instead of the IGP metric to reach the next hop. If the BGP synchronization rule is causing a route not to be installed in the routing table, the text “not synchronized” will be displayed behind the route.

If a best path has been selected by BGP for the prefix, but not installed in the routing table due to the presence of a competing route with a better administrative presence, the route will be marked as a “RIB-failure” in the BGP table. To list all BGP routes that have not been installed in the routing table due to a RIB failure, use the **show ip bgp rib-failure** command.

Lab Debrief Notes

Use these notes sections to write down the primary learning points that are discussed during the Lab Debrief.

Lab 5-3: Alternate Solutions

[illegible]

[illegible]

[illegible]

Lab 5-3: Important Commands and Tools

[illegible]

Lab 5-3: References

If you need more information on the commands and their options, you can go to the following sections of <http://www.cisco.com>.

- Cisco Systems, Inc. *Cisco IOS IP Routing Protocols Command Reference*. San Jose, California, November 2008:
http://www.cisco.com/en/US/docs/ios/iproute/command/reference/irp_book.html
- Cisco Systems, Inc. IP Routing Troubleshooting TechNotes, Border Gateway Protocol (BGP):
http://www.cisco.com/en/US/tech/tk365/tsd_technology_support_troubleshooting_technotes_list.html#anchor1

Lab 5-4: Router Performance

Complete this lab activity to practice what you learned in the related module.

Activity Objective


In this activity, you will experience the challenges of troubleshooting various problems related to router performance. After completing this activity, you will be able to meet these objectives:

- Diagnose and resolve problems related to router performance
- Document troubleshooting progress, configuration changes, and problem resolution

Information Packet

The figure illustrates what you will accomplish in this activity.

Visual Objective



Required Resources

These are the resources and equipment that are required to complete this activity:

- One PC with Internet access per team member
- Four Cisco Catalyst 3560 Series Switches per team
- Six Cisco 1841 Integrated Services Routers per team
- Three Client PCs per team
- One Server per team

Job Aids

These job aids are available to help you complete the lab activity.

- Trouble tickets
- Troubleshooting log

Trouble Ticket Q: Problems with Connectivity

It is Monday morning, and as soon as you enter your office in headquarters, you receive a call from your colleague from the branch office. He tells you that client applications report errors while connecting to the corporate server for large file transfers. In addition, he has received complaints from sales staff about slow or no connection at all to partner servers outside the corporate network, which reside in the IP address blocks 10.16.0.0/12 and 10.32.0.0/12.

You remember a colleague of yours, working on night shifts in that office, who currently studies Cisco CCNP® materials and has full access to branch and central office routers. What could he have done? As usual, you must rely on your great troubleshooting skills to resolve this problem quickly.

Your intuition tells you that you have to deal with the connectivity to networks 10.16.0.0/12 and 10.32.0.0/12 first.

Note	In this trouble ticket, you should make configuration changes on routers BRO1, BRO2, CRO1, and CRO2 only.

Instructions

Together with your team members, create a troubleshooting plan to divide the work, assign each team member appropriate roles, and coordinate device access between team members. Together, work on Trouble Ticket Q to resolve the issues. Document your progress in the following Troubleshooting Log in order to help facilitate efficient communication within the team and to have an overview of your troubleshooting process for reference during the Lab Debrief discussions.

You are allowed a total of one hour to complete the trouble ticket. After one hour, the instructor will debrief the lab. The main objective for the troubleshooting labs in this course is to give you an opportunity to practice structured troubleshooting. Fixing the problems is secondary to practicing proper processes and procedures.

You can use the hidden Cisco IOS Software command **ttcp** to simulate a heavy traffic load instead of performing large file transfers between client CLT2 and server SRV1 to test. The Test TCP (TTCP) utility consists of a client and server side. Because this lab revolves around router performance, the switches can be used as test points.

Note	The ttcp command is a hidden, unsupported, privileged mode command. As such, its availability may vary from one Cisco IOS Software release to another, and may not exist in some releases. Some platforms, for instance, require the Cisco IOS Enterprise feature set in order to perform this activity. More information about this utility can be found at the following link: http://www.cisco.com/en/US/tech/tk801/tk36/technologies_tech_note09186a0080094694.shtml

To initiate a TTCP connection and generate test traffic from switch ASW1 at headquarters to switch BSW1 at the branch follow the procedure outlined here.

First, go to the device that will perform the server side of the connection, in this case switch BSW1, and issue the following command sequence:

```
BSW1#ttcp
transmit or receive [receive]:
receive packets asynchronously [n]:
perform tcp half close [n]:
receive buflen [32768]:
bufalign [16384]:
bufoffset [0]:
port [5001]:
sinkmode [y]:
rcvwndsize [32768]:
ack frequency [0]:
delayed ACK [y]:
show tcp information at end [n]:

ttcp-r: buflen=32768, align=16384/0, port=5001
rcvwndsize=32768, delayedack=yes tcp
```

As you can see from this example, you issue the **ttcp** command and then accept all default values. After you complete this process, TTCP will wait for the transmitting (client) side to make a connection.

Now, go to the device that will perform the client side of the connection, in this case switch ASW1, and issue the following command sequence:

```
ASW1#ttcp
transmit or receive [receive]: transmit
Target IP address: 10.1.1.163.193
calculate checksum during buffer write [y]:
perform tcp half close [n]:
send buflen [32768]:
send nbuf [2048]:
bufalign [16384]:
bufoffset [0]:
port [5001]:
sinkmode [y]:
buffering on writes [y]:
show tcp information at end [n]:

ttcp-t: buflen=32768, nbuf=2048, align=16384/0, port=5001 tcp
-> 10.1.1.163.193
ttcp-t: connect
```

As you can see from the example, you should specify that this side is the transmitting side of the connection because the default is to run in receive mode. In addition, the IP address of the receiving side (BSW1) needs to be specified to initiate the connection. TTCP now starts transferring data between the transmitting and receiving side. You cannot issue any commands on the console of either device until the transfer finishes. The transmission can be interrupted at any point in time from the transmitting side using the key combination **Ctrl-Shift-6**.

Caution This utility can very easily overload a router with test traffic. It is not recommended that you use it on production devices. Read the TTCP documentation before using the TTCP utility.

A recommended approach to this lab is to follow a troubleshooting process that includes the following high-level tasks:

- Generate test traffic using the **ttcp** command as described previously or by transferring files between client PC CLT2 and server SRV1.
- Use the **ping** command to measure the performance between headquarters and the branch office, for example ping from client PC CLT2 to server SRV1.
- Examine the key performance indicators, such as interfaces, CPU, and memory on the routers and watch for symptoms associated with performance problems.
- Examine the routers for features and configurations that deviate from the baseline configurations and attempt to find the root cause of the problems.
- Address the issues causing the performance problems and test to verify that the performance has improved.

Lab Setup

The instructor will provide you with directions to prepare the lab equipment for this lab. After the instructor indicates that the lab is fully prepared, you are ready to start troubleshooting.

Troubleshooting Log

Use this log to document your actions and results during the troubleshooting process.

Trouble Ticket Q

Your task is to diagnose the performance problems on the network and if possible, resolve them.

Note Refer to the Activity Verification items at the end of this log to verify that you have successfully completed this task and lab.

[illegible]

Lab Debrief Notes

Use these notes sections to write down the primary learning points that are discussed during the Lab Debrief.

Lab 5-4: Alternate Solutions

[illegible]

[illegible]

[illegible]

Lab 5-4: Important Commands and Tools

[illegible]

Lab 6-1: Introduction to Network Security

Complete this lab activity to practice what you learned in the related module.

Activity Objective


In this activity, you will troubleshoot various problems related to network security. After completing this activity, you will be able to meet these objectives:

- Diagnose and resolve problems related to router and switch access
- Diagnose and resolve problems related to packet filtering using access lists
- Document troubleshooting progress, configuration changes, and problem resolution

Information Packet

The figure illustrates what you will accomplish in this activity:

Visual Objective



Required Resources

These are the resources and equipment that are required to complete this activity:

- One PC with Internet access per team member
- Four Cisco Catalyst 3560 Series Switches per team
- Six Cisco 1841 Integrated Services Routers per team
- Three client PCs per team
- One server per team

Job Aids

These job aids are available to help you complete the lab activity.

- Trouble tickets
- Troubleshooting log

Introduction: Increased Network Security

Your organization has recently experienced a number of security incidents. Luckily, the damage was relatively limited, but it has served to warn you that there is a problem. A security officer has been appointed, who is busy doing risk analysis and is drafting a security policy. It has also been decided to start implementing parts of the policy to address the most serious issues and refine the policy and implementation as an ongoing process.

The following is a summary of the sections from the security policy that are relevant to the network:

- Any traffic that is not specifically allowed on the network should be blocked.
- The network management team is responsible for deciding to what level of detail the rules should be implemented. A balance should be struck between security and manageability and scalability of the solution. The implementation will be regularly audited and reviewed for compliance with the policies.
- Management and control traffic should be permitted, but only as necessary.
- Device access should be authenticated against a central username and password database.

The following statements provide more detail about the rules that should be applied to the various zones of the network.

Guest access should be provided to external contractors, partners and other guests based on the following rules:

- Guests are only allowed to access the Internet and should not have access to any internal machines.
- Guests are not allowed to host any of the well-known services (TCP and UDP ports below 1024) on their machines.
- Any other TCP- or UDP-based services and applications can be used.

Access from the office LANs is limited by the following rules:

- Sending SMTP-based email is not allowed. All email should be sent through the corporate mail server.
- Office users are not allowed to host any of the well-known services (TCP and UDP ports below 1024) on their machines.
- Any other TCP- or UDP-based services and applications can be used.

Access from the branch offices to server SRV1 is restricted to the following:

- All traffic necessary to provide name services to the clients
- All traffic necessary to provide file services to the clients
- All management traffic from network devices

Internet access is restricted as follows:

- Internet traffic is restricted to HTTP- and HTTPS-based services.
- Internet access for users from the Branch 1 Office VLAN is suspended temporarily. Several of the security incidents originated from this LAN and for that reason it has been decided to deny Internet access to these users for the moment.

This policy is still under development. Further restrictions may be added at any point. Exceptions to the policy can be made after approval from the security officer.

Note	The network engineering team is still busy implementing these policies. Not all points have been fully implemented on all routers. However, the network support team has been instructed that any changes that they make during maintenance and troubleshooting should be validated against these policies.

Because of pressure from management to increase the security on the network, normal implementation and change procedures have been relaxed to speed up the implementation. Over the weekend, the engineering team has been working very hard to implement several measures, but they did have only limited time for testing. Amongst the features that they implemented are various access lists and a TACACS+ server for centralized authentication services.

Your team should be ready to address and resolve any issues that may show up when people start coming into the office on Monday.

Trouble Ticket R: Internet Not Reachable from Client PC CLT1

One of the headquarters office users who uses client PC CLT1 has reported that he cannot browse to <http://www.isp3.local>.

Your task is to diagnose this problem and restore connectivity to <http://www.isp3.local> and any other connectivity that this user is entitled to according to the security policy. You must ensure that no other traffic is allowed other than the traffic specifically permitted by the security policy.

Trouble Ticket S: Internet Not Reachable from Client PC CLT3

An external contractor in the branch office who uses client PC CLT3 has reported that he cannot browse to <http://www.isp3.local>.

Your task is to diagnose this problem and restore connectivity to <http://www.isp3.local>. You must ensure that no other traffic is allowed other than the traffic specifically permitted by the security policy.

Trouble Ticket T: Client PC CLT2 Has No Network Connectivity

One of the branch office users who uses CLT PC CLT2 reports that the network is broken and he cannot connect anywhere. He reports that not even **ping** works. (This user worked as a part-time server administrator in his previous job and thinks he knows a lot about networking.)

Your task is to diagnose this problem and restore all connectivity that this user is entitled to according to the security policy.

Instructions

Together with your team members, create a troubleshooting plan to divide the work, assign each team member appropriate roles, and coordinate device access between team members. Together, work on Trouble Tickets R, S, and T to resolve the issues. Document your progress in the following Troubleshooting Logs in order to help facilitate efficient communication within the team and to have an overview of your troubleshooting process for reference during the Lab Debrief discussions.

You are allowed a total of 45 minutes to complete as many of the trouble tickets as you can. After 45 minutes, the instructor will debrief the lab and review all trouble tickets and their solutions. The main objective for the troubleshooting labs in this course is to give you an opportunity to practice structured troubleshooting. Fixing the problems is secondary to practicing proper processes and procedures.

Lab Setup

The instructor will provide you with directions to prepare the lab equipment for this lab. After the instructor indicates that the lab is fully prepared, you are ready to start troubleshooting.

Troubleshooting Log

Use this log to document your actions and results during the troubleshooting process.

Trouble Ticket R

Your task is to diagnose the Internet connectivity problem experienced by the user on client PC CLT1 and restore connectivity to <http://www.isp3.local> and any other connectivity that this user is entitled to according to the security policy. Ensure that no other traffic is allowed other than the traffic specifically permitted by the security policy.

Note Refer to the Activity Verification items at the end of this log to verify that you have successfully completed this task.

[illegible]

Troubleshooting Log

Use this log to document your actions and results during the troubleshooting process.

Trouble Ticket S

Your task is to diagnose the problem experienced by the user on client PC CLT3 and restore connectivity to <http://www.isp3.local>. Ensure that no other traffic is allowed other than the traffic specifically permitted by the security policy.

Note Refer to the Activity Verification items at the end of this log to verify that you have successfully completed this task.

[illegible]

Troubleshooting Log

Use this log to document your actions and results during the troubleshooting process.

Trouble Ticket T

Your task is to diagnose the problems experienced by the user on client PC CLT2 and restore all connectivity that this user is entitled to according to the security policy.


Note Refer to the Activity Verification items at the end of this log to verify that you have successfully completed this task.

[illegible]

Lab 6-1: Sample Troubleshooting Flows

Troubleshooting TACACS+


The figure illustrates an example of a method that you could follow to diagnose and resolve problems related to the TACACS+ protocol, which can be used for AAA purposes.



When you have problems logging in remotely on a router or switch, one of the first things to check is the type of authentication that is used on the device. If the authentication that is used is TACACS+, you need to check the configuration of both the TACACS+ server and the router or switch to see if the basic parameters match.

Sample TACACS+ Troubleshooting Flow

- The **show running-config** command displays the configuration of both the generic AAA framework and the specific TACACS+ parameters on the device.



© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0-105

The **show running-config** command shows the type of authentication that is used by the device; in this case, it is TACACS+.

In the running configuration, you can also find the mandatory commands that are configured on the device to communicate with the TACACS+ server:

```
tacacs-server host
tacacs-server key
```

The figure shows an example configuration for TACACS+. When you have verified these items in the configuration of the device and you still cannot authenticate, you need to check the following:

- Ensure that you have configured the client device on the TACACS+ server.
- Ensure that the username and password that you are trying to use have been configured on the TACACS+ server.
- Confirm that you have IP connectivity from the client device to the TACACS+ server.
- Confirm that the keys configured on the server and the client match.

Sample TACACS+ Troubleshooting Flow (Cont.)

- Use **debug aaa authentication** to troubleshoot authentication.
- This output shows a successful authentication.

```
Router#debug aaa authentication
AAA Authentication debugging is on
*May 26 12:15:10.039: AAA/BIND(00000005): Bind i/f
*May 26 12:15:10.043: AAA/AUTHEN/LOGIN (00000005): Pick method list 'default'
*May 26 12:15:16.103: AAA: parse name=tty322 idb type=-1 tty=-1
*May 26 12:15:16.103: AAA: name=tty322 flags=0x11 type=5 shelf=0 slot=0 adapter=0
port=322 channel=0
*May 26 12:15:16.103: AAA/MEMORY: create_user (0x4691F25C) user='vesko' ruser='NULL'
ds0=0 port='tty322'
rem_addr='192.168.1.2' authn_type=ASCII service=ENABLE priv=15 initial_task_id='0',
vrf= (id=0)
*May 26 12:15:16.103: AAA/AUTHEN/START (3485450623): port='tty322' list=''
action=LOGIN service=ENABLE
*May 26 12:15:16.103: AAA/AUTHEN/START (3485450623): non-console enable - default to
enable password
*May 26 12:15:16.103: AAA/AUTHEN/START (3485450623): Method=ENABLE
*May 26 12:15:16.103: AAA/AUTHEN (3485450623): Status=GETPASS
*May 26 12:15:16.843: AAA/AUTHEN/CONT (3485450623): continue_login (user='(undef)')
*May 26 12:15:16.843: AAA/AUTHEN (3485450623): Status=GETPASS
*May 26 12:15:16.843: AAA/AUTHEN/CONT (3485450623): Method=ENABLE
*May 26 12:15:16.851: AAA/AUTHEN (3485450623): Status=PASS
```

© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0-106

There are two debug commands that are often used to troubleshoot authentication problems. In the figure, you see the output of **debug aaa authentication** command. The specific example in the figure shows a successful authentication.

Sample TACACS+ Troubleshooting Flow (Cont.)

- Use **debug tacacs** to troubleshoot TACACS+-related problems.
- This output shows a problem with the TACACS+ key.

```
Router#debug tacacs
TACACS access control debugging is on
<... output omitted ...>
*May 26 12:22:17.107: TPLUS: received bad AUTHEN packet: length = 6, expected
59464
*May 26 12:22:17.107: TPLUS: Invalid AUTHEN packet (check keys).
*May 26 12:22:17.107: TPLUS (00000006)/0/REQ_WAIT/46430CDC: timed out
*May 26 12:22:17.111: TPLUS: Authentication start packet created for 6()
*May 26 12:22:17.111: TPLUS (00000006)/0/REQ_WAIT/46430CDC: timed out, clean up
*May 26 12:22:17.111: TPLUS (00000006)/0/46430CDC: Processing the reply packet
*May 26 12:22:19.111: TPLUS: Queuing AAA Authentication request 6 for
processing
*May 26 12:22:19.111: TPLUS: processing authentication start request id 6
*May 26 12:22:19.111: TPLUS: Authentication start packet created for 6()
*May 26 12:22:19.111: TPLUS: Using server 192.168.1.5
*May 26 12:22:19.119: TPLUS: received bad AUTHEN packet: length = 6, expected
48883
*May 26 12:22:19.119: TPLUS: Invalid AUTHEN packet (check keys).
*May 26 12:22:19.119: TPLUS (00000006)/0/REQ_WAIT/46430CDC: timed out
```


© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0-107

Another useful command is **debug tacacs**. In the output in the figure, you can see this command as it displays an unsuccessful authentication attempt. The most important things to notice in this output are the TACACS+ server IP address (which is 192.168.1.5 in this case) and the fact that there is an invalid authentication packet (which includes a suggestion to check the keys used between server and client).

Troubleshooting Console Connections

The figure illustrates an example of a method that you could follow to diagnose and resolve problems related to accessing the console port of a device to access the Cisco IOS Software CLI.



You only need to check a few things when you have problems accessing a device via the console port. The terminal settings are one of the first things that you should check. If the settings match the settings of your terminal emulation and you are working on a lab environment, you could consider reloading the device. However, in a production environment, reloading is generally not a feasible approach.

If you cannot reload, or if reloading did not help, you can also try to access the device via Telnet or SSH. If you succeed, you will have the option to verify the console settings and to address any potential problems that you might spot in the configuration.

If you do not have remote access to the device, you can try the procedure used to recover a lost password in order to get access to the device without loading the startup configuration. If this attempt works, you can check the startup configuration of the device and fix any potential configuration issues on the console.

Of course, the problem might also be caused by a hardware issue such as a bad cable or failed console port.

Verify Console Configuration

- The **show running-config** command displays information about the console 0 configuration.
- The **show line console 0** command gives you more information about the operational status of the console line.

```
Router#show running-config | begin line con 0
line con 0
  exec-timeout 0 10

Router#show line console 0
  Tty Line Typ Tx/Rx A Modem Roty AccO AccI Uses Noise Overruns Int
*    0    0 CTY - - - - - 0 1 0/0 -
Line 0, Location: "", Type: ""
Length: 24 lines, Width: 80 columns
Baud rate (TX/RX) is 9600/9600, no parity, 2 stopbits, 8 databits
Status: PSI Enabled, Ready, Active, Automore On
Timeouts:      Idle EXEC Idle Session Modem Answer Session Dispatch
               00:00:10 never          none not set

<... output omitted ...>
Time since activation: 00:00:14
Editing is enabled.
History is enabled, history size is 20.
```


© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0-109

You can use the **show running-config** or **show line console 0** commands in order to check the settings of the console line. In the example in the figure, you can see that the exec-timeout of the device is set to 0 minutes and 10 seconds, which will log you out of the device after being idle for 10 seconds on the console. Clearly, this setting will be quite disruptive. The default exec-timeout setting is 10 minutes.

Troubleshooting the ICMP

The figure illustrates an example of a method that you could follow to diagnose and resolve problems related to the ICMP.



When a host fails to respond to a ping, this failure does not always mean that you do not have IP connectivity to that host. It is possible that an access list is blocking ICMP messages in between the source and the destination. By using the **ping** and **traceroute** commands, you may be able to determine where the problem is located.

Sample ICMP Troubleshooting Flow

The **ping** and **tracert** on the network host give you useful information regarding on which hop the problem might be. Then you can run **debug ip icmp** on the suspected device.

```
C:\Documents and Settings\Administrator>ping 10.1.152.1
```

```
Pinging 10.1.152.1 with 32 bytes of data:
```

```
Reply from 10.1.160.125: Destination net unreachable.  
Reply from 10.1.160.125: Destination net unreachable.  
Reply from 10.1.160.125: Destination net unreachable.  
Reply from 10.1.160.125: Destination net unreachable.
```

The **ping** and **tracert** commands give you a hint to continue troubleshooting on 10.1.160.125.

```
C:\Documents and Settings\Administrator>tracert 10.1.152.1
```

```
Tracing route to srv1.tshoot.local [10.1.152.1]  
over a maximum of 30 hops:
```

```
  1  10.1.160.125  reports: Destination net unreachable.
```

```
Trace complete.
```

Debugging 10.1.160.125 shows that this router is sending administratively prohibited messages to the source.

```
BRO2#debug ip icmp
```

```
May 26 06:06:48.306 PDT: ICMP: dst (10.1.152.1) administratively prohibited unreachable sent to 10.1.160.65  
May 26 06:06:49.310 PDT: ICMP: dst (10.1.152.1) administratively prohibited unreachable sent to 10.1.160.65  
May 26 06:06:50.310 PDT: ICMP: dst (10.1.152.1) administratively prohibited unreachable sent to 10.1.160.65  
May 26 06:06:51.310 PDT: ICMP: dst (10.1.152.1) administratively prohibited unreachable sent to 10.1.160.65
```

© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0-111

In the figure, **ping** and **tracert** commands are issued on a network host. A report from 10.1.160.125 indicates that the destination network is unreachable; giving you a hint to check the configuration on this device.

You can check several things on 10.1.160.125:

- You can use the command **debug ip icmp**. This command shows that this router is sending an “administratively prohibited unreachable” message to 10.1.160.65, which is the source of the ping. This message is another clue that you should check for configured access lists on the router.

Sample ICMP Troubleshooting Flow (Cont.)

The **show ip interface** and **show access-lists** commands display important information.

```
BR01#show ip interface
FastEthernet0/0 is up, line protocol is up
Internet address is 10.1.192.2/30
Broadcast address is 255.255.255.255
Address determined by non-volatile memory
MTU is 1500 bytes
Helper address is not set
Directed broadcast forwarding is disabled
Multicast reserved groups joined: 224.0.0.10
Outgoing access list is icmp
Inbound access list is not set
Proxy ARP is enabled
ICMP redirects are always sent
.... / .... output truncated

BR01#show access-lists
Extended IP access list icmp
10 deny tcp 10.1.160.64 0.0.0.63 any eq smtp
20 permit tcp 10.1.160.192 0.0.0.63 gt 1024 any
30 permit udp 10.1.160.192 0.0.0.63 gt 1024 any
40 permit icmp 10.1.160.64 0.0.0.63 host 10.1.152.1 host-unreachable
50 permit icmp 10.1.160.64 0.0.0.63 host 10.1.152.1 host-unknown
60 permit icmp 10.1.160.64 0.0.0.63 host 10.1.152.1 net-unreachable
70 permit icmp 10.1.160.64 0.0.0.63 host 10.1.152.1 network-unknown
80 deny icmp any any (12 matches)
90 permit ip any any (21 matches)
```

An access list called "icmp" is applied outbound to interface FastEthernet 0/0.

A line does not appear in access list icmp that allows echo messages to reach host 10.1.152.1. You may see the counter in line 80 incrementing if you ping 10.1.152.1. Add **75 permit icmp 10.1.160.64 0.0.0.63 host 10.1.152.1 echo.**


© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0-112

- You can use the command **show ip interface** to check for access lists that are being configured on router interfaces. In the figure, an access list is configured on interface FastEthernet 0/0 in the outbound direction.
- The command **show access-lists** gives you more detailed information about the access lists that are configured on the device. You may also see the number of packets that match the different statements.

Troubleshooting DNS

The figure illustrates an example of a method that you could follow to diagnose and resolve problems related to the DNS.



Very often problems with an Internet connection are caused by DNS problems. Most often that means connection to the DNS server or servers has been lost, or the gateway router is not configured for DNS. You may use the **ping** or **telnet** commands in order to troubleshoot a DNS problem.

Sample DNS Troubleshooting Flow

The **ping** and **telnet** commands give you useful information for troubleshooting. In the following outputs you can see the network host and network router unsuccessfully trying to reach a web page. They are unsuccessful because of a DNS problem.

Output from host:

```
C:\Documents and Settings\Administrator>ping www.isp3.local
Ping request could not find host www.isp3.local. Please check the name
and try again.
```

Output from router:

```
BRO1#ping www.isp3.local

Translating "www.isp3.local"...domain server (10.1.152.1)
% Unrecognized host or address, or protocol not running.

BRO1#telnet 172.34.224.1 53
Trying 172.34.224.1, 53 ...
% Destination unreachable; gateway or host down
```

© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0—114

In the output, the network host and network router is trying to reach a web page unsuccessfully. They are unsuccessful because of a DNS problem. The fact that the router is showing the IP address of the name server (10.1.152.1) implies that the name is configured on the router, but there is no connection to the server.

Sample DNS Troubleshooting Flow (Cont.)

The **ping** and **telnet** commands give you useful information for troubleshooting. In the following outputs you can see the network host and network router successfully try to reach a web page.

Output from host:

```
C:\Documents and Settings\Administrator>ping www.isp3.local

Pinging www.isp3.local [172.34.224.1] with 32 bytes of data:

Reply from 172.34.224.1: bytes=32 time=10ms TTL=122
Reply from 172.34.224.1: bytes=32 time=10ms TTL=122
```

Output from router:

```
BRO1#ping www.isp3.local

Translating "www.isp3.local"...domain server (10.1.152.1) [OK]

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.34.224.1, timeout is 2 seconds:
!!!!

BRO1#telnet 172.34.224.1 53
Trying 172.34.224.1, 53 ... Open
```

© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0—115

Here you can see a successful connection to the DNS server. Using Telnet to port 53 shows the line **Trying 172.34.224.1, 53 ... Open**. The word *Open* means that the connection is successful. Using the **telnet** command with the port is a useful troubleshooting tool.

Sample DNS Troubleshooting Flow (Cont.)

Use **show running config** command to check DNS configuration on the router. Use **show access-lists** for statements that are blocking DNS.

```
CRO1#show running-config
.....
ip domain lookup
...
ip name-server 10.1.152.1
.....

CRO1#show access-lists
Extended IP access list dns
 10 deny icmp 10.1.0.0 0.0.255.255 host 10.1.252.1
 20 permit icmp 10.1.0.0 0.0.255.255 any
 30 permit udp 10.1.0.0 0.0.255.255 host 10.1.252.1 eq domain
 40 deny udp 10.1.0.0 0.0.255.255 any
 50 deny tcp 10.1.0.0 0.0.255.255 host 10.1.152.1 eq ftp
 60 deny tcp 10.1.160.64 0.0.0.63 any eq smtp
 70 deny icmp 10.1.160.64 0.0.0.63 host 10.1.152.9 echo
 80 permit tcp 10.1.0.0 0.0.255.255 gt 1023 host 10.1.252.1
```

Enables DNS-based host name-to-address translation. This command is enabled by default.

Specifies the address of one or more name servers.


The third and eighth lines of the access list have the wrong IP address and need to be fixed to permit DNS queries to the DNS server 10.1.152.1.

When the **ping** and **telnet** commands show unsuccessful attempts to reach the DNS server, you need to check the running configuration of the router for DNS configuration and for available access lists that may block DNS queries. The output shows the router is properly configured to use DNS and shows the IP address of the DNS server.

The configured access list shows a line that has the wrong IP address and needs to be fixed to permit DNS queries to the DNS server 10.1.152.1.

Troubleshooting HTTP

The figure illustrates an example of a method that you could follow to diagnose and resolve problems related to web services using HTTP.



Troubleshooting HTTP is very similar to troubleshooting DNS. You may use the same **ping** and **telnet** commands; just change the port number to 80 (http).

You will also need to check the connectivity and configuration of the DNS server because a lost connection to a DNS server is very often the reason the HTTP connection is lost. You may use the steps shown in the figure to troubleshoot DNS.

Sample HTTP Troubleshooting Flow (Cont.)

The **ping** and **telnet** commands give you useful information for troubleshooting. In the following outputs you can see the network host and network router trying to reach a web page unsuccessfully.

Output from host:

```
C:\Documents and Settings\Administrator>ping www.isp3.local
Ping request could not find host www.isp3.local. Please check
the name and try again.
```

Output from router:

```
BR01#ping www.isp3.local

Translating "www.isp3.local"...domain server (10.1.152.1)
% Unrecognized host or address, or protocol not running.

BR01#telnet 172.34.224.1 80
Trying 172.34.224.1, 80 ...
% Destination unreachable; gateway or host down
```

© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0-118

The output from the **ping** and **telnet** commands should be already familiar to you from the DNS troubleshooting section.

Sample HTTP Troubleshooting Flow (Cont.)

The **ping** and **telnet** commands give you useful information for troubleshooting. In the following outputs you can see a network host and network router successfully trying to reach a web page via ping or Telnet to port 80.

Output from host:

```
C:\Documents and Settings\Administrator>ping www.isp3.local

Pinging www.isp3.local [172.34.224.1] with 32 bytes of data:

Reply from 172.34.224.1: bytes=32 time=10ms TTL=122
Reply from 172.34.224.1: bytes=32 time=10ms TTL=122
```

Output from router:

```
BR01#ping www.isp3.local

Translating "www.isp3.local"...domain server (10.1.152.1) [OK]

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.34.224.1, timeout is 2 seconds:
!!!!

```

```
BR01#telnet 172.34.224.1 80
Trying 172.34.224.1, 80 ... Open
```


© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0-119

The successful ping or telnet to port 80 of a web page does not ensure that you have HTTP connectivity. If you can resolve the web page to IP address, you have a proper DNS configuration. If you still cannot open the web page, an access list is probably blocking the HTTP protocol.

Troubleshooting Access Lists


The figure illustrates an example of a method that you could follow to diagnose and resolve problems related to access lists.



Many things could go wrong or not exactly as desired when you configure and apply access lists. Here are the most important ones to focus on:

- One of the first things that you must do is to know what you want the access list to do. Consider which networks, hosts, or protocols that you want to permit or deny. Then decide if you should use a standard or extended access list. Using a named access list will give you more flexibility when you change statements and more visibility to the purpose of the access list. Be careful when replacing or changing statements and be aware of the effects that a change may cause.
- Be careful when selecting the source and destination networks, hosts, or protocol numbers. Sometimes you might have selected the proper direction of the networks, but mismatched the destination and source port numbers, which makes the access list ineffective.
- The order of the statements is very important. Cisco IOS Software checks the statements from beginning to end and does not check further if there is match. Therefore, you should be careful not to permit or deny something that you already have configured in a more detailed or compact statement that precedes the new statement.
- Do not forget that on the end of every access list there is an explicit *deny* statement. So verify that the networks or protocols you want to permit are included in some permit statements in the lines.
- Carefully choose the wildcard mask. An incorrect wildcard mask might lead to permitting unwanted networks or hosts or denying wanted networks or hosts.

Sample Access Lists Troubleshooting Flow



© 2009 Cisco Systems, Inc. All rights reserved.


TSHOOT v1.0-121

After you have correctly configured the access list, it is time to apply it to the interface. The access list does not go into effect until it is applied. You need to select carefully the router that you want to apply the access list to, because the correct selection is important for the proper and effective functioning of your network.

Next, choose the appropriate interface or interfaces to apply the access list to. In redundant topologies, you might block one interface but permit the traffic through another interface that might be unwanted.

The next important thing to check is the direction that you want the access list to work on the interface. As stated previously, you might apply one access list in the inbound direction and one in the outbound direction. The direction is very important for the proper functioning of the access list.

Sample Access Lists Troubleshooting Flow



© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0-122

Another important thing that you must do when you configure access lists is to be sure that you permit or deny only the intended traffic. That means that if you want to permit some networks, be careful not to permit other networks that should be restricted. In addition, if you want to deny some networks, ensure that you are not also denying other normally permitted networks.

A lost connection or failed application does not always mean that there is a problem with the access lists. If there is a problem with an application, check the application itself.

Before applying the access lists, be aware of the possible effects. For example if you have a remote connection to a device, it is possible to lose the connection after you apply the access list. In addition, it is possible for an access list to cause a network outage if it is configured or applied incorrectly.

Verify Configuration of Access Lists

The **show access list** displays all access lists that are configured on the device.

```
IR02#show access-lists
Standard IP access list INTERNAL-NETWORKS
 10 permit 10.1.128.0, wildcard bits 0.0.127.255 (4 matches)
```

The **show ip interfaces *intf-id*** command displays essential interface information and the access lists that were applied.

```
IR02#show ip interface FastEthernet 0/1
FastEthernet0/1 is up, line protocol is up
 Internet address is 172.24.244.81/29
 Broadcast address is 255.255.255.255
 Address determined by non-volatile memory
 MTU is 1500 bytes
 Helper address is not set
 Directed broadcast forwarding is disabled
 Multicast reserved groups joined: 224.0.0.10
 Outgoing access list is INTERNAL-NETWORKS
 Inbound access list is not set
```

© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0—123

The figure shows two important commands that enable you to configure and troubleshoot Cisco IOS access lists.

Lab Debrief Notes

Use these notes sections to write down the primary learning points that are discussed the Lab Debrief.

Lab 6-1: Alternate Solutions

[illegible]

[illegible]

[illegible]

[illegible]

Lab 6-1: References

If you need more information on the commands and their options, you can go to the following sections of <http://www.cisco.com>.

- Cisco Systems, Inc. *Cisco IOS Security Command Reference*. San Jose, California, July 2009: http://www.cisco.com/en/US/docs/ios/security/command/reference/sec_book.html

Lab 6-2: Cisco IOS Security Features

Complete this lab activity to practice what you learned in the related module.

Activity Objective


In this activity, you will troubleshoot various problems related to network security. After completing this activity, you will be able to meet these objectives:

- Diagnose and resolve problems related to switch-based Cisco IOS security features.
- Diagnose and resolve problems related to router-based Cisco IOS security features.
- Document troubleshooting progress, configuration changes, and problem resolution

Information Packet

The figure illustrates what you will accomplish in this activity:

Visual Objective



Required Resources

These are the resources and equipment that are required to complete this activity:

- One PC with Internet access per team member
- Four Cisco Catalyst 3560 Series Switches per team
- Six Cisco 1841 Integrated Services Routers per team
- Three client PCs per team
- One server per team

Job Aids

These job aids are available to help you complete the lab activity.

- Trouble tickets
- Troubleshooting log

Introduction: Improving Network Security

As part of the ongoing process of improving network security, the security officer at your organization has decided that a number of additional security measures need to be taken to improve the security of the infrastructure. The following measures will be implemented according to the security policy:

- Implementation of SSHv2 as the only allowed remote access method for all routers and switches.
- Implementation of stateful packet inspection on routers IRO1 and IRO2.
- Protection from rogue DHCP servers on all VLANs. The office VLAN in the branch office (VLAN 17) will be implemented first as a pilot. After successful implementation of this VLAN, the solution will be rolled out on all other VLANs.
- Protection against users plugging in switches, which might attempt to become the root for the STP.
- Protection against MAC address flooding attacks on all switches.
- Added protection against worms by blocking specific TCP ports (starting with ports 25 and 135) and unnecessary broadcast and multicast traffic on all VLANs. The guest VLAN in the branch office (VLAN 19) will be implemented first as a pilot. After successful implementation of this VLAN, the solution will be rolled out on all other VLANs.
- The restriction on Internet access for users in the Branch 1 Office VLAN has been lifted and they will all be allowed to access the Internet again.

Again, it has been decided to implement these features over the weekend, but this time you will be allowed more time for testing on Sunday, in order not to disrupt the business when users start coming into the office on Monday. You will be allowed to roll back the implementation of any features that you cannot get to work and would disrupt the business on Monday. However, you will need to have a good explanation of the reason why you decided to roll back the changes and hold them off until the next scheduled maintenance interval.

The engineering team has just finished their implementation and you are ready to start your testing. You send some of your junior team members out to start performing some initial tests and report any issues that they find. Soon the first trouble tickets start coming in.

Trouble Ticket U: No Connectivity from Client PCs CLT2 and CLT3

The first report that you receive is that both client PCs CLT2 and CLT3 have no network connectivity at all. According to your colleague, the PCs cannot even obtain an IP address.

Your task is to diagnose this problem and restore any connectivity that the users on these PCs are entitled to according to the security policy. Ensure that no other traffic is allowed other than the traffic specifically permitted by the security policy. If you do not succeed in resolving the problem, you are allowed to roll back any newly implemented security features as necessary.

However, the security features that were implemented during earlier security implementations should not be removed.

Trouble Ticket V: No Connectivity from Client PC CLT1

The second report you receive concerns client PC CLT1. This PC has also lost all network connectivity. This is remarkable, because client PC CLT1 is not in any of the VLANs that the rogue DHCP and worm protection measures have been implemented on.

Your task is to diagnose this problem and restore any connectivity that the user on client PC CLT1 is entitled to according to the security policy. Ensure that no other traffic is allowed other than the traffic specifically permitted by the security policy. If you do not succeed in resolving the problem, you are allowed to roll back any newly implemented security features as necessary. However, the security features that were implemented during earlier security implementations should not be removed.

Trouble Ticket W: No Connectivity to server SRV1

It looks like your team will have a lot of work to do this afternoon. The next report that you receive involves server SRV1. Your colleague reports that he cannot reach server SRV1 from any of the switches or routers. Because this server is the DNS server for your network, this implies that you will not be able to connect to any network device by name until this problem is resolved.

You remember that the NIC of this server was replaced just a week ago and you wonder if the replacement has anything to do with this problem.

Your task is to restore network connectivity to server SRV1 (within the constraints laid down in the security policy). If you do not succeed in resolving the problem, you are allowed to roll back any newly implemented features as necessary. However, the security features that were implemented during earlier security implementations should not be removed.

Trouble Ticket X: Lost Remote Connectivity to All Routers

To roll out SSHv2 on all routers and switches, one of your colleagues from engineering has created a small script, which executed the following commands on all routers:

```
ip ssh version 2
ip ssh source-interface Loopback0
line vty 0 15
transport input ssh
```

On the switches, a similar script was executed, setting the source interface to the management VLAN instead of a loopback interface:

```
ip ssh version 2
ip ssh source-interface Vlan128
line vty 0 15
transport input ssh
```

After execution of these configuration commands, the script has saved the configurations on all devices.

When your colleague tries to test the connectivity to the routers and switches, he notices that he cannot connect to the devices using SSH. Because he cannot figure out what has happened, he asks for your help.

It is your task to diagnose this problem and ensure SSH connectivity to all routers and switches.

Trouble Ticket Y: Port Security Problems on Switch BSW1

The same colleague who has just informed you about the SSH problem, also asks for your help with a different problem. A week ago, while he was preparing for the security implementation, he tried to create a macro on switch BSW1 that enables a number of security features. The macro itself worked without any issues, and the security features were applied to the port. However, during a failover test the next day, switch BSW1 lost connectivity to routers BRO1 and BRO2, disrupting all communications between the branch office and headquarters. He rolled back the changes and managed to restore connectivity. However, he would still like to be able to enable the security features using this macro and asks for your help in testing. He tells you that to enable the security features you should log in to switch BSW1 and execute the following commands:

```
interface range FastEthernet 0/1 - 24
macro apply SECURE-PORTS
```

Your colleague claims that these commands work as intended, but after a reload of router BRO1 or BRO2 problems start being experienced.

It is your task to verify and diagnose this problem and, if possible, resolve it. Ensure that failover between routers BRO1 and BRO2 works as intended by the network design.

Instructions

Together with your team members, create a troubleshooting plan to divide the work, assign each team member appropriate roles, and coordinate device access between team members. Together, work on Trouble Tickets U, V, W, and X and Y to resolve the issues. Document your progress in the following Troubleshooting Logs in order to help facilitate efficient communication within the team and to have an overview of your troubleshooting process for reference during the Lab Debrief discussions.

You are allowed a total of two and a half hours to complete as many of the trouble tickets as you can. After two and a half hours, the instructor will debrief the lab and review all trouble tickets and their solutions. The main objective for the troubleshooting labs in this course is to give you an opportunity to practice structured troubleshooting. Fixing the problems is secondary to practicing proper processes and procedures.

Lab Setup

The instructor will provide you with directions to prepare the lab equipment for this lab. After the instructor indicates that the lab is fully prepared, you are ready to start troubleshooting.

Troubleshooting Log

Use this log to document your actions and results during the troubleshooting process.

Trouble Ticket U

Your task is to diagnose the problems experienced on client PCs CLT2 and CLT3 and restore any connectivity that the users on these PCs are entitled to according to the security policy. Ensure that no other traffic is allowed other than the traffic specifically permitted by the security policy. If you do not succeed in resolving the problem, you are allowed to roll back any newly implemented security features as necessary. However, the security features that were implemented during earlier security implementations should not be removed.

Note Refer to the Activity Verification items at the end of this log to verify that you have successfully completed this task.

[illegible]

Troubleshooting Log

Use this log to document your actions and results during the troubleshooting process.

Trouble Ticket V

Your task is to diagnose the problems experienced on client PC CLT1 and restore any connectivity that the user on this PC is entitled to according to the security policy. Ensure that no other traffic is allowed other than the traffic specifically permitted by the security policy. If you do not succeed in resolving the problem, you are allowed to roll back any newly implemented security features as necessary. However, the security features that were implemented during earlier security implementations should not be removed.

Note Refer to the Activity Verification items at the end of this log to verify that you have successfully completed this task.

[illegible]

Troubleshooting Log

Use this log to document your actions and results during the troubleshooting process.

Trouble Ticket W

Your task is to restore network connectivity to server SRV1 (within the constraints laid down in the security policy). If you do not succeed in resolving the problem, you are allowed to roll back any newly implemented features as necessary. However, the security features that were implemented during earlier security implementations should not be removed.

Note Refer to the Activity Verification items at the end of this log to verify that you have successfully completed this task.

[illegible]

Troubleshooting Log

Use this log to document your actions and results during the troubleshooting process.

Trouble Ticket X

It is your task to diagnose the problems that users experienced when they attempted to connect to the routers. In addition, it is your task to ensure SSH connectivity to all routers and switches.

Note Refer to the Activity Verification items at the end of this log to verify that you have successfully completed this task.

[illegible]

Troubleshooting Log

Use this log to document your actions and results during the troubleshooting process.

Trouble Ticket Y

It is your task to verify and diagnose the failover problem experienced after applying the macro on switch BSW1 and, if possible, resolve it. Ensure that failover between routers BRO1 and BRO2 works as intended by the network design.


Note Refer to the Activity Verification items at the end of this log to verify that you have successfully completed this task.

[illegible]

Lab 6-2: Sample Troubleshooting Flows


Troubleshooting DHCP

The figure illustrates an example of a method that you could follow to diagnose and resolve problems related to DHCP.



Usually, you would start troubleshooting the Layer 1 connectivity between the adjacent network equipment and host. A typical symptom that could lead you to start examining Layer 1 connectivity would be that the Ethernet controller on the PC is disconnected.

Sample DHCP Troubleshooting Flow




© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0—126

The next step in troubleshooting should be verification of basic DHCP configuration. By default, the DHCP server and DHCP relay agent are enabled on the Cisco switches but are not configured. If the DHCP server and DHCP relay are enabled but not configured in your situation, enable the DHCP server and relay agent using **service dhcp** global configuration mode command.

In this situation, the DHCP server and the DHCP clients are on different networks or subnets, so you must configure the switch with the **ip helper-address address** interface configuration command. The general rule is to configure the command on the Layer 3 interface closest to the client. The address used in the **ip helper-address** command can be a specific DHCP server IP address (SRV1). Check the DHCP packets forwarding configuration on the switch and be sure that the configured IP address matches the address of SRV1.

Sample DHCP Troubleshooting Flow



© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0-127

Use the **show ip dhcp snooping** command to verify the DHCP snooping configuration. The DHCP snooping should be enabled globally using the **ip dhcp snooping** global configuration command.

Sample DHCP Troubleshooting Flow (Cont.)

The **show ip dhcp snooping** command displays configuration of the DHCP snooping feature.

```
BSW1#show ip dhcp snooping
Switch DHCP snooping is enabled
DHCP snooping is configured on following VLANs:
1001
DHCP snooping is operational on following VLANs:
1001
DHCP snooping is configured on the following L3 Interfaces:

Insertion of option 82 is enabled
  circuit-id format: vlan-mod-port
  remote-id format: MAC
Option 82 on untrusted port is not allowed
Verification of hwaddr field is enabled
Verification of giaddr field is enabled
DHCP snooping trust/rate is configured on the following Interfaces:
```

Interface	Trusted	Rate limit (pps)
FastEthernet0/3	yes	unlimited
FastEthernet0/4	yes	unlimited

© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0-128

From the output shown in the figure, you can see that DHCP snooping is enabled on VLAN 1001. DHCP option 82 is turned on and FastEthernet0/3 and FastEthernet0/4 interfaces are in trusted state.

Also, DHCP snooping should be configured on the proper VLAN (VLAN 1001) using the command **ip dhcp snooping vlan 1001**. The switch should insert the DHCP option-82 field in forwarded DHCP request messages to the DHCP server (default behavior). If the insertion of DHCP option 82 is disabled, use the **ip dhcp snooping information option** global configuration command.

Sample DHCP Troubleshooting Flow (Cont.)

The **show ip dhcp snooping database** command displays the current operating state of the database agent and statistics associated with the transfers.

```
BSW1#sh ip dhcp snooping database
Agent URL : flash:/DHSCP-SDB
Write delay Timer : 300 seconds
Abort Timer : 300 seconds

Agent Running : No
Delay Timer Expiry : Not Running
Abort Timer Expiry : Not Running

Last Succeeded Time : 19:23:25 UTC Fri May 22 2009
Last Failed Time : None
Last Failed Reason : No failure recorded.

Total Attempts      :      18   Startup Failures :      0
Successful Transfers :      18   Failed Transfers :      0
Successful Reads     :      18   Failed Reads    :      0
Successful Writes    :      18   Failed Writes   :      0
Media Failures       :          0
```

© 2009 Cisco Systems, Inc. All rights reserved.


TSHOOT v1.0-129

Verify the DHCP snooping binding database agent using **show ip dhcp snooping database** on BSW1. If disabled, use the **ip dhcp snooping database flash:/filename** global configuration command. The output in the figure shows important transport statistics and the state of the database. You can see 18 successful transfers, 18 successful reads and writes, and no fails.

Ports connected to SRV1 should be in a trusted state. If not, use the **ip dhcp snooping trust** interface configuration command to change the state of the interfaces to “trusted.” All other interfaces connected to DHCP clients should be “untrusted” (by default).


Troubleshooting Cisco IOS Firewall

The figure illustrates an example of a method that you could follow to diagnose and resolve problems related to stateful packet inspection on the Cisco IOS Firewall.



Ping is a very popular tool that uses ICMP. ICMP relies on IP to perform its tasks, and a failing ping between devices is evidence of lack of Layer 3 connectivity.

The next step is to verify the configuration of your access lists and determine whether these access lists are applied to the right interfaces in the right direction.



The **show access-lists** and **show ip interface *intf-id*** commands are very useful commands that can be used to check your access list configuration on the router.

Verify Configuration of Access Lists

The **show access list** command displays all access lists that are configured on the device.

```
IRO2#show access-lists
Standard IP access list INTERNAL-NETWORKS
 10 permit 10.1.128.0, wildcard bits 0.0.127.255 (4 matches)
```


The **show ip interfaces *intf-id*** command displays essential interface information and the access lists that were applied.

```
IRO2#show ip interface fast 0/1
FastEthernet0/7 is up, line protocol is up
 Internet address is 172.24.244.81/29
 Broadcast address is 255.255.255.255
 Address determined by non-volatile memory
 MTU is 1500 bytes
 Helper address is not set
 Directed broadcast forwarding is disabled
 Multicast reserved groups joined: 224.0.0.10
 Outgoing access list is INTERNAL-NETWORKS
 Inbound access list is not set
```

© 2009 Cisco Systems, Inc. All rights reserved. TSHOOT v1.0—132

From the output shown in the figure, you can see the configuration of an access list named **INTERNAL-NETWORK**, and this access list is applied outbound on the outbound direction of FastEthernet0/1 on IRO0. No access list is applied on inbound direction on the same interface.

Sample Cisco IOS Firewall Troubleshooting Flow



© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0-133

After you are certain of your access list configuration, you should check the Context-Based Access Control (CBAC) configuration on the router, the policy, and proper direction of the inspect statement. Be sure that the inspect statement has an opposite direction to that of the access list, configured to protect your network from outside.

Verify Firewall Policy

The **show ip inspect all** command displays firewall policy in detail.

```
IR01# show ip inspect all
Session audit trail is disabled
Session alert is enabled
one-minute (sampling period) thresholds are [400:20000] connections
max-incomplete sessions thresholds are [400:20000]
max-incomplete tcp connections per host is 100000. Block-time 0
minute
tcp synwait-time is 30 sec -- tcp finwait-time is 5 sec
tcp idle-time is 3600 sec -- udp idle-time is 30 sec
dns-timeout is 5 sec
Outgoing inspection rule is FWL
smtp max-data 20000 alert is on audit-trail is off timeout 3600
ftp alert is on audit-trail is off timeout 3600
tcp alert is on audit-trail is off timeout 3600
udp alert is on audit-trail is off timeout 30
Inbound access list is 105
Outgoing access list is not set
Established Sessions
Session 49AA929C (10.1.160.61:14320)=>(172.23.224.1:23) tcp SIS_OPEN
Half-open Sessions
Session 467479EC (10.1.160.65:20150)=>(172.23.224.1:23) tcp
SIS_OPENING matches)
```

© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0-134

In this example, you can see different DoS parameters configured on router BRO1. The inspection rule named FWL is applied in the outgoing direction and access list 105 is applied in the inbound direction. There is one established session between 10.1.160.61 and 172.23.224.1 and one half-opened session between 10.1.160.65 and 172.23.224.1.

Verify Firewall Policy(Cont.)

The **show ip inspect statistics** command displays statistics of the firewall policy.

```
IR01#show ip inspect statistics
Packet inspection statistics [process switch:fast switch]
tcp packets: [43:281]
udp packets: [75:30]
smtp packets: [0:0]
ftp packets: [0:0]
dns packets: [0:0]
Interfaces configured for inspection 1
Current session counts (estab/half-open/terminating) [11:7:0]
Maxever session counts (estab/half-open/terminating) [16:12:1]
```

© 2009 Cisco Systems, Inc. All rights reserved.


TSHOOT v1.0-135

The example in the figure displays important statistics of the firewall policy. You can see that 43 TCP process switched packets and 281 TCP fast switched packets pass through the firewall. The current number of established sessions is 11, and you have seven half-opened sessions. In addition, the output displays the maximum number of established sessions (16 in this example) and half-opened sessions (12 in this example).

Troubleshooting SSH

The figure illustrates an example of a method that you could follow to diagnose and resolve problems related to the SSH protocol.

Sample SSH Troubleshooting Flow




© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0—136

If you do not have SSH remote access to the device first, you should check Layer 4 connectivity to the management address of the device. You can use **ping** for this purpose. If your ping is successful, you do not have a Layer 3 problem in the network.

Sample SSH Troubleshooting Flow



© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0—137

Check whether SSH is enabled on the device.

Verify SSH Server

The **show ssh** command is a great tool that you can use to check the state of the SSH server on the device.

SSH server disabled

```
BRO2#show ssh
%No SSHv2 server connections running.
%No SSHv1 server connections running.
```

SSH server enabled


```
BRO2#show ssh
Connection    Version    Encryption    State          Username
1             2          3DES          Session Started cisco
```

© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0-138

The **show ssh** command displays information for the SSH server status on the device. The first example in the figure shows that the SSH server is disabled on BRO2. In the second example in the figure, SSHv2 is enabled and encryption is 3DES. If SSH is disabled, you can use **ip ssh version 2** global configuration mode command to enable it. If SSH configuration commands are rejected as illegal commands, you have not successfully generated an RSA key pair for your device.

Sample SSH Troubleshooting Flow



© 2009 Cisco Systems, Inc. All rights reserved.

TSHOOT v1.0-139

Before generating an RSA key pair, the hostname and domain should be configured on the device. If it is not configured, use **hostname *name*** and **ip domain-name** global configuration mode commands. When you are sure that the hostname and domain are properly configured on the device, use the **crypto key generate rsa** command to generate an RSA key pair and enable SSH server on the router using **ip ssh version 2** global configuration mode command.

Lab Debrief Notes

Use these notes sections to write down the primary learning points that are discussed during the Lab Debrief.

Lab 6-2: Alternate Solutions

[illegible]

[illegible]

Lab 6-2: Procedure and Communication Improvements

[illegible]

[illegible]

Lab 6-2: References

If you need more information on the commands and their options, you can go to the following sections of <http://www.cisco.com>.

- Cisco Systems, Inc. Command References for Cisco Catalyst LAN Switches: Go to Product Support (<http://www.cisco.com/web/psa/products/index.html>), select Switches, select LAN Switches and then the product family that you are working with. The Command References can then be found under the “Reference Guides” section.
- Cisco Systems, Inc. *Catalyst 3560 Switch Software Configuration Guide, Rel. 12.2(44)SE: Configuring DHCP Features and IP Source Guard*. San Jose, California, January 2008: http://www.cisco.com/en/US/docs/switches/lan/catalyst3560/software/release/12.2_44_se/configuration/guide/swdhcp82.html
- Cisco Systems, Inc. *Understanding and Troubleshooting HSRP Problems in Catalyst Switch Networks*. San Jose, California, May 2009: http://www.cisco.com/en/US/tech/tk648/tk362/technologies_tech_note09186a0080094afd.shtml
- Cisco Systems, Inc. *Spanning Tree Protocol Root Guard Enhancement*. San Jose, California, August 2005: http://www.cisco.com/en/US/tech/tk389/tk621/technologies_tech_note09186a00800ae96b.shtml
- Cisco Systems, Inc. *Context-Based Access Control (CBAC) Introduction and Configuration*. San Jose, California, June 2008: http://www.cisco.com/en/US/products/sw/secursw/ps1018/products_tech_note09186a0080094e8b.shtml
- Cisco Systems, Inc. *Cisco IOS Classic Firewall/IPS: Configuring Context-Based Access Control (CBAC) for Denial-of-Service Protection*. San Jose, California, June 2008: http://www.cisco.com/en/US/products/sw/secursw/ps1018/products_configuration_example09186a00808b7200.shtml
- Cisco Systems, Inc. *Configuring Secure Shell on Routers and Switches Running Cisco IOS*. San Jose, California, June 2007: http://www.cisco.com/en/US/tech/tk583/tk617/technologies_tech_note09186a00800949e2.shtml
- Cisco Systems, Inc. *Secure Shell (SSH) FAQ*. San Jose, California, February 2006: http://www.cisco.com/en/US/tech/tk583/tk617/technologies_q_and_a_item09186a0080267e0f.shtml

Lab 7-1: Troubleshooting Complex Environments

Complete this lab activity to practice what you learned in the course.

Activity Objective


In this activity, you will troubleshoot various problems in a complex network environment. After completing this activity, you will be able to meet these objectives:

- Diagnose and resolve problems related to any feature, protocol, or technology that could be encountered in a complex, integrated enterprise network.
- Document troubleshooting progress, configuration changes, and problem resolution.

Information Packet

The figure illustrates what you will accomplish in this activity:

Visual Objective



Required Resources

These are the resources and equipment that are required to complete this activity:

- One PC with Internet access per team member
- Four Cisco Catalyst 3560 Series Switches per team
- Six Cisco 1841 Integrated Services Routers per team
- Three client PCs per team
- One server per team

Job Aids


These job aids are available to help you complete the lab activity.

- Trouble tickets
- Troubleshooting logs
- Change logs

Introduction: The Enterprise Network

The network that you will be troubleshooting in this lab strongly resembles the network that you have been working on in all the other labs in this course. In general, the physical connections, VLAN structure, and IP address plan are the same. However, there are some differences, so you cannot simply trust your baselines from the previous labs.

It is not necessary to perform a full network survey before starting to troubleshoot. The major differences between the design for this lab and the design used in the previous labs are documented in this introduction section. Ask your instructor for clarification when you have doubts about the design, implementation, or policies used on the lab network.




The physical topology is largely the same as it was in the previous labs. Note the following differences:

- Routers CRO1 and CRO2 do not have redundant connections to switches CSW1 and CSW2. Router CRO1 is connected to switch CSW1 and router CRO2 is connected to switch CSW2.
- Some of the WAN connections were removed, and a new connection was added. Only the connections shown in the diagram are used.
 - The WAN connection between router BRO1 and router CRO1 is a direct serial link, using PPP as the encapsulation.

- The link between router BRO1 and router CRO2 uses Frame Relay. The PVC between these routers is identified by data link connection identifier (DLCI) 112 on router BRO1 and DLCI 121 on router CRO2.
- The link between router BRO2 and router IRO2 is new and uses Frame Relay. The PVC between these two routers is identified by DLCI 114 on router BRO2 and DLCI 122 on router IRO2.


- Switch BSW1 has been configured as a multilayer switch.



The logical topology is also slightly different than in previous labs. Note the following differences:

- Instead of using routed ports from switches CSW1 and CSW2 to routers CRO1 and CRO2, a transit VLAN has been defined (VLAN 130).
- The LANs in the branch have been separated.
 - The guest LAN (VLAN 19) uses Layer 2 switching on switch BSW1. Guest traffic from client PC CLT3 is switched to router BRO2, which provides WAN access to router IRO2 and from there to the Internet. Traffic from the guest network is not allowed to enter the headquarters LAN.
 - The office LAN (VLAN 17) and other VLANs use router BRO1 to access the WAN, the headquarters LAN, and the Internet.
 - Traffic between VLANs (except VLAN 19) in the branch is switched on Layer 3 by switch BSW1.
 - Because routers BRO1 and BRO2 perform different roles and guest traffic is separated from the other user traffic, there is no first-hop redundancy for the branch office LANs.
 - Router BRO2 acts as a DHCP server for the guest VLAN 19 and switch BSW1 acts as a DHCP server for all other branch office VLANs.

Lab 7-1: Routing Design



Although it does have some similarities, the routing design of the network is different from the routing design in previous labs. The following routing protocols and mechanisms are used:

- In the headquarters LAN, OSPF is used as the routing protocol. Three different areas have been defined.
 - The backbone area 0.0.0.0 is formed by the transit VLAN 130 between routers CSW1, CSW2, CRO1 and CRO2. The loopback interfaces of these routers are also part of area 0.0.0.0.
 - The transit VLAN 129 between routers CSW1, CSW2, IRO1 and IRO2 is part of area 10.1.192.0. The loopback interfaces of routers IRO1 and IRO2 are also part of this area.
 - All other access VLANs are part of area 10.1.128.0.
- On the WAN and in the branch, EIGRP is used as the routing protocol. Routers CRO1, CRO2, BR01 and BSW1 are all using EIGRP in AS 100.
- Routers CRO1 and CRO2 redistribute routing information between OSPF and EIGRP.
- The WAN link between routers CRO2 and BR01 should be used as a backup only, in case the primary link between CRO1 and BR01 fails. Routing policies have been implemented to prefer the path across the WAN via CRO1 over the path via CRO2.
- BGP is used to route to the Internet via two redundant ISPs.
 - The AS numbers and IP addresses of the service provider routers ISP1 and ISP2 are the same as in Lab 5-4.
 - NAT is not used. The prefix 10.1.128.0/17 is advertised to both routers ISP1 and ISP2. From both service providers only the default route is accepted. A routing policy is in place to filter all other prefixes that the providers might send.
- Between routers IRO2 and BR02 static routing is used. A default route is configured on router BR02 pointing to IRO2. On IRO2, a static route to the subnet of the guest VLAN and a static route to the loopback IP address of router BR02 have been configured.

Regarding policies for this network, consider the following guidelines:

- Any routing policies that are implemented should stay in effect.
- Where redundant connections are present, the network should reconverge when a link or device failure occurs.
- Any security mechanisms and policies that are implemented should stay in effect.

If your solution restores connectivity, but does not comply with the policies mentioned, your changes will be considered a workaround and points will be subtracted from your score.

Note	If you have any questions about the design, implementation or policies of the lab network, ask your instructor for clarification.

Trouble Ticket A: No Connectivity from CLT1 to SRV1

The user at headquarters who uses client PC CLT1 on VLAN 17 has complained that he cannot use the file services on server SRV1.

Your task is to restore connectivity from client PC CLT1 to server SRV1 and ensure that the user can view the directory “\\SRV1\Public” and upload a file to it.

Trouble Ticket B: No Internet Access from CLT1

Many users on the network are experiencing problems when accessing the Internet. You have a report from an office user, who uses client PC CLT1, claiming that he cannot browse to <http://www.isp3.local> (172.34.224.1).

Your task is to restore the connectivity from client PC CLT1 to the Internet and ensure that the user can connect to the site <http://www.isp3.local> using a web browser.

Trouble Ticket C: No Connectivity Between Headquarters and Branch Office

The office users in the branch office have complained that they cannot connect to server SRV1 at headquarters and they cannot browse the Internet. Additionally, the core WAN routers CRO1 and CRO2 cannot be reached via Telnet or SSH from the branch office. You can use client PC CLT2 in VLAN 17 for testing purposes.

Your task is to restore connectivity between headquarters and the branch office. This task includes ensuring that you can use Telnet to connect to routers CRO1 and CRO2 and ensuring that the user on client PC CLT2 can view the directory “\\SRV1\Public” and upload a file to it.

Trouble Ticket D: No Internet Access for Guest Users

A guest user in the branch office who uses PC CLT3 on VLAN 19 has complained that he cannot browse to the website `www.isp3.local`.

Your task is to restore connectivity from client PC CLT3 to the Internet and ensure that the user can connect to the site `http://www.isp3.local` using a web browser.

Network Maintenance: Verify Network Operation

After resolving the problems reported by the users, you should verify the operational state of the network and look for unreported problems, or issues that might cause problems in the future. At least the following areas should be investigated:

- **Redundancy and failover:** Analyze which redundant paths are available in the network and ensure that the network re-converges after a component or link failure if possible.
- **Security policy compliance:** Verify that none of the implemented security mechanisms was disabled during the troubleshooting process and verify that the implemented security features have been deployed in a consistent manner across devices.
- **Network management:** Verify that all network management features and protocols that are implemented on the devices are operating correctly.

Instructions

Together with your team members, create a troubleshooting plan to divide the work, assign each team member appropriate roles and coordinate device access between team members. Together, work on Trouble Tickets A, B, C and D to resolve the issues. Document your progress in the following Troubleshooting Logs in order to help facilitate efficient communication within the team and to have an overview of your troubleshooting process for reference during the Lab Debrief discussions. Document any changes that you make to the configurations of the devices in the following Change Logs. Your logs (or copies) have to be handed over to the instructor for scoring purposes.

You are allowed a total of three hours to complete as many of the trouble tickets as you can. After three hours, the instructor will debrief the lab and review all trouble tickets and their solutions. The main objective for the troubleshooting labs in this course is to give you an opportunity to practice structured troubleshooting and maintenance.

In this final lab, your performance will be scored on several aspects. These aspects include problem resolution, but also the use of proper processes and procedures.

The Activity Verification section for each trouble ticket includes terminal objectives and details regarding the points that can be scored for different aspects of the solution.

Lab Setup

The instructor will provide you with directions to prepare the lab equipment for this lab. After the instructor indicates that the lab is fully prepared, you are ready to start troubleshooting.

Trouble Ticket A Troubleshooting Log

Use this log to document your actions and results during the troubleshooting process. Provide this log (or a copy) to the instructor.

[illegible]

Trouble Ticket A Change Log

Use this log to document any changes you made during the troubleshooting process. Provide this log (or a copy) to the instructor.

[illegible]

Activity Verification

You have completed this trouble ticket when you attain these results.

Trouble Ticket A

The problem multiplier for this ticket is 6. The maximum total number of points scored for this ticket is 150.

- **Result (10 points):** You have successfully restored connectivity from client PC CLT1 to server SRV1. As proof of your solution, demonstrate to the instructor that client PC CLT1 can transfer a file to the directory “\\SRV1\Public”.
- **Solution (5 points):** You have addressed the root cause or causes of the problem, not implemented a workaround. Give your instructor your completed change log as proof that you have addressed the problem.
- **Process (5 points):** You have clearly documented your process, your solution, and any changes that you have made to the device configurations. Give your instructor your completed troubleshooting log as proof that you have documented everything appropriately.
- **Timing (5 points):** You have restored connectivity in a timely manner. (The maximum of 5 points is scored if connectivity has been restored within half an hour. For each additional half hour required to restore connectivity, a point is subtracted. For example, you would score 4 points for restoring connectivity within 31 and 60 minutes.)

Total score: _____ points x multiplier 6 = _____ points

Trouble Ticket B Troubleshooting Log

Use this log to document your actions and results during the troubleshooting process. Provide this log (or a copy) to the instructor.

[illegible]

Trouble Ticket B Change Log

Use this log to document any changes you made during the troubleshooting process. Provide this log (or a copy) to the instructor.

[illegible]

Activity Verification

You have completed this lab when you attain these results.

Trouble Ticket B

The problem multiplier for this ticket is 10. The maximum total number of points scored for this ticket is 250.

- **Result (6 points):** You have successfully restored connectivity from client PC CLT1 to IP address 172.34.224.1 (www.isp3.local). As proof of your solution, demonstrate to the instructor that client PC CLT1 can ping IP address 172.34.224.1.
- **Result (4 points):** You have successfully restored connectivity from client PC CLT1 to server www.isp3.local. As proof of your solution, demonstrate to the instructor that client PC CLT1 can browse to http://www.isp3.local.
- **Solution (5 points):** You have addressed the root cause or causes of the problem, not implemented a workaround. Give your instructor your completed change log as proof that you have addressed the problem.
- **Process (5 points):** You have clearly documented your process, your solution, and any changes that you have made to the device configurations. Give your instructor your completed troubleshooting log as proof that you have documented everything appropriately.
- **Timing (5 points):** You have restored connectivity in a timely manner. (The maximum of 5 points is scored if connectivity has been restored within half an hour. For each additional half hour required to restore connectivity, a point is subtracted).

Total score: _____ points x multiplier 10 = _____ points

Trouble Ticket C Troubleshooting Log

Use this log to document your actions and results during the troubleshooting process. Provide this log (or a copy) to the instructor.

[illegible]

Trouble Ticket C Change Log

Use this log to document any changes you made during the troubleshooting process. Provide this log (or a copy) to the instructor.

[illegible]

Activity Verification

You have completed this lab when you attain these results.

Trouble Ticket C

The problem multiplier for this ticket is 8. The maximum total number of points scored for this ticket is 200.

- **Result (3 points):** You have successfully restored connectivity from client PC CLT2 to router CRO1. As proof of your solution, demonstrate to the instructor that you can use Telnet from client PC CLT2 to connect to the loopback IP address 10.1.220.1 of router CRO1.
- **Result (3 points):** You have successfully restored connectivity from client PC CLT2 to router CRO2. As proof of your solution, demonstrate to the instructor that you can use Telnet from client PC CLT2 to connect to the loopback IP address 10.1.220.2 of router CRO2.
- **Result (4 points):** You have successfully restored connectivity from client PC CLT2 to server SRV1. As proof of your solution, demonstrate to the instructor that client PC CLT2 can transfer a file to the directory "\\SRV1\Public".
- **Solution (5 points):** You have addressed the root cause or causes of the problem, not implemented a workaround. Give your instructor your completed change log as proof that you have addressed the problem.
- **Process (5 points):** You have clearly documented your process, your solution, and any changes that you have made to the device configurations. Give your instructor your completed troubleshooting log as proof that you have documented everything appropriately.
- **Timing (5 points):** You have restored connectivity in a timely manner. (The maximum of 5 points is scored if connectivity has been restored within half an hour. For each additional half hour required to restore connectivity, a point is subtracted).

Total score: _____ points x multiplier 8 = _____ points

Trouble Ticket D Troubleshooting Log

Use this log to document your actions and results during the troubleshooting process. Provide this log (or a copy) to the instructor.

[illegible]

Trouble Ticket D Change Log

Use this log to document any changes you made during the troubleshooting process. Provide this log (or a copy) to the instructor.

[illegible]

Activity Verification

You have completed this lab when you attain these results.

Trouble Ticket D

The problem multiplier for this ticket is 6. The maximum total number of points scored for this ticket is 150.

- **Result (8 points):** You have successfully restored connectivity from client PC CLT3 to IP address 172.34.224.1 (www.isp3.local). As proof of your solution, demonstrate to the instructor that client PC CLT3 can ping IP address 172.34.224.1.
- **Result (2 points):** You have successfully restored connectivity from client PC CLT3 to server www.isp3.local. As proof of your solution, demonstrate to the instructor that client PC CLT3 can browse to http://www.isp3.local.
- **Solution (5 points):** You have addressed the root cause or causes of the problem, not implemented a workaround. Give your instructor your completed change log as proof that you have addressed the problem.
- **Process (5 points):** You have clearly documented your process, your solution, and any changes that you have made to the device configurations. Give your instructor your completed troubleshooting log as proof that you have documented everything appropriately.
- **Timing (5 points):** You have restored connectivity in a timely manner. (The maximum of 5 points is scored if connectivity has been restored within half an hour. For each additional half hour required to restore connectivity a point is subtracted).

Total score: _____ points x multiplier 6 = _____ points

Network Maintenance Process Log

Use this log to document your actions and results during the maintenance process. Provide this log (or a copy) to the instructor.

[illegible]

Network Maintenance Change Log

Use this log to document any changes you made during the maintenance process. Provide this log (or a copy) to the instructor.

[illegible]

Activity Verification

You have completed this lab when you attain these results.

Network Maintenance

The problem multiplier for this process is 10. The maximum total number of points scored for this ticket is 250.

- **Result (5 points):** You have ensured that Internet connectivity is maintained when any of the two ISP connections fails. As proof of your solution, initiate a continuous ping to 172.34.224.1 on one of the client PCs and demonstrate to the instructor that connectivity is regained after link or device failure. The instructor will select the link or device to be disabled.
- **Result (5 points):** You have ensured that connectivity from the branch office network to the headquarters network is maintained when any of the two WAN connections fails. As proof of your solution, initiate a continuous ping to server SRV1 on client PC CLT2 and demonstrate to the instructor that connectivity is regained after link or device failure. The instructor will select the link or device to be disabled.
- **Result (5 points):** You have discovered and addressed additional issues that were not part of the reported problems. Summarize the issues that you found and the way you addressed the issues and enter the information in your log for the instructor.
- **Solution (5 points):** The changes that you have made comply with the policies implemented on the network. Give your instructor your completed change log as proof that you have addressed the problem.
- **Process (5 points):** You have clearly documented your verification process, the issues found, and any changes that you have made to the device configurations. Give your instructor your completed your network maintenance process log as proof that you have documented everything appropriately.

Total score: _____ points x multiplier 10 = _____ points

Lab Score:

Your total score of the lab can be calculated here at the end of the debrief lesson:

- Trouble Ticket A: _____ points
- Trouble Ticket B: _____ points
- Trouble Ticket C: _____ points
- Trouble Ticket D: _____ points
- Network Maintenance: _____ points
- Grand total: _____ points

Lab 7-1: Sample Troubleshooting Flows

This integrated capstone lab covers all technologies that were practiced in the previous labs. Therefore, no specific additional troubleshooting flows are provided for this lab. Refer to the Sample Troubleshooting Flows sections in previous labs for examples of troubleshooting procedures for specific technologies.

Lab Debrief Notes

Use these notes sections to write down the primary learning points that are discussed during the Lab Debrief.

Lab 7-1: Alternate Solutions

[illegible]

Lab 7-1: Alternate Methods and Processes

[illegible]

[illegible]

Lab 7-1: Important Commands and Tools

[illegible]

Lab 7-1: References

If you need more information on the commands and their options, you can go to the following sections of <http://www.cisco.com>.

Cisco Systems, Inc. *Cisco IOS Debug Command Reference*. San Jose, California, February 2009: http://www.cisco.com/en/US/docs/ios/debug/command/reference/db_book.html

Answer Key

The correct answers and expected solutions for the activities that are described in this guide appear here.

Lab 2-1 Answer Key: Introduction to Troubleshooting

When you complete this activity, your documented solutions for the trouble ticket will be similar to the results here, with differences that are specific to your device or workgroup:

Note	If your solution to the problem described in the trouble ticket consists of commands other than the ones listed here, bring your alternative solution to the attention of the instructor and the group during the Lab Debrief discussion.

The problem is caused by misconfiguration of the routing protocol on routers CRO1 and CRO2. EIGRP is configured on all routers in the network, but the routers CRO1 and CRO2 at the headquarters and the routers BRO1 and BRO2 at Branch 1 did not become neighbors across the WAN. The underlying cause is that the network statements for EIGRP on routers CRO1 and CRO2 have not been correctly configured.

You can remedy this issue by reconfiguring EIGRP on routers CRO1 and CRO2. The following commands will restore the configuration on router CRO1:

```
router eigrp 1
  no network 10.1.193.0 0.0.0.0
  no network 10.1.194.0 0.0.0.0
  network 10.1.193.1 0.0.0.0
  network 10.1.194.1 0.0.0.0
  network 10.1.194.5 0.0.0.0
```

In addition, the following commands will restore the configuration on router CRO2:

```
router eigrp 1
  no network 10.1.193.0 0.0.0.0
  no network 10.1.194.0 0.0.0.0
  network 10.1.193.5 0.0.0.0
  network 10.1.194.9 0.0.0.0
  network 10.1.194.13 0.0.0.0
```

Although restoring one of the two routers is sufficient action to take to restore connectivity, you will not have redundancy in the WAN, despite the fact that the design is built for redundant WAN connections.

Student Notes

Use this Student Notes section to write down any alternate troubleshooting methods and additional troubleshooting commands that you learned during the labs and lab reviews.

[illegible]

Student Notes

Use this Student Notes section to write down any alternate troubleshooting methods and additional troubleshooting commands that you learned during the labs and lab reviews.

[illegible]

Lab 3-1 Answer Key: Maintenance and Troubleshooting Tools

This lab only contains review and planning tasks, not any configuration or troubleshooting tasks; therefore, configurations are not listed in this section. After reviewing the configured maintenance and troubleshooting tools, the filled-in table should reflect the following:

Device	Configured feature	Target server	Target tool or application
ASW1	Syslog	SRV1	Syslog server
	DNS	SRV1	DNS server
	Configuration archive	SRV1	TFTP server
	SNMP traps	SRV1	
	NTP	IRO1, IRO2	NTP server
CSW1	Syslog	SRV1	Syslog server
	DNS	SRV1	DNS server
	Configuration archive	SRV1	TFTP server
	SNMP traps	SRV1	
	NTP	IRO1, IRO2	NTP server
CSW2	Syslog	SRV1	Syslog server
	DNS	SRV1	DNS server
	Configuration archive	SRV1	TFTP server
	SNMP traps	SRV1	
	NTP	IRO1, IRO2	NTP server
IRO1	Syslog	SRV1	Syslog server
	DNS	SRV1	DNS server
	Configuration archive	SRV1	TFTP server
	SNMP traps	SRV1	
	NetFlow	SRV1	
IRO2	NTP	ISP1, ISP2	NTP server
	Syslog	SRV1	Syslog server
	DNS	SRV1	DNS server
	Configuration archive	SRV1	TFTP server
	SNMP traps	SRV1	
IRO2	NetFlow	SRV1	
	NTP	ISP1, ISP2	NTP server

Device	Configured feature	Target server	Target tool or application
CRO1	Syslog	SRV1	Syslog server
	DNS	SRV1	DNS server
	Configuration archive	SRV1	TFTP server
	SNMP traps	SRV1	
	NetFlow	SRV1	
	NTP	IRO1, IRO2	NTP server
CRO2	Syslog	SRV1	Syslog server
	DNS	SRV1	DNS server
	Configuration archive	SRV1	TFTP server
	SNMP traps	SRV1	
	NetFlow	SRV1	
	NTP	IRO1, IRO2	NTP server
BRO1	Syslog	SRV1	Syslog server
	DNS	SRV1	DNS server
	Configuration archive	SRV1	TFTP server
	SNMP traps	SRV1	
	NTP	IRO1, IRO2	NTP server
BRO2	Syslog	SRV1	Syslog server
	DNS	SRV1	DNS server
	Configuration archive	SRV1	TFTP server
	SNMP traps	SRV1	
	NTP	IRO1, IRO2	NTP server
BSW1	Syslog	SRV1	Syslog server
	DNS	SRV1	DNS server
	Configuration archive	SRV1	TFTP server
	SNMP traps	SRV1	
	NTP	IRO1, IRO2	NTP server

Student Notes

Use this Student Notes section to write down any alternate troubleshooting methods and additional troubleshooting commands that you learned during the labs and lab reviews.

[illegible]

Student Notes

Use this Student Notes section to write down any alternate troubleshooting methods and additional troubleshooting commands that you learned during the labs and lab reviews.

This image shows a single sheet of white paper with horizontal ruling lines. The lines are evenly spaced and run across the width of the page. There are no margins, text, or other markings on the paper.

Lab 4-1 Answer Key: Layer 2 Connectivity and Spanning Tree

When you complete this activity, your documented solutions for the trouble tickets will be similar to the results here, with differences that are specific to your device or workgroup:

Note	If you have solved the problems described in the trouble tickets by using commands other than the ones listed here, bring your alternative solution to the attention of the instructor and the group during the Lab Debrief discussion.

Trouble Ticket A

This trouble ticket consists of several problems that need to be solved before connectivity is restored.

On switch ASW1, the spanning-tree mode had been changed to MST, causing the uplink ports to be placed in “broken (BKN)” spanning-tree state. This resulting state effectively blocks all traffic on the uplinks to switches CSW1 and CSW2.

You can remedy this issue by configuring the following command on switch ASW1:

```
spanning-tree mode rapid-pvst
```

In addition to this problem on switch ASW1, there are two separate problems with the uplinks between switches ASW1, CSW1 and CSW2, respectively. In order to regain connectivity for the clients, it is enough to find and resolve one of the two issues, but to regain the redundancy that is inherent in the physical network design, you need to diagnose and resolve both issues. Proper verification should uncover both issues.

- On switch CSW1, the trunk encapsulation on the EtherChannel towards the access switches has been changed to Inter-Switch Link (ISL) encapsulation. This change causes all Layer 2 traffic on these links (including Cisco Discovery Protocol packets) to fail. However, the links stay up and there are no errors recorded on the interfaces since ISL and 802.1Q are both using a valid Ethernet frame format. To remedy this situation, you should configure the following commands:

```
interface Port-channel 1
    switchport trunk encapsulation dot1q
```

- On switch CSW2, the list of allowed VLANs has been removed from the physical interfaces that are members of the EtherChannel between switch CSW2 and switch ASW1. The removal of the list of allowed VLANs causes an inconsistency between the configuration on the Port-channel interface and the interfaces FastEthernet 0/3 and 0/4 that are members of the EtherChannel. This inconsistency, in turn, causes the interfaces FastEthernet 0/3 and 0/4 to be suspended and the Port-channel interface to go down. To resolve this situation and restore the consistency between the configuration of the Port-channel interface and the FastEthernet interfaces, configure the following commands:

```
interface range FastEthernet 0/5 - 6
    switchport trunk allowed vlan 17-19,128
```

After issuing these commands and re-enabling the links between the access switch ASW1 and the core switches CSW1 and CSW2, client CLT1 will regain connectivity to the rest of the network.

One final issue remains: You cannot use Telnet to connect to switch ASW1 from server SRV1 (or from any other point in the network) because VLAN 128, the management VLAN, was removed from switch ASW1. As a result, the VLAN interface on switch ASW1 for VLAN 128 will be down. By issuing the following commands, the VLAN interface will become operational again and connectivity to the management address of switch ASW1 will be restored:

```
vlan 128
name MGMT
```

Trouble Ticket B

The problem in this trouble ticket is caused by the removal of VLAN 19 (B1S1-GUEST) from the list of allowed VLANs on the trunks between the switch BSW1 and the routers BRO1 and BRO2.

You can remedy this issue by configuring the following commands on BSW1:

```
interface range FastEthernet 0/3-4
switchport trunk allowed vlan add 19
```

After adding VLAN 19 to the list of allowed VLANs, the PC CLT3 should be able to get an IP address from router BRO1 via DHCP and use a browser to connect to the website <http://www.isp3.local>.

Trouble Ticket C

The problem in this trouble ticket is not caused by changes on the equipment in your pod, but by a configuration change in the service provider network, which causes BPDUs to be sent to switch CSW1. The BPDU guard feature is enabled on switch CSW1 to protect against exactly this type of behavior and, therefore, CSW error-disables the port leading to ISP1. You could solve the problem if you remove the BPDU guard feature from the port leading to ISP1, but given that this feature was enabled in the baseline configurations, this solution is not the correct solution. Therefore, your only available option is to escalate the problem to the ISP and request that they research the situation and stop the BPDUs from being sent to switch CSW1.

Student Notes

Use this Student Notes section to write down any alternate troubleshooting methods and additional troubleshooting commands that you learned during the labs and lab reviews.

[illegible]

Student Notes

Use this Student Notes section to write down any alternate troubleshooting methods and additional troubleshooting commands that you learned during the labs and lab reviews.

This image shows a single sheet of white paper with horizontal ruling lines. The lines are evenly spaced and run across the width of the page. There are no margins, text, or other markings on the paper.

Lab 4-2 Answer Key: Layer 3 Switching and First-Hop Redundancy

When you complete this activity, your documented solutions for the trouble tickets will be similar to the results here, with differences that are specific to your device or workgroup:

Note	If you have solved the problems described in the trouble tickets by using commands other than the ones listed here, bring your alternative solution to the attention of the instructor and the group during the Lab Debrief discussion.

Trouble Ticket D

The problem in this trouble ticket is caused by the fact that IP routing has been enabled on the switch. By enabling routing, the switch will become a Layer 3 switch instead of a Layer 2 switch and, as a result, it will depend on its routing table to route packets, including packets that it generates itself. This situation means that the configured default gateway will no longer be used to route packets. The simplest way for you to solve the problem is to disable the Layer 3 switching capabilities and restore the switch as a pure Layer 2 switch by issuing the following command:

```
no ip routing
```

This solution is the preferred solution because it conforms to the original design and baseline configuration.

However, as an alternative, it is also possible to keep the switch in Layer 3 mode, even though the switch does not actually route between VLANs. Still, because it uses the routing table to forward its own packets, it will need a default route to replace the default gateway configuration that it used in Layer 2 mode. You can achieve this result by issuing the following command:

```
no ip default-gateway 10.1.159.254
ip route 0.0.0.0 0.0.0.0 10.1.159.254
```

Trouble Ticket E

This ticket consists of two separate issues:

- Connectivity between clients in VLAN 17 B1S1-OFFICE and server SRV1 when router CRO1 is rebooted
- Connectivity between server SRV1 and switch B1S1

The problem with clients in VLAN 17 is caused by someone using the real IP address of router CRO1 as the default gateway rather than the GLBP virtual IP address. The default gateway for clients is assigned via DHCP by router CRO1. To resolve this problem, you should change the default gateway address that is handed out by router CRO1 to clients as follows:

```
ip dhcp pool B1S1-OFFICE
default-router 10.1.160.126
```

After you make this change, the IP address on client CLT2 should be released and renewed to force the client to update its default gateway.

The second problem is caused by mismatched GLBP parameters between routers CRO1 and CRO2 for VLAN 128. You should change the GLBP group number and virtual IP address on router CRO2 to match router CRO1, as follows:

```
interface FastEthernet0/1.128
  no glbp 28 ip 10.1.163.245
  no glbp 28 preempt
  glbp 128 ip 10.1.163.254
  glbp 128 preempt
```

Trouble Ticket F

This problem consists of two separate issues: To begin with, there is no Layer 2 connectivity between switches CSW1 and CSW2 in the newly created test VLAN 44. Secondly, the configured key-strings in the key chains of switches CSW1 and CSW2 are mismatched.

The lack of Layer 2 connectivity between switches CSW1 and CSW2 in VLAN 44 is caused by the fact that VLAN 44 is not allowed on the trunk between switches CSW1 and CSW2, nor is it allowed to pass through via the trunks between the core switches and access switch ASW1. You can resolve this issue by configuring the following command on switch CSW2:

```
interface Port-channel 10
  switchport trunk allowed vlan add 44
```

You do not need to allow VLAN 44 on the trunks to switch ASW1 unless you would like to use a client on switch ASW1 to test HSRP in VLAN 44.

The authentication failure between switches CSW1 and CSW2 is caused by a mismatch between the key-strings that are configured on switch CSW1 (on switch CSW1, the key-string uses a lowercase letter “l” for the “i” in the string “Clsc0”, while switch CSW2 uses a number “1” for the “i” in the string “Clsc0”). Therefore, you can solve this problem by changing the key on switch CSW1 to match switch CSW2 or vice versa. For instance, configure the following command on switch CSW1:

```
key chain TEST
  key 1
    key-string Clsc0
```

You could even consider choosing a new key-string altogether. Whichever key you decide to use, it is important that you carefully document the new string that was chosen.

Trouble Ticket G

This problem is caused by a Layer 2 connectivity error between routers BRO1 and BRO2. The underlying issue is that VLAN 1000—the VLAN chosen to test HSRP—is configured as the native VLAN on switch BSW1. This configuration causes switch BSW1 to pass the VLAN 1000 frames that it receives from router BRO1 as untagged frames to router BRO2. Router BRO2 expects the frames in VLAN 1000 to be tagged and, therefore, discards the frames. There are several solutions to this problem. Essentially, both sides of the trunks need to either tag—or not tag—the traffic for VLAN 1000 in a consistent manner (or a different VLAN should be selected for this test altogether).

One of the solutions is for you to configure routers BRO1 and BRO2 to associate untagged frames with VLAN 1000, and the corresponding sub-interface, by configuring VLAN 1000 as the native VLAN using the following command:

```
interface FastEthernet0/1.1000
 encapsulation dot1Q 1000 native
```

Alternatively, you can configure BSW1 to tag all traffic on 802.1Q trunks, including native VLAN traffic:

```
vlan dot1q tag native
```

Student Notes

Use this Student Notes section to write down any alternate troubleshooting methods and additional troubleshooting commands that you learned during the labs and lab reviews.

[illegible]

Student Notes

Use this Student Notes section to write down any alternate troubleshooting methods and additional troubleshooting commands that you learned during the labs and lab reviews.

[illegible]

Lab 5-1 Answer Key: Layer 3 Connectivity and EIGRP

When you complete this activity, your documented solutions for the trouble tickets will be similar to the results here, with differences that are specific to your device or workgroup:

Note	If you have solved the problems described in the trouble tickets by using commands other than the ones listed here, bring your alternative solution to the attention of the instructor and the group during the lab debrief discussion.

Trouble Ticket H

Trouble Tickets H, I, and J are all affected by the problems that are introduced in Trouble Ticket I. Therefore, the solution Trouble Ticket H lists only the commands that you need to complete the resolution of Trouble Ticket H, after you have resolved the problems introduced in Trouble Ticket I.

The only problem that remains on router BRO1 after the connectivity across the WAN has been restored is that the IP address for subinterface FastEthernet 0/1.29, which belongs to the CCTV VLAN 29, is misconfigured. It is configured as 10.1.164.126/26, while the CCTV subnet is 10.1.163.64/26. You should change the IP address on BRO1 as follows:

```
interface FastEthernet 0/1.29
ip address 10.1.163.126 255.255.255.192
```

There is no need for you to change the network statement for EIGRP because the EIGRP network statements on router BRO1 include an entry **network 10.1.160.0 0.0.3.255**, which covers IP address 10.1.163.126.

If you have simply assumed that network 10.1.164.64/26 should have been used for the CCTV VLAN, contrary to what the trouble ticket states, and you have adapted the configurations accordingly, you have chosen a solution that is not considered correct.

Trouble Ticket I

The main problem in this trouble ticket is the fact that the configuration of router CRO1 has been loaded on router CRO2 and vice versa. There are several solutions to this problem. Either you can restore connectivity to the headquarters LAN and then load the correct configurations from the archive on the TFTP server, or you fully reconfigure both routers based on the available documentation.

To restore connectivity to the headquarters LAN on router CRO1, you should correct the IP address on interface FastEthernet 0/0 (or interface FastEthernet 0/1) and enable it under the EIGRP routing protocol. You can achieve this change by issuing the following commands:

```
interface FastEthernet 0/0
ip address 10.1.192.2 255.255.255.252
!
router eigrp 1
network 10.1.192.2 0.0.0.0
```

After this step, you can copy an archived backup configuration from the TFTP server SRV1 to the startup configuration and reload router CRO1.

Similarly, issue the following commands to restore LAN connectivity for router CRO2:

```
interface FastEthernet 0/0
  ip address 10.1.192.6 255.255.255.252
!
router eigrp 1
  network 10.1.192.6 0.0.0.0
```

After you make this configuration change, you can also copy an archived backup configuration from the TFTP server SRV1 to the startup configuration of router CRO2 and reload the router.

As a less disruptive alternative, you can use the **configure replace** to replace the current running configuration with the archived backup configuration, without the need to reload. However, if timed and coordinated properly, the routers can be reloaded with minimal disruption to the network operation.

If you do not have a good backup configuration to roll back to, you can also fully reconfigure routers CRO1 and CRO2 based on the documentation and diagrams. The following list of commands represents a minimal list of the changes you need to make to reconfigure router CRO1 to match the original baseline configuration:

```
hostname CRO1
!
interface Loopback0
  ip address 10.1.220.1 255.255.255.255
!
interface FastEthernet0/0
  ip address 10.1.192.2 255.255.255.252
!
interface FastEthernet0/1
  ip address 10.1.192.10 255.255.255.252
!
interface Serial0/0/0.121 point-to-point
  ip address 10.1.194.1 255.255.255.252
!
interface Serial0/0/0.122 point-to-point
  ip address 10.1.194.5 255.255.255.252
!
interface Serial0/0/1
  description 128 kbps leased line to BRO1
  ip address 10.1.193.1 255.255.255.252
!
router eigrp 1
  no network 10.1.192.6 0.0.0.0
  no network 10.1.192.14 0.0.0.0
```

```

no network 10.1.193.5 0.0.0.0
no network 10.1.194.9 0.0.0.0
no network 10.1.194.13 0.0.0.0
no network 10.1.220.2 0.0.0.0
network 10.1.192.2 0.0.0.0
network 10.1.192.10 0.0.0.0
network 10.1.193.1 0.0.0.0
network 10.1.194.1 0.0.0.0
network 10.1.194.5 0.0.0.0
network 10.1.220.1 0.0.0.0

```

In a similar manner, you can reconfigure CRO2 to match the original baseline configuration by issuing the following commands:

```

hostname CRO2
!
interface Loopback0
 ip address 10.1.220.2 255.255.255.255
!
interface FastEthernet0/0
 ip address 10.1.192.6 255.255.255.252
!
interface FastEthernet0/1
 ip address 10.1.192.14 255.255.255.252
!
interface Serial0/0/0.121 point-to-point
 ip address 10.1.194.9 255.255.255.252
!
interface Serial0/0/0.122 point-to-point
 ip address 10.1.194.13 255.255.255.252
!
interface Serial0/0/1
 description 128 kbps leased line to BRO2
 ip address 10.1.193.5 255.255.255.252
!
router eigrp 1
 no network 10.1.192.2 0.0.0.0
 no network 10.1.192.10 0.0.0.0
 no network 10.1.193.1 0.0.0.0
 no network 10.1.194.1 0.0.0.0
 no network 10.1.194.5 0.0.0.0
 no network 10.1.220.1 0.0.0.0

```

```
network 10.1.192.6 0.0.0.0
network 10.1.192.14 0.0.0.0
network 10.1.193.5 0.0.0.0
network 10.1.194.9 0.0.0.0
network 10.1.194.13 0.0.0.0
network 10.1.220.2 0.0.0.0
```

After you have restored the configurations of routers CRO1 and CRO2, the configuration of router BRO1 should also be changed. The IP addresses on the Frame Relay subinterfaces on this router were reconfigured to match the changed configurations of routers CRO1 and CRO2. In addition to that, the interface Serial 0/0/1 was shut down on router BRO1. To reverse these changes, the following commands can be configured:

```
interface Serial0/0/1
no shutdown
!
interface Serial0/0/0.111
description PVC to CRO1
ip address 10.1.194.2 255.255.255.252
!
interface Serial0/0/0.112
description PVC to CRO2
ip address 10.1.194.10 255.255.255.252
```

Trouble Ticket J

In order to resolve this ticket completely, the WAN connectivity problems introduced in Trouble Ticket I need to be diagnosed and resolved first.

Trouble Tickets H, I, and J are all affected by the problems that are introduced in Trouble Ticket I. Therefore, the solution Trouble Ticket J lists only the commands that you need to complete the resolution of Trouble Ticket J, after you have resolved the problems introduced in Trouble Ticket I.

The problem in this ticket is caused by a wrong next-hop IP address for the default route configured on router IRO2. Because both routers IRO1 and IRO2 advertise a default route with equal metrics, traffic is load balanced to the two service providers Internet Service Provider 1 and Internet Service Provider 2. Therefore, not all sessions to destinations on the Internet are affected by this problem. This problem is not a control plane problem, but a data plane problem. The default route is distributed to all routers because the next-hop is a valid IP address, but when traffic arrives at IRO2 it cannot be forwarded correctly.

To correct the problem, you should change the default route on IRO2 to point to the correct IP address of router ISP2. You can achieve this change by issuing the following commands:

```
no ip route 0.0.0.0 0.0.0.0 172.24.244.85 track 1
ip route 0.0.0.0 0.0.0.0 172.24.244.86 track 1
```

Student Notes

Use this Student Notes section to write down any alternate troubleshooting methods and additional troubleshooting commands that you learned during the labs and lab reviews.

[illegible]

Student Notes

Use this Student Notes section to write down any alternate troubleshooting methods and additional troubleshooting commands that you learned during the labs and lab reviews.

[illegible]

Lab 5-2 Answer Key: OSPF and Route Redistribution

When you complete this activity, your documented solutions for the trouble tickets will be similar to the results here, with differences that are specific to your device or workgroup:

Note	If you have solved the problems described in the trouble tickets by using commands other than the ones listed here, bring your alternative solution to the attention of the instructor and the group during the lab debrief discussion.

Trouble Ticket K

This trouble ticket revolves around two different problems: an issue that prevents the redistribution from OSPF routes into EIGRP from working correctly and a second issue that causes the redistribution in the other direction, from EIGRP into OSPF, to fail. Both issues need to be solved in order to resolve this ticket.

The redistribution from OSPF routes into EIGRP fails because no seed metric is specified for the redistribution. The default metric for EIGRP redistribution is set to “infinity” and this setting causes the routes not to redistribute into EIGRP. To resolve this problem, a correct seed metric needs to be specified for the redistribution from OSPF to EIGRP. The exact value of the seed metric is not extremely important. If you have chosen a different metric than the metric presented here, your solution is still likely to be correct, as long as it has restored the connectivity for client CLT2. The values chosen here are the metric values that are typical for a directly connected Fast Ethernet interface. If you have doubts about the chosen values for the EIGRP seed metric, bring your solution to the attention of the instructor and the group during the lab debrief discussion.

You can achieve correct redistribution of the OSPF routes into EIGRP by issuing the following commands on routers CRO1 and CRO2:

```
router eigrp 1
 redistribute ospf 100 metric 1544 2000 255 1 1500
```

The second problem in this trouble ticket is caused by the omission of the keyword **subnets** from the **redistribute** command that is responsible for the redistribution from EIGRP into OSPF. As a result, OSPF will only redistribute classful routes. Because all routes in the branch office are subnets of network 10.0.0.0, none of these routes will be redistributed.

To correct this problem issue the following commands on routers CRO1 and CRO2:

```
router ospf 100
 redistribute eigrp 1 metric 100 subnets
```

Although it is not strictly necessary, it is good habit for you to specify a seed metric when you configure redistribution.

At this point, connectivity from client PC CLT2 to server SRV1 is restored. The connectivity to the server <http://www.isp3.local> is dependent on the successful resolution of Trouble Ticket C.

Trouble Ticket L

The problem in this trouble ticket is caused by the fact that router BRO2 does not succeed in establishing an OSPF neighbor relationship with routers CRO1 and CRO2. Two factors prevent BRO2 from successfully becoming neighbors with these routers. One factor is a mismatch between the area that is configured on routers CRO1 and CRO2 (which is area 11) and the area that is configured on router BRO2 (which is area 111). The second factor is that area 11 is configured as a totally stubby area on routers CRO1 and CRO2, while on router BRO2, the area is not configured as a stubby area. The introduction to the trouble ticket states that area 11 should be used for branch office 1 and the area should be configured as a totally stubby area.

You can resolve these problems by issuing the following commands on router BRO2:

```
router ospf 100
  area 11 stub
  no network 10.1.160.0 0.0.3.255 area 111
  no network 10.1.193.6 0.0.0.0 area 111
  no network 10.1.194.6 0.0.0.0 area 111
  no network 10.1.194.14 0.0.0.0 area 111
  no network 10.1.221.2 0.0.0.0 area 111
  network 10.1.160.0 0.0.3.255 area 11
  network 10.1.193.6 0.0.0.0 area 11
  network 10.1.194.6 0.0.0.0 area 11
  network 10.1.194.14 0.0.0.0 area 11
  network 10.1.221.2 0.0.0.0 area 11
```

At this point, connectivity from client PC CLT3 to server SRV1 is restored. The connectivity to the server <http://www.isp3.local> is dependent on the successful resolution of Trouble Ticket M.

Trouble Ticket M

The problem in this ticket is caused by mismatched hello and dead timers on the transit VLAN 129 between routers IRO1 and IRO2, which have been tuned to use 5-second hellos and a 15-second dead timer, and routers CSW1 and CSW2, which use the default hello time of 10 seconds and dead time of 40 seconds. The trouble ticket introduction does not clearly state which timer values should be used. The solution presented here is to change the timers back to the default values on routers IRO1 and IRO2. Changing the values on routers CSW1 and CSW2 to match routers IRO1 and IRO2 is a valid solution as well. However, if you decide to use 5-second hellos and a 15-second dead timer on VLAN 129, you should consider changing the timers on all other interfaces in the network as well for consistency.

To reset the hello and dead timer to the default values on router IRO1, issue the following commands:

```
interface FastEthernet 0/1
  default ip ospf hello-interval
  default ip ospf dead-interval
```

In a similar manner, you can change the timers on router IRO2 by issuing the following commands:

```
interface FastEthernet 0/0
  default ip ospf hello-interval
  default ip ospf dead-interval
```

Trouble Ticket N

Two separate issues contribute to the problem in this ticket. The first issue is that all interfaces have been configured to be passive by default. All the interfaces that are intended as transit interfaces are specifically excluded from the default passive interface configuration. On switch CSW2 interface, VLAN 111 has not been configured as an exception.

Issuing the following commands on switch CSW2 can solve this problem:

```
router ospf 100
  no passive-interface Vlan 111
```

The second problem is caused by a mismatched key-string for the OSPF MD5 authentication between switches CSW1 and CSW2 on VLAN 111. On switch CSW1, the key has been defined as “cisco” (ending with a space) and on switch CSW2, it is configured as “cisco” (without the space at the end).

To resolve this problem you can change the key-string on switch CSW1 by issuing the following commands:

```
interface Vlan 111
  no ip ospf message-digest-key 1
  ip ospf message-digest-key 1 md5 cisco
```

Given that the authentication in this scenario was only configured as a proof-of-concept test, defining a different password is also considered a correct solution.

If you could not resolve the issue, removing the authentication entirely is also considered a valid option, as specified in the trouble ticket text.

Student Notes

Use this Student Notes section to write down any alternate troubleshooting methods and additional troubleshooting commands that you learned during the labs and lab reviews.

[illegible]

Student Notes

Use this Student Notes section to write down any alternate troubleshooting methods and additional troubleshooting commands that you learned during the labs and lab reviews.

[illegible]

Lab 5-3 Answer Key: Border Gateway Protocol

When you complete this activity, your documented solutions for the trouble tickets will be similar to the results here, with differences that are specific to your device or workgroup:

Note	If you have solved the problems described in the trouble tickets by using commands other than the ones listed here, bring your alternative solution to the attention of the instructor and the group during the lab debrief discussion.

Trouble Ticket O

The reason that the peering between routers IRO1 and ISP1 is not established correctly is because an incorrect AS number has been configured for neighbor router ISP1 under the **router bgp** configuration on router IRO1. During session establishment, router IRO1 receives AS number 65525 in the “OPEN” message from router ISP1, which does not match AS number 65255, the AS number that has been configured in the neighbor statement for router ISP1. As a result of this mismatch, router IRO1 immediately closes the session. Verification of the documentation reveals that the correct AS number for router ISP1 is 65525.

To correct this problem, issue the following commands on router IRO1:

```
router bgp 64568
neighbor 192.168.224.254 remote-as 65525
```

A second problem in this ticket is the fact that router IRO1 does not locally inject prefix 172.17.76.0/22 into its BGP table. A network statement for this prefix has been configured under the **router bgp** process on router IRO1, but the required matching route in the routing table is not present. As a result, the prefix is not injected in the BGP table on router IRO1. On router IRO2, a similar network statement under the **router bgp** process exists, but the configuration of router IRO2 also contains a static route to the **null 0** interface for prefix 172.17.76.0/22 in order to provide the required matching route in the routing table.

To resolve this issue, a static route can be configured on router IRO1, similar to the configuration on router IRO2. You can achieve this configuration by issuing the following command:

```
ip route 172.17.76.0 255.255.252.0 Null0
```

Careful testing is required to find and resolve this issue. The problem is difficult to find because it will not affect the connectivity from client PC CLT1 as long as router IRO2 is up because router IRO2 will inject the prefix 172.17.76.0/22 into its BGP table and advertise it to router IRO1. Proper failover testing will reveal this problem.

Trouble Ticket P

The problem described in this ticket is caused by a mistake in the prefix list that filters the networks that are advertised to router ISP2. Instead of permitting the prefix 172.17.76.0/22 and subnets thereof, it permits prefix 172.16.76.0/22 and subnets. As a result, the “implicit deny” at the end of the prefix list denies prefix 172.17.76.0/22 and causes this prefix not to be advertised to router ISP2. The lack of a route to this prefix in the Internet routers prevents return traffic from server <http://www.isp3.local> to reach client PC CLT1 on subnet 172.17.76.0/24.

To resolve this problem, you need to change the prefix list on router IRO2 as follows:

```
no ip prefix-list LOCAL-ROUTES seq 5 permit 172.16.76.0/22 le 24
ip prefix-list LOCAL-ROUTES seq 5 permit 172.17.76.0/22 le 24
```

Removing the prefix list from the neighbor statement to router ISP2 will also restore the connectivity from client PC CLT1 to <http://www.isp3.local> via Internet Service Provider 2. However, this solution introduces the possibility that the company AS becomes a transit AS for traffic from Internet Service Provider 2 to Internet Service Provider 1. Consequently, this solution is not considered a valid solution, unless other measures are implemented to prevent the advertisement of nonlocal routes to router ISP2.

Student Notes

Use this Student Notes section to write down any alternate troubleshooting methods and additional troubleshooting commands that you learned during the labs and lab reviews.

[illegible]

Student Notes

Use this Student Notes section to write down any alternate troubleshooting methods and additional troubleshooting commands that you learned during the labs and lab reviews.

[illegible]

Lab 5-4 Answer Key: Router Performance

When you complete this activity, your documented solutions for the trouble ticket will be similar to the results here, with differences that are specific to your device or workgroup:

Note	If you have solved the problems described in the trouble tickets by using commands other than the ones listed here, bring your alternative solution to the attention of the instructor and the group during the Lab Debrief discussion.

Trouble Ticket Q

This trouble ticket consists of two issues: high CPU utilization on routers CRO1 and CRO2, caused by the use of process switching as the switching mode and a large unnecessary access list, and memory exhaustion on routers BRO1 and BRO2, caused by an inappropriate BGP configuration.

On routers CRO1 and CRO2, you should issue the following commands to disable process switching and re-enable Cisco Express Forwarding:

```
ip cef
interface FastEthernet 0/0
  ip route-cache cef
interface Serial 0/0/1
  ip route-cache cef
```

The first command re-enables Cisco Express Forwarding globally. The second command is used to restore Cisco Express Forwarding on the interfaces. Apply this command to each misconfigured interface.

The next step is for you to remove the huge and unnecessary access list from routers CRO1 and CRO2 by issuing the following commands:

```
no ip access-list standard huge-acl
interface FastEthernet 0/0
  no ip access-group huge-acl in
  no ip access-group huge-acl out
interface Serial 0/0/1
  no ip access-group huge-acl in
  no ip access-group huge-acl out
```

The next step is to solve the memory exhaustion problems on routers BRO1 and BRO2 caused by the huge number of BGP prefixes sent to these routers. First, restore the default value of the BGP scanner on routers BRO1, BRO2, CRO1, and CRO2:

```
router bgp 65000
  bgp scan-time 60
```

Next, you have to resolve the issue of the excessive number of BGP prefixes that are sent to routers BRO1 and BRO2. The most straightforward method to address this problem is to configure BGP route aggregation on routers CRO1 and CRO2 for the 10.16.0.0/12 and 10.32.0.0/12 address blocks. In order to suppress the advertisement of the more-specific prefixes, you should use the **no-summary** command option. Issue the following commands on routers CRO1 and CRO2:

```
router bgp 65000
  aggregate-address 10.16.0.0 255.240.0.0 summary-only
  aggregate-address 10.32.0.0 255.240.0.0 summary-only
```

As an alternative, you can also configure a prefix list or route map on routers BRO1 and BRO2 to drop all prefixes except for the two major blocks as the updates are received. However, this method is considered less efficient because it still causes the updates to be sent to routers BRO1 and BRO2 even if these routers discard them immediately after receiving them.

Student Notes

Use this Student Notes section to write down any alternate troubleshooting methods and additional troubleshooting commands that you learned during the labs and lab reviews.

[illegible]

Student Notes

Use this Student Notes section to write down any alternate troubleshooting methods and additional troubleshooting commands that you learned during the labs and lab reviews.

[illegible]

Lab 6-1 Answer Key: Introduction to Network Security

When you complete this activity, your documented solutions for the trouble tickets will be similar to the results here, with differences that are specific to your device or workgroup:

Note	If you have solved the problems described in the trouble tickets by using commands other than the ones listed here, bring your alternative solution to the attention of the instructor and the group during the lab debrief discussion.

Trouble Ticket R

There are two separate issues in this trouble ticket. The access-list on routers IRO1 and IRO2 that filters traffic from the Internet is misconfigured, causing HTTP and HTTPS return traffic to be dropped. This misconfiguration affects connectivity from all devices to the Internet, which means that this ticket will need to be resolved to finish resolving Trouble Tickets B and C.

The second problem, which you will discover while tracking the path of the traffic from client PC CLT1 to the Internet, is that the exec-timeout on the console of router IRO1 has been set to 1 second. As a result, you will be logged out of the session on the console of router IRO1 anytime you stop typing for a second. This setting only affects the console sessions, not Telnet or SSH. As a result, you can log in through means of Telnet or SSH to change the exec timeout on the console to a more reasonable setting. For example, to change the timeout to the same setting that is used on other routers issue the following commands to set it to one hour (60 minutes):

```
line console 0
exec-timeout 60 0
```

Given the more restrictive security settings, you can argue that the timeout should be changed to something shorter, such as 5 minutes. If you decide to do so, you should change this setting in the same way on all other devices for sake of consistency.

The problem with the access list is that the lines that are supposed to permit HTTP and HTTPS access from all subnets (except the branch office 1 subnet) have been configured to match destination TCP ports 80 and 443. However, because this access list matches return traffic coming from the Internet, these lines (which are numbered 140 through 170) should match source port 80 and 443 instead.

To correct this problem issue the following commands on routers IRO1 and IRO2:

```
ip access-list extended IN-FROM-INTERNET
no 140
140 deny tcp any eq www 10.1.160.64 0.0.0.63
no 150
150 deny tcp any eq 443 10.1.160.64 0.0.0.63
no 160
160 permit tcp any eq www 10.1.128.0 0.0.127.255
no 170
170 permit tcp any eq 443 10.1.128.0 0.0.127.255
```

Various other solutions are possible, but this solution is the one that stays the closest to the original implementation. It is important that you do not only change lines 160 and 170, because if you do, clients on subnet 10.1.160.64/26 will also gain access to the Internet, which is expressly forbidden by the security policy.

Trouble Ticket S

This ticket also consists of two issues. The main problem in this ticket revolves around an access list problem on routers CRO1 and CRO2 that causes DNS traffic from the branch office to the headquarters DNS server to be dropped. Clearly, this problem will affect all connectivity from client PC CLT2 and client PC CLT3. Users can still initiate connections based on the IP address, but not based on hostnames. As a result, this ticket will need to be resolved before ticket C can be fully resolved.

The second problem will again be discovered during the troubleshooting process. AAA has been configured on routers CRO1 and CRO2 to authenticate against a central database on server SRV1 using the TACACS+ protocol. On router CRO2, the key that is used to secure the TACACS traffic between the router and the server has been misconfigured. This misconfiguration causes all attempts to log on to router CRO2 to fail.

To solve this issue, you will first have to perform a password recovery procedure on router CRO2 so that you can change the configuration. (In this case, the procedure is not used to change the passwords themselves, but to access privileged mode and configure the correct key for the TACACS+ communication.)

To recover the password on router CRO2 the following steps should be taken from a console connection on the router.

First, the router needs to be powered off and back on. Within 60 seconds of power-up, send a “break” signal by pressing the appropriate key for your terminal emulation program. This step will put you in ROM Monitor mode. From this point, issue the following command sequence on router CRO2:

```
confreg 0x2142
reset
```

The router will now boot and ignore the configuration stored in NVRAM based on the changed setting of the configuration register. Wait until the router has fully booted and you are presented with the following prompt:

Would you like to enter the initial configuration dialog? [yes/no]:

Answer **No** to this question or press **Ctrl-C** to abort this dialog. Press **Return** to get started, and you should then enter the following command sequence:

```
enable
copy startup-config running-config
configure terminal
interface FastEthernet 0/0
no shutdown
interface FastEthernet 0/1
no shutdown
interface Serial 0/0/0
```



```
no shutdown
interface Serial 0/0/1
no shutdown
config-register 0x2102
```

At this point, you have restored the original configuration of router CRO2 and you are still in configuration mode. Now you can configure the correct TACACS+ key by issuing the following command:

```
tacacs-server key cisco
```

The correct key value of “cisco” can be retrieved from the TACACS+ server, or alternatively, you could have guessed this key by examining the configuration of router CRO1, which contains the same key.

The DNS problem is caused by a mistyped IP address in the first two lines of the access list on routers CRO1 and CRO2 that is supposed to permit DNS traffic to the IP address of server SRV1, which is 10.1.152.1. Instead, the access list has been configured to permit traffic to a different IP address: 10.1.252.1.

You can resolve this issue by issuing the following commands on routers CRO1 and CRO2:

```
ip access-list extended LIMIT-HQ-ACCESS
no 10
10 permit udp 10.1.128.0 0.0.127.255 host 10.1.152.1 eq domain
no 20
20 permit tcp 10.1.128.0 0.0.127.255 host 10.1.152.1 eq domain
```

Various other solutions are possible, but this solution is the one that stays the closest to the original implementation. If you have questions about your solution, bring them up during the lab debrief discussion.

Trouble Ticket T

Similar to tickets R and S, this ticket also consists of an access list issue, which causes the pings from client PC CLT2 to fail, and a password problem, which prevents you from accessing router BRO1 during troubleshooting.

Be aware that after resolving Trouble Tickets A and B, client PC CLT2 will have regained connectivity to server SRV1. The only problem that remains is that you cannot ping anywhere from this PC. (In addition, Internet access is still not working, but this situation is in compliance with the security policy and it should not be changed).

When you try to access router BRO1, you notice that you cannot enter privileged mode using the enable secret password “cisco”. To recover from this situation you should follow the same password recovery procedure used on router CRO2 in Trouble Ticket B. Only, this time you do not conclude the procedure by changing the TACACS+ key, but you conclude the procedure by changing the enable secret password by issuing the following command:

```
enable secret cisco
```

The issue that causes pings from client PC CLT2 to fail is that the access lists on routers BRO1 and BRO2 do not permit the ICMP echo and echo-reply messages. To enable this capability, you can add two more lines to the access list that is applied to the office VLAN on routers BRO1 and BRO2. You can accomplish this change to the access list by issuing the following commands:

```
ip access-list extended LIMIT-OFFICE-ACCESS
72 permit icmp 10.1.160.64 0.0.0.63 any echo
74 permit icmp 10.1.160.64 0.0.0.63 any echo-reply
```

You can argue whether adding these extra lines complies with the security policy, but given the fact that all other configured access lists also permit these messages, the conclusion must be that adding the extra lines is allowed.

Student Notes

Use this Student Notes section to write down any alternate troubleshooting methods and additional troubleshooting commands that you learned during the labs and lab reviews.

[illegible]

Student Notes

Use this Student Notes section to write down any alternate troubleshooting methods and additional troubleshooting commands that you learned during the labs and lab reviews.

This image shows a single page of white paper with horizontal ruling lines. The lines are evenly spaced and run across the width of the page, providing a guide for handwriting or typing. There are no margins, text, or other markings on the page.

Lab 6-2 Answer Key: Cisco IOS Security Features

When you complete this activity, your documented solutions for the trouble tickets will be similar to the results here, with differences that are specific to your device or workgroup:

Note	If you have solved the problems described in the trouble tickets by using commands other than the ones listed here, bring your alternative solution to the attention of the instructor and the group during the lab debrief discussion.

Trouble Ticket U

This trouble ticket contains a number of separate issues. The fact that client PC CLT2 cannot obtain an IP address is caused by the fact that DHCP snooping has been configured for VLAN 17 on switch BSW1, but the uplink ports that lead to routers BRO1 and BRO2 have not been configured as trusted ports. This configuration prevents any DHCP requests from being forwarded on these ports and, as a result, client PC CLT2 will not obtain an IP address via DHCP from its assigned DHCP server, router BRO1.

To resolve this issue you should configure the uplink ports on switch BSW1 as trusted ports by issuing the following commands:

```
interface range FastEthernet 0/3 - 4
ip dhcp snooping trust
```

After you have made this change, client PC CLT2 should be able to receive an IP address via DHCP and get access to the network. You should be able to ping most routers and switches from this client at this point, but to obtain access to server SRV1 and to the Internet, several other issues need to be resolved.

The reason the client PC CLT3 cannot obtain an IP address via DHCP is not related to the DHCP snooping feature, which is only enabled for VLAN 17 and not for VLAN 19, which client PC CLT3 is a member of. This problem is caused by the VLAN access map named PROTECT-AGAINST-WORMS that has been configured for VLAN 19 to block potential worm traffic. This VLAN access map drops all broadcast and multicast traffic, except for the traffic that is explicitly permitted by the access-list named NECESSARY-BROADCASTS. This access list allows GLBP traffic, but does not contain a line that permits DHCP traffic. To add the DHCP traffic as an exception to the rule that all unnecessary broadcasts and multicasts are dropped, you can issue the following commands on switch BSW1:

```
ip access-list extended NECESSARY-BROADCASTS
permit udp any any eq bootps
permit udp any any eq bootpc
```

After you make this change, client PC CLT3 should be able to receive an IP address via DHCP and get access to the network. Again, you should be able to ping most routers and switches from the client at this point. However, unless one of your team members has resolved the issue as part of one of the other trouble tickets, client PC CLT3 will not be able to browse to <http://www.isp3.local>.

Because server SRV1 serves as the DNS server for your network, you will not be able to connect to any device by name until trouble ticket W is resolved. However, at this point, you cannot connect to the web server by IP address either. Browsing to <http://172.34.224.1> also fails. This problem is caused by two errors in the Cisco IOS Firewall configuration on routers IRO1 and IRO2. The inspection policy on routers IRO1 and IRO2 is erroneously applied in the inbound direction on the interface facing the Internet instead of in the outbound direction. Secondly, the TCP protocol is not defined in the inspection policy. As a result, the TCP-based HTTP sessions are not inspected and return web traffic is dropped by routers IRO1 and IRO2. To correct the inspection policy on router IRO1, you can issue the following commands:

```
ip inspect name INTERNET-TRAFFIC tcp
interface FastEthernet 0/0
no ip inspect INTERNET-TRAFFIC in
ip inspect INTERNET-TRAFFIC out
```

The configuration on router IRO2 needs to be adjusted in a similar manner, by issuing the following commands:

```
ip inspect name INTERNET-TRAFFIC tcp
interface FastEthernet 0/1
no ip inspect INTERNET-TRAFFIC in
ip inspect INTERNET-TRAFFIC out
```

From this point on, you should be able to ping and browse to IP address 172.34.224.1. Connecting by name will not be possible until you complete Trouble Ticket W.

Trouble Ticket V

The fact that client PC CLT1 cannot obtain an IP address via DHCP is caused by an underlying Layer 2 problem. Switch CSW1 serves as the DHCP server for VLAN 17, which client PC CLT1 is a member of. On switch CSW1, the root guard feature is enabled on both Port-channel interfaces 1 and 10. Because switch CSW2 claims to be the root for VLAN 17 (which is legitimate according to the design), spanning tree is in the “broken” state for VLAN 17 on both Port-channel interfaces. As a result, no VLAN 17 traffic from switch ASW1 can reach switch CSW1 and DHCP fails. To correct this situation you should remove the root guard feature from the Port-channel 10 interface by issuing the following commands on switch CSW1:

```
interface Port-channel 10
no spanning-tree guard root
```

This change restores the connectivity between client PC CLT1 and its DHCP server, switch CSW1, and you should be able to obtain an IP address on client PC CLT1 after this. Assuming that you have fixed the Cisco IOS Firewall issues from the previous trouble ticket, you should now also be able to ping and browse to IP address 172.34.224.1 (www.isp3.local).

You even can argue that the spanning-tree root guard feature should also be removed from Port-channel 1 on both switches CSW1 and CSW2. The security policy states, “Protection against users plugging in switches, which might attempt to become the root for the Spanning Tree Protocol,” which can be interpreted to mean that the root guard feature should only be enabled on user-facing ports and not on interswitch links.

Trouble Ticket W

The main problem in this ticket is caused by the fact that the port security feature has been enabled on the port on switch CSW1 that connects to server SRV1, and the wrong secure MAC address has been configured on this port. This configuration causes switch CSW1 to record a security violation when it starts receiving frames from server SRV1, and it will error-disable the port. To correct this issue, you should remove the misconfigured secure MAC address from the port. To make this change, issue the following commands on switch CSW1:

```
interface FastEthernet 0/10
shutdown
no switchport port-security mac-address 0000.0cc0.ffee
no shutdown
```

In addition, you may want to add the dynamically learned MAC address of server SRV1 to the configuration manually as a static secure address. The guidelines in the security policy are not clear enough to know for sure if this change is intended.

After making these changes, server SRV1 should be reachable again and, as a result, DNS starts functioning. From this point onward, you should be able to test connectivity by name again.

Trouble Ticket X

The problem in this ticket is caused by the fact that SSH requires a key pair to be generated on the routers and switches. The SSH configuration itself is correct, but without a public and private key the device is not ready to accept SSH connections. To generate a key pair to be used for SSH, you can issue the following command sequence on all devices:

```
crypto key generate rsa general-keys modulus 1024
```

Although this command is executed in global configuration mode, the key data is not stored in the startup configuration, but in a separate section of NVRAM. However, to preserve the key data across reboots, it is still necessary to issue the **copy running-config startup-config** command to save the key data.

The security policy did not specify the key length to be used. In the example, a length of 1024 bits was selected, but other lengths are equally valid.

Trouble Ticket Y

The problem in this ticket is caused by the fact that the macro is used to enable the port security feature on all ports, including the uplink ports to routers BRO1 and BRO2. This enablement creates an issue, because the routers use GLBP, which makes use of special virtual MAC addresses. When one of the routers is rebooted and reinitializes, it will initially claim the same virtual MAC address that is also used by the other router. The port-security feature records the same secured MAC address on both uplink ports, perceives this action as a security violation, and disables the uplink ports as a result. Consequentially, port security cannot be used on these two ports.

To resolve this issue, instead of applying the macro to all ports, you should apply the macro to all ports except the two uplink ports Fa 0/3 and Fa 0/4. If you have already applied the macro, you can also disable port-security on these two ports and re-enable them as follows:

```
interface range FastEthernet 0/3 - 4
shutdown
no switchport port-security
default switchport port-security maximum
no macro description
no shutdown
```


Student Notes

Use this Student Notes section to write down any alternate troubleshooting methods and additional troubleshooting commands that you learned during the labs and lab reviews.

[illegible]

Student Notes

Use this Student Notes section to write down any alternate troubleshooting methods and additional troubleshooting commands that you learned during the labs and lab reviews.

This image shows a blank sheet of white paper with horizontal ruling lines. The lines are evenly spaced and extend across the width of the page, providing a guide for writing. There are no margins, text, or other markings on the paper.

Lab 7-1 Answer Key: Troubleshooting Complex Environments

Caution Do not look at these answers before you have completed Lab 7-1. Looking at the suggested answers will invalidate your lab score and both you and your team will lose an important opportunity to assess the troubleshooting skills acquired during this course.

When you complete this activity, your documented solutions for the trouble tickets will be similar to the results here, with differences that are specific to your device or workgroup:

Note If you have solved the problems described in the trouble tickets by using commands other than the ones listed here, bring your alternative solution to the attention of the instructor and the group during the lab debrief discussion.

Trouble Ticket A

The connectivity problem between client PC CLT1 and server SRV1 was caused by a mismatch in configuration between the Port-channel 10 interface and interfaces FastEthernet 0/3 and 0/4 on switch CSW1. On the Port-channel interface, an allowed VLAN list was configured, which was missing on the physical interfaces. This problem causes the physical interfaces to be suspended and the Port-channel interface to go down as a result. To resolve this matter issue the following commands on switch CSW1:

```
interface range FastEthernet 0/3 - 4
  switchport trunk allowed vlan 17-19,21-23,25-27,33-35,37-39
  switchport trunk allowed vlan add 41-43,111,112,128-130
```

Instead of adding these commands to interfaces FastEthernet 0/3 and 0/4, you can also restore connectivity between client PC CLT1 and server SRV1 by removing the allowed VLAN list from the Port-channel 10 interface. However, this solution does not comply with the established policies on the network, because allowed VLAN lists are configured on all other trunks between the switches. As a result, using this solution will cause you to lose points for policy compliance in the Network Maintenance section of this lab.

Trouble Ticket B

The solution of this ticket consists of two parts. First, the system MTU of switch CSW2 needs to be reset to 1500 in order to ensure that the OSPF adjacencies between switch CSW2 and its neighbors are established. To change the system MTU issue the following command:

```
system mtu 1500
```

After you issue the command, the switch needs to be reloaded to have the change take effect.

The second issue that causes the connectivity from client PC CLT1 to the Internet to fail is the duplicate OSPF router ID between router IRO2 and router CRO2. This problem can be resolved by changing the OSPF router ID on CRO2 to the IP address that is configured on its loopback. To make this change, configure the following commands on router CRO2:

```
router ospf 1
  router-id 10.1.220.2
```

After you change the router ID, you need to reset the OSPF process using the **clear ip ospf process** command.

The obvious choice for the router ID on router CRO2 is the IP address of its loopback interface 10.1.220.2. Choosing any other unique router ID, or changing the router ID on router IRO2 to a unique value will also restore connectivity and is technically considered to be a solution. However, because all other routers use the IP address of their Loopback 0 interface as their OSPF router ID, this solution is not considered to be in compliance with established policies, and as a result, you will lose points in the Network Maintenance section of this lab.

Trouble Ticket C

This ticket introduced two problems. Only one of the two problems needs to be resolved to score the points for the ticket. However, to achieve network redundancy, both problems need to be addressed, which is necessary to score the points for the redundancy section in the network maintenance task.

To solve the PPP CHAP authentication problem, the AAA configuration on router BRO1 needs to be adapted to fall back to local authentication for PPP if RADIUS authentication fails. To make this change, issue the following command:

```
aaa authentication ppp default group radius local
```

To solve the EIGRP neighbor problem between routers BRO1 and CRO2, the control plane policy on router CRO2 needs to be changed to include EIGRP in the class defined for routing protocol traffic. The easiest way for you to change the policy is to add a line to the already defined extended access-list named ROUTING-PROTOCOLS by issuing the following commands:

```
ip access-list ex ROUTING-PROTOCOLS
permit eigrp any any
```

Any other solution that adds EIGRP to the class “ROUTING-TRAFFIC” is also valid.

Trouble Ticket D

There are two issues in this ticket. The first issue is that the configuration register has been set to 0x2100 on router BRO2 causing it to boot to ROM monitor mode instead of booting the Cisco IOS Software. To remedy this problem, two steps need to be taken: the router needs to be manually booted from the Cisco IOS Software in flash memory and the configuration register value needs to be reset to its default value of 0x2102. To boot router BRO2 from ROM monitor mode issue the following command:

```
boot flash:c1841-advipservicesk9-mz.124-23.bin
```

Replace the name of the Cisco IOS image file with the name of the file found in the flash of router BRO2. (You can use the **dir flash:** command to list the files in flash from ROM monitor mode.)

To reset the configuration register to its default value of 0x2102, issue the following command after the router has fully booted:

```
config-register 0x2102
```

Alternatively, the value can be changed from ROM monitor before booting the Cisco IOS Software by issuing the command:

```
confreg 0x2102
```

The second problem in this ticket is that VLAN 19 has been configured as a remote SPAN VLAN on switch BSW1. To remedy this problem, issue the following commands on switch BSW1:

```
vlan 19
no remote-span
```

Network Maintenance

The only unresolved problem that was not covered in the previous trouble tickets is the failing peering between routers IRO1 and ISP1. This problem is caused by the fact that the routes received from router ISP1 are not correctly filtered. Instead of only accepting the default route, the route map that is supposed to filter the routes permits all routes received from router ISP1. This situation causes the maximum number of prefixes that is allowed to be received from router ISP1 to be exceeded, and the BGP session is closed almost immediately after it is established. To remedy this issue remove the line that permits all routes from route map FROM-AS-65525 on router BRO1 by issuing the following command:

```
no route-map FROM-AS-65525 permit 20
```

Removing the maximum number of allowed prefixes from the peering to router ISP1 also enables the peering. However, this solution does not conform to the established policies and you will lose points for the policy compliance section of this task.

Student Notes

Use this Student Notes section to write down any alternate troubleshooting methods and additional troubleshooting commands that you learned during the labs and lab reviews.


[illegible]

Student Notes


Use this Student Notes section to write down any alternate troubleshooting methods and additional troubleshooting commands that you learned during the labs and lab reviews.

[illegible]

Lab Topology




Lab 3-1: WAN Physical Topology




34P_749


Lab 3-1: Branch LAN Physical Topology



Lab 3-1: Headquarters Logical Topology




Lab 3-1: WAN Logical Topology




334P_746


Lab 3-1: Branch Logical Topology




Lab 5-2: OSPF Area Design




Lab 5-2: Branch Office Migration Design



Lab 5-3: BGP Design



<https://t.me/learningnets>



Lab 7-1: Routing Design

