

CISCO (642-832) TSHOOT

Cisco Certified Network Professional



**Smarter
Training**

LearnSmart's CCNP TSHOOT exam manual aims to equip network professionals with all of the skill sets necessary to pass Troubleshooting and Maintaining CISCO IP Networks (642-832), a qualifying exam for the CCNP certification. By studying this guide, candidates will become familiar with network concepts found in the exam, including:

- Maintaining and Monitoring Network Performance
- Troubleshooting IVR
- Troubleshooting Gateway Redundancy Protocols
- Troubleshooting Network Derogation Problems
- And more!

Give yourself the competitive edge necessary to further your career as a network professional and purchase this exam manual today!

CCNP TSHOOT (642-832) LearnSmart Exam Manual

Copyright © 2011 by LearnSmart, LLC.

Product ID: 012461

Production Date: November 16, 2011

All rights reserved. No part of this document shall be stored in a retrieval system or transmitted by any means, electronic, mechanical, photocopying, recording, or otherwise, without written permission from the publisher. No patent liability is assumed with respect to the use of the information contained herein.

Warning and Disclaimer

Every effort has been made to make this document as complete and as accurate as possible, but no warranty or fitness is implied. The publisher and authors assume no responsibility for errors or omissions. The information provided is on an “as is” basis. The authors and the publisher shall have neither liability nor responsibility to any person or entity with respect to any loss or damages arising from the information contained in this document.

LearnSmart Cloud Classroom, LearnSmart Video Training, Printables, Lecture Series, Quiz Me Series, Awdeco, PrepLogic and other PrepLogic logos are trademarks or registered trademarks of PrepLogic, LLC. All other trademarks not owned by PrepLogic that appear in the software or on the Web Site (s) are the property of their respective owners.

Volume, Corporate, and Educational Sales

Favorable discounts are offered on all products when ordered in quantity. For more information, please contact us directly:

1-800-418-6789

solutions@learnsmartsystems.com

International Contact Information

International: +1 (813) 769-0920

United Kingdom: (0) 20 8816 8036

Table of Contents

Domain 1: Maintain and Monitor Network Performance	8
What is Network Maintenance?	8
Network Maintenance Models	8
Network Administrator Responsibilities	9
Network Administrator Tools	10
<i>CLI Tools</i>	10
<i>GUI Tools</i>	12
Network Documentation Management Tools	16
<i>Trouble Ticket Reporting System</i>	16
<i>Wiki</i>	17
Added-Value Network Monitoring Tools	17
<i>SNMP Server</i>	18
<i>NetFlow Server</i>	19
Troubleshooting Methodology	21
<i>Root Cause Problem Resolution</i>	25
The Integration of Maintenance and Troubleshooting	25
A Closer Look at IOS CLI Troubleshooting Tools	26
<i>IOS Output Filtering Tools</i>	26
<i>Output Redirection</i>	29
<i>Troubleshooting Tools</i>	29
<i>Hardware Diagnostics</i>	32
Domain 2: Troubleshoot Multi-Protocol System Networks	36
Potential Layer 2 Switch Problems	36
<i>Hardware Problems</i>	36
<i>VLAN Configuration Problems</i>	36
<i>Trunk Configuration Problems</i>	36
Layer 2 Troubleshooting Techniques	36
<i>Show interfaces</i>	37
<i>Show interfaces counters</i>	38
<i>Show interfaces errors</i>	38
<i>Show mac-address-table</i>	39
<i>Show mac-address-table address</i>	39
<i>Clear mac-address-table dynamic</i>	39

<i>Show vlan brief</i>	40
<i>Show interfaces switchport</i>	40
<i>Show interfaces trunk</i>	41
<i>Traceroute mac</i>	41
Spanning Tree Troubleshooting Techniques	42
<i>Spanning Tree Stages</i>	43
<i>Show spanning tree vlan</i>	44
<i>Show spanning tree <interface> detail</i>	44
Common Spanning Tree Problems	45
<i>MAC Address Table Corruption</i>	45
<i>Broadcast Storm</i>	45
EtherChannel Troubleshooting Techniques	45
Troubleshooting InterVLAN Routing (IVR)	47
<i>Show ip route</i>	48
<i>Show ip arp</i>	48
<i>How MLS works from a Layer 2 and Layer 3 point of View:</i>	49
<i>Show ip cef</i>	49
<i>Show adjacency</i>	50
<i>Troubleshooting SVIs</i>	50
Troubleshooting Gateway Redundancy Protocols	50
<i>Hot Standby Routing Protocol (HSRP)</i>	51
<i>show standby <group></i>	51
<i>show standby brief</i>	52
<i>debug standby</i>	52
<i>Virtual Router Redundancy Protocol (VRRP)</i>	53
<i>show vrrp brief</i>	54
<i>Gateway Load Balancing Protocol (GLBP)</i>	54
<i>show glbp brief</i>	55
Troubleshooting Network Derogation Problems	55
<i>Straight-through vs. crossover</i>	55
<i>Speed/Duplex Mismatch</i>	56
<i>Switch Forwarding Processes</i>	57
Troubleshooting Switch Supervisor Redundancy	57
Troubleshooting IP Routing Protocols	58

<i>Basic Routing Concepts</i>	58
<i>Administrative Distance</i>	59
<i>Basic Routing Concepts</i>	60
Troubleshooting EIGRP	64
<i>EIGRP Neighbors</i>	64
<i>EIGRP Tables</i>	64
<i>EIGRP Metric</i>	65
Essential EIGRP Troubleshooting Commands	65
<i>Show ip eigrp interfaces</i>	65
<i>Show ip eigrp neighbors</i>	66
<i>Show ip eigrp topology</i>	66
<i>Show ip eigrp traffic</i>	67
<i>Debug ip eigrp</i>	67
<i>Debug eigrp packet</i>	68
Troubleshooting OSPF	68
<i>OSPF Tables</i>	69
<i>OSPF Metric</i>	69
<i>Purpose of the Designated Router (DR)</i>	69
<i>Purpose of OSPF Areas</i>	69
<i>Link State Algorithm Types</i>	70
<i>OSPF Network Types</i>	71
<i>OSPF Hellos</i>	71
<i>Adjacency Status States</i>	71
<i>Essential OSPF Troubleshooting Commands</i>	72
Troubleshooting BGP	78
<i>IGP vs. EGP</i>	78
<i>BGP Operation</i>	78
<i>BGP Metric</i>	78
<i>Troubleshooting BGP Neighbors</i>	79
<i>Essential BGP Troubleshooting Commands</i>	79
Troubleshooting Route Redistribution	84
<i>Why an Administrator would Redistribute Routes</i>	84
<i>The Boundary Router</i>	85
<i>Essential Route Redistribution Troubleshooting Commands</i>	85

Troubleshooting DHCP	86
<i>Purpose of DHCP</i>	86
<i>DHCP Operation</i>	87
<i>DHCP Messages</i>	88
<i>Common DHCP Problems</i>	88
<i>DHCP Troubleshooting Commands</i>	89
Troubleshooting NAT	91
<i>Purpose of NAT</i>	91
<i>Four Types of NAT</i>	91
<i>NAT Topology</i>	91
<i>Common NAT Problems</i>	92
<i>NAT over VPN</i>	92
<i>Latency due to NAT Processing Overhead</i>	92
<i>Applications that Do Not Play Well with NAT</i>	93
<i>NAT hiding True Source or Destination</i>	93
<i>NAT troubleshooting Commands</i>	93
Troubleshooting VoIP	94
<i>Voice Network Requirements</i>	95
<i>Voice Network Requirements</i>	95
<i>Common Voice Problems</i>	96
<i>Not Receiving Necessary Services (i.e. CDP, DHCP or TFTP)</i>	97
<i>Misconfiguration of QoS</i>	97
<i>Misconfiguration of Voice VLANs on Access Ports</i>	98
<i>Insufficient Power from PoE Source</i>	98
<i>Voice Troubleshooting Commands</i>	99
Troubleshooting Video	100
<i>Troubleshooting Multicasting for Video Applications</i>	101
<i>Common Video Problems</i>	102
<i>Multicast Video Troubleshooting Commands</i>	103
Troubleshooting Wireless Networks	104
<i>Wireless Components</i>	104
<i>Wireless Architectures</i>	105
Troubleshooting Advanced Services	107
<i>NetFlow Overview</i>	107

<i>NetFlow Troubleshooting Commands</i>	107
<i>Overview of IP Service Level Agreement (IP SLA)</i>	108
<i>IP SLA Troubleshooting Commands</i>	108
<i>Overview of Network-based Application Recognition (NBAR)</i>	109
<i>NBAR Troubleshooting Commands</i>	110
Troubleshooting Security	110
<i>Securing Router and Switch Planes</i>	111
<i>Management Plane Security</i>	111
<i>Management Plane Security Troubleshooting Commands</i>	112
<i>Control Plane Security</i>	112
<i>Securing Spanning Tree</i>	113
<i>Routing Service Security</i>	114
<i>Gateway Redundancy Authentication</i>	115
<i>Securing DHCP</i>	115
<i>Securing ARP</i>	115
<i>Control Plane Security Troubleshooting Commands</i>	116
<i>Data Plane Security</i>	117
<i>Access Control Lists</i>	118
<i>IOS Firewall Feature Set</i>	118
<i>VPN Tunnels</i>	118
<i>802.1x Authentication</i>	118
<i>Data Plane Security Troubleshooting Commands</i>	118
Troubleshooting IPv6	120
<i>IPv6 Overview</i>	120
<i>IPv6 Header</i>	121
<i>IPv6 Traffic Types</i>	121
<i>IPv6 Routing Methods</i>	121
<i>IPv6 to IPv4 Communication</i>	122
<i>IPv6 Address Format</i>	122
<i>IPv6 OSPFv3 Troubleshooting Commands</i>	123
<i>IPv6 RIPng Troubleshooting Commands</i>	126
Practice Questions	128
Answers & Explanations	135

Domain 1: Maintain and Monitor Network Performance

What is Network Maintenance?

The first thing you should learn about troubleshooting Cisco networks is how to help avoid problems in the first place. The TSHOOT exam expects the test taker to know how to properly maintain route and switch equipment so it has the latest revision to avoid any kind of software related problems due to bugs or inoperable images. A router or switch that is kept up to date with the latest general deployment IOS image has the following advantages:

- Stability
- Fewer bugs
- Easier and quicker to repair

So effectively, network maintenance is doing the preemptive steps required to keep a network running and at the same time, meeting the current and near-term needs of the organization. Examples of network maintenance include:

- Hardware and software installation/configuration according to best-practices
- Troubleshooting network problems proactively and through trouble tickets
- Monitoring using network tools such as syslog and SNMP
- Tuning the network for additional performance
- Planning for network expansion
- Network documentation for both current and future needs
- Compliance with legal regulations and organizational policies
- Network security

As hinted earlier, there are two different troubleshooting triggers:

- **Proactive** – planned network maintenance base on general upkeep and future expansion.
- **Reactive** – unplanned maintenance commonly triggered by responding to problems found by network monitoring tools or through trouble tickets from end-users.

Network Maintenance Models

There are several, well-known network maintenance models that you need to know for the TSHOOT exam. Those models include:

- **Fault management, Configuration management (FCAPS)** – an ISO model that specifies management for Accounting, Performance, and Security.
- **IT Infrastructure Library (ITIL)** – a United Kingdom developed model that provides detailed checklists, tasks and processes that can be easily fitted into most businesses.
- **Telecommunications Management Network (TMN)** – an ITU-T developed model for management of telecommunication networks. It is a variation of FCAPS that focuses more on voice and telecom networks.

- **Cisco Lifecycle Services** – outlines a Cisco network according to the technology lifecycle for the hardware. This is a cycle based plan that defines the following phases:
 - ▶ Preparation
 - ▶ Planning
 - ▶ Design
 - ▶ Implementation
 - ▶ Operation

Because of the five distinct lifecycle phases, it is also known as the **PPDIO** model.

Network Administrator Responsibilities

The role of a Cisco network administrator is to handle the maintenance and upkeep of the current network. This includes network related tasks such as:

- MAC-D: Moves, Adds, Changes and Deletions
- Writing network documentation related to the upkeep and future growth of the network
- Disaster Recovery (DR) planning and testing
- Network monitoring
- Troubleshooting hardware/software problems and incompatibility issues from end-users
- Scheduled maintenance including Change Control documentation
- Fallback planning in the case of a maintenance failure

Change Control documentation is a critical task in regards to network maintenance. Most large organizations have change control policies which have levels of risk to the operation of the business. These levels commonly state days of the week and times that maintenance can be performed based on the level of importance a piece of network equipment has.

Often times, Change Control systems are implemented where the network administrator requests that network maintenance be performed. The Change Control System then notifies the proper people within the organization to openly discuss the maintenance plan and approve or deny the change.

The importance of proper Network Documentation cannot be stressed enough in a large organization with many locations, device types and level of importance to an organization. Documentation tasks include:

- Create configuration “boilerplates” or “templates” that should be universally used on like-devices
- Organize a configuration history for easier troubleshooting in the event of a configuration error
- Create a production and spare equipment inventory to easily identify network hardware in use and available as a spare in case of a failure. Having serial numbers is handy when creating trouble tickets with Cisco TAC
- Create a PSTN and other WAN circuit inventory (including circuit ID and ISP/PSTN contacts)
- Create an IP address assignment spreadsheet for easily identifying subnets in use and available subnets for future growth

- Create network diagrams to better grasp traffic flow. This is also very useful for new engineers to quickly learn the network setup. Network diagrams should show both physical and logical characteristics.
- Create an Out-of-band communications spreadsheet that lists how to access devices remotely in the case of a major failure.
- Create a Network Communications plan that details the “who, what, when, where and why” when a network outage occurs.

Configuration Boilerplates (or templates) are either full or partial network configurations that an engineer is completely satisfied with and would like to be used on all equipment that performs identical tasks such as an access-layer switch configuration. These boilerplates will have clearly identified fill-in-the-blank areas to customize the configuration on another device.

A **Network Communications Plan** is typically a shared spreadsheet used by the organization that specifies the following when a network outage occurs:

- Who should be contacted in the internal network department when a specific outage occurs?
- Defines the critical value of a piece of network equipment to the organization. Based on this, the plan will specify the times that people should be contacted depending on importance of the device.
- Who the IT managers are in the case where an outage needs to be escalated.
- Vendor names, numbers and email addresses in the case where external help is required to fix a problem.

Network Administrator Tools

Network maintenance tools vary widely both in terms of monetary cost to the organization as well as their usefulness to the administrator. Let’s briefly look at the more common network administrator tools based on tools found built-in to Cisco devices using either the command line interface (CLI), graphical user interface (GUI) or popular external application tools used on networks.

CLI Tools

The Cisco IOS provides many useful tools to troubleshoot hardware and software problems. Most administrators will find that a large percentage of their time will be spent troubleshooting from the CLI.

The **show command** is useful to view various statistics of an IOS device including things such as interfaces, routing protocols, CPU/memory, startup/running configurations and logging.

The **debug command** is enabled or disabled on an IOS device to provide real-time information coming from the Cisco hardware. For example, if you are having problems with your EIGRP routing process, you can issue the following command and interrupt the output as shown below:

```

Router# debug ip eigrp
IP-EIGRP: Processing incoming UPDATE packet
IP-EIGRP: Ext 192.168.3.0 255.255.255.0 M 386560 - 256000 130560 SM
360960 - 256000 104960
IP-EIGRP: Ext 192.168.0.0 255.255.255.0 M 386560 - 256000 130560 SM
360960 - 256000 104960
IP-EIGRP: Ext 192.168.3.0 255.255.255.0 M 386560 - 256000 130560 SM
360960 - 256000 104960
IP-EIGRP: 172.24.43.0 255.255.255.0, - do advertise out Ethernet0/1
IP-EIGRP: Ext 172.24.43.0 255.255.255.0 metric 371200 - 256000 115200
IP-EIGRP: 192.135.246.0 255.255.255.0, - do advertise out Ethernet0/1
IP-EIGRP: Ext 192.135.246.0 255.255.255.0 metric 46310656 - 45714176
596480
IP-EIGRP: 172.24.40.0 255.255.255.0, - do advertise out Ethernet0/1
IP-EIGRP: Ext 172.24.40.0 255.255.255.0 metric 2272256 - 1657856
614400
IP-EIGRP: 192.135.245.0 255.255.255.0, - do advertise out Ethernet0/1
IP-EIGRP: Ext 192.135.245.0 255.255.255.0 metric 40622080 - 40000000
622080
IP-EIGRP: 192.135.244.0 255.255.255.0, - do advertise out Ethernet0/1

```

You should also be familiar with a fairly new IOS feature that is similar to debug in the fact that the IOS monitors real-time events. The difference is that the IOS has intelligence built-in to automatically identify problems and correct them without human intervention. This feature is called the **Embedded Event Manager (EEM)**.

The Cisco IOS can be used to create a **Configuration History** by copying the changed configurations to a TFTP server and labeling them based on the date the change was made. To accomplish this, you need a dedicated TFTP server. Then it's just a matter of changing the configuration name to include the date and device name as shown using a TFTP server at the IP of 192.168.10.100:

```

Router10#copy run tftp
Address or name of remote host []? 192.168.10.100 Destination
filename [Router10-config]? 2010-05-01-Router10.txt
!! 860 bytes copied in 2.228 secs (362 bytes/sec)

```

Another great tool that can be used with troubleshooting devices remotely using Telnet or SSH is the **terminal monitor** command. This command lets you view console messages as if you were directly connected to the device's console port. This information can be extremely valuable when troubleshooting. This is great, but sometimes the messages quickly scroll across the screen. An alternative is to use the **logging buffered <severity>** command. This writes the log messages to the IOS device memory. The memory is finite in nature and when full, it will rewrite the oldest messages first. The <severity> specifies the type of messages you want to be stored to the buffer. The severity levels range from 0 to 7.

This table shows the levels and names:

Logging Severity Number	Logging Severity Name
0	Emergencies
1	Alerts
2	Critical
3	Errors
4	Warnings
5	Notifications
6	Informational
7	Debugging

Figure 1: Logging Severity Levels

Please note that the default logging severity level is 7 (Debugging). Once enabled, the buffer can be viewed using the CLI by issuing the **show log** command.

GUI Tools

Cisco also provides several free and add-on GUI tools for maintaining and troubleshooting equipment. These tools include:

- Cisco Works
- The Cisco Configuration Professional (CCP)
- The Cisco Configuration Assistant (CCA)
- The Cisco Network Assistant (CNA)
- The Cisco Security Device Manager (SDM)

Here is a screenshot of the Cisco Configuration Assistant GUI tool used to monitor Cisco Smart Business Communication Systems Devices such as the UC500 platform as seen here:

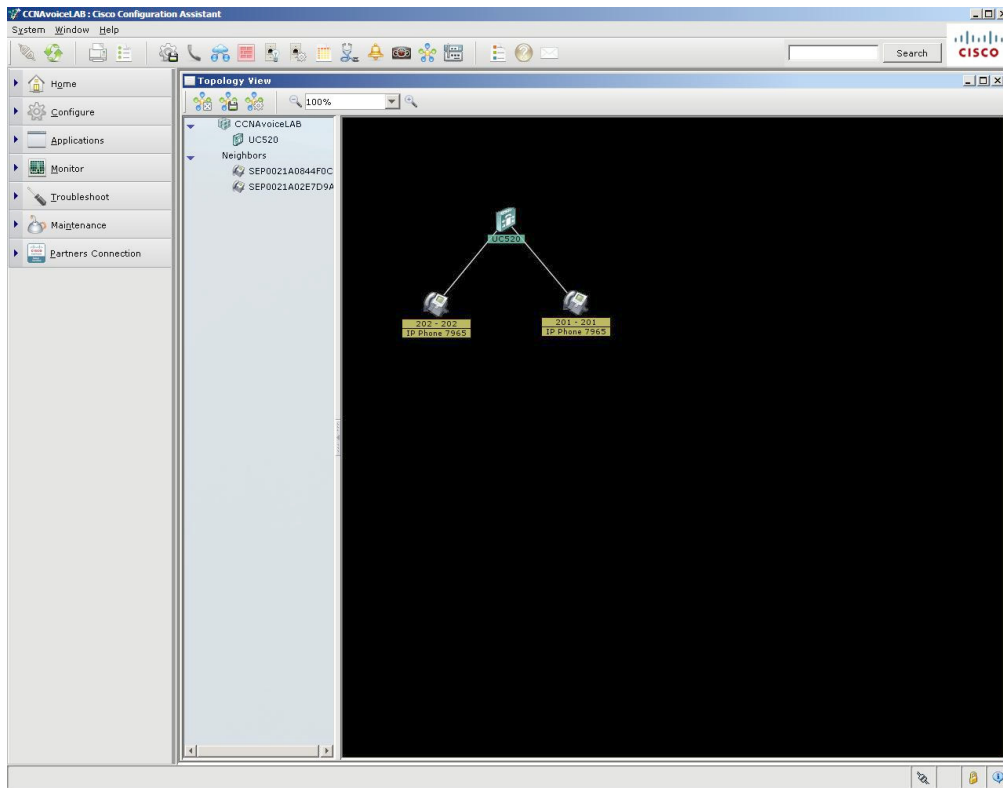


Figure 2: The Cisco Configuration Assistant

External Tools

As mentioned above, A TFTP server is a great external tool that can be used with the IOS (or GUI) to maintain configuration files as well as keep track of current and past IOS images used on the network. Other tools can also be used by administrators to help with maintenance and troubleshooting of a network including:

- FTP Servers
- NTP Servers
- Syslog Servers

Let's look at each of these tools individually to better understand how they can be used to help maintain a Cisco Network.

FTP Servers

FTP servers can be used to backup configuration files automatically. This differs from the above TFTP backup example shown earlier, which was a manual process. To configure an automatic backup, you need an FTP server and login credentials in the form of a username and password that is already configured on the FTP server. Once the FTP server is setup, an administrator can use the CLI of an IOS device to backup configuration files on a regular basis such as every night.

Here's an example of how this is configured:

```
Router10#config t
Enter configuration commands, one per line. End with CNTL/Z.
Router10(config)# ip ftp username andy
Router10(config)# ip ftp password cisco
Router10(config)#archive
Router10(config-archive)#path ftp://192.168.10.100/Router10-config
Router10(config-archive)#time-period 1440
```

This configuration sets an FTP username (andy), password (cisco) and then sets the path where the FTP server is (192.168.10.100) as well as the directory the file should be backed up to. The time-period command is set for 1440 seconds, which means that an archive will be kicked off every 24 hours from the time the command was set. Also note that this same archive feature can be used with a TFTP server with the difference being that the TFTP server does not require any kind of username/password authentication to upload the files.

You can then view your archived configurations on the IOS by issuing the following command:

```
Router10# show archive
      The next archive file will be named ftp:/ /192.168.10.100 74/
Router10-config-3
Archive #  Name
0
1      ftp:/ /192.168.10.100 74/Router10-config-1
2      ftp:/ /192.168.10.100 74/Router10-config-2 <- Most Recent
3
4
5
6
7
8
9
10
11
12
13
14
```

If you've discovered that you've made a configuration mistake that might be difficult to correct manually, you can easily fallback to an archived configuration. Note that by rolling back to the last archived configuration, you are completely replacing the current configuration with the archived one. There is no merging of configuration files. Here is how to fallback on an IOS device:

```
Router10#configure replace ftp://192.168.10.100/Router10-config-2

This will apply all necessary additions and deletions to replace
the current running configuration with the contents of the specified
configuration file, which is assumed to be a complete configuration, not
a partial configuration. Enter Y if you are sure you want to proceed. ?
[no]: y

Loading Router10-config-2 from 192.168.10.100 (via FastEthernet0/0):!!!

[ OK - 3113/4096 bytes]
```

NTP Servers

New network administrators will quickly come to the conclusion that a vital part of troubleshooting network problems using logging information is to insure that the times are identically set on network equipment. You can use log messages on different devices and compare logs when problems occur. Correlating events is much easier when clocks and log timestamps match up. You should configure the Network Time Protocol (NTP) on all of your devices to make sure that times are identical and do not drift which can occur when setting the clock locally.

NTP allows routers to point to a device acting as an NTP server. This can either be another IOS device or a separate server configured for NTP. Keep in mind that the NTP server might be used by devices in different time zones. Therefore, each device has its own time zone configuration set which indicates how many hours its time zone differs from Greenwich Mean Time (GMT). An example is -6, which is the GMT offset for the Central time zone in the USA. Here's an example of how to configure basic NTP services on an IOS device that participates in daylight savings:

```
Router(config)#clock timezone CST -6
Router(config)#ntp server 10.249.1.100
Router(config)#clock summer-time CDT recurring 2 Sun Mar 2:00 1 Sun Nov
2:00
```

Once configured, you can verify that NTP is working by issuing this command:

```
Router#show ntp status
Clock is synchronized, stratum 3, reference is 10.249.1.100 nominal
freq is 250.0000 Hz, actual freq is 249.9981 Hz, precision is 2**18
reference time is CDB93371.D03D5B6B (08:43:29.813 EDT Sat May 1 2010)
clock offset is -6.7340 msec, root delay is 83.56 msec root dispersion
is 57.02 msec, peer dispersion is 15.29 msec
Router#
```

Syslog Servers

Syslog servers are used as a central repository for all networking logging information that is generated by your equipment. This information can be very valuable for determining network issues. We've reviewed the terminal monitor and logging buffered commands previously. The logging command can also be used to offload messages to a dedicated logging server. This way, there is one central location to check all logs and since servers can store much more information, your logs can be saved much further back in time as opposed to logging messages locally which has a much smaller storage footprint. Configuring a syslog server is very easy as shown here:

```
Router(config)# logging 10.200.11.11
```

Now you can verify the syslog server at 10.200.11.11 to make sure that log messages are being sent. A very popular syslog server is Kiwi Syslog. Here is a screenshot that show the logging messages coming into the server:

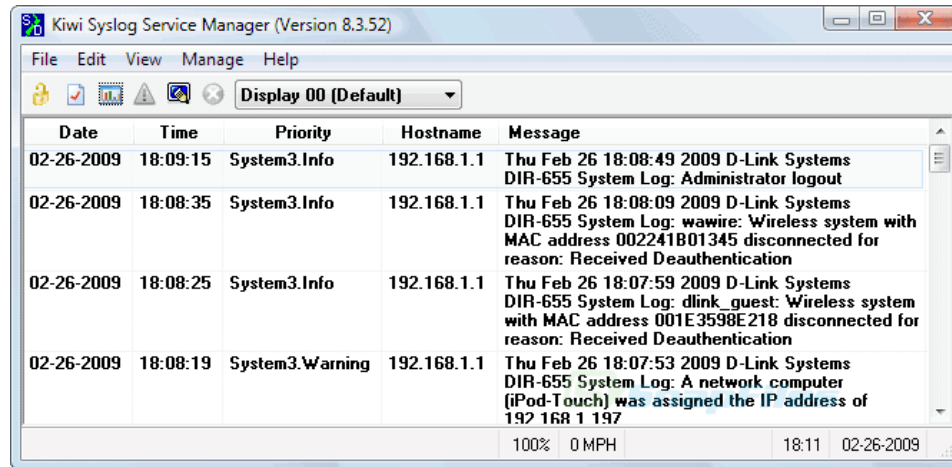


Figure 3: Kiwi Syslog Server

Network Documentation Management Tools

There are a few documentation management tools that a TSHOOT candidate should be aware of. These tools are Trouble Ticket Reporting Systems and Wiki's. Let's briefly look at each of them:

Trouble Ticket Reporting System

In large organizations, when a customer has a network problem, they typically will call a centralized IT helpdesk number. The helpdesk will do the following:

- Document the reported problem from the user's perspective.
- Take the customer's contact information such as the customer's name, telephone number, description of the device (and device ID if available), IP address, MAC address, building name and network jack ID.
- Perform initial troubleshooting to see if the problem can be quickly resolved.

If the helpdesk employee determines that a network engineer should look into the problem, a trouble ticket will be created and sent off to the network administration team to do more advanced troubleshooting until the problem is resolved.

There are several trouble ticket reporting systems on the market today. One popular ticket reporting system is called Remedy. Here is a screenshot of what their reporting system looks like:

The screenshot displays the Remedy IT Service Management for the Enterprise interface. At the top, there is a navigation bar with 'Current mode: Modify', 'Save', 'New search', 'New request', 'Modify all', 'Advanced search', 'Clear', 'Set to defaults', 'Status history', and 'Home'. Below this is a table showing one case:

Case ID+	Priority*	Status*	Category*	Type*	Item*	Summary*	Individual+
HD000000011155	High	Assigned	Default	Default	Default	Memory Utilization is 22.86... Demo	

Below the table are buttons for 'Report', 'Select all', 'Deselect all', and 'Delete'. The main header reads 'Remedy IT Service Management for the Enterprise'. The 'Help Desk Case' form includes the following fields:

- Summary***: Memory Utilization is 22.86%, crossed warning (10) or critical (70) threshold. Sev
- Status***: Assigned
- Description***: (Empty)
- Category***: Default
- Case ID+**: HD000000011155
- Type***: Default
- Case Type***: Incident
- Item***: Default
- Pending**: (Empty)
- Closure Code**: (Empty)
- Escalated?**: Yes

The 'Requester Information' section includes:

- Login****: Demo
- Name****: Demo
- Phone**: (Empty)
- Region**: (Empty)
- Site**: (Empty)
- Department**: (Empty)
- Office**: (Empty)
- Source***: NMP
- Urgency**: High

At the bottom, there are buttons for 'Save', 'Print Case', 'Bulletin Board', 'Reminders', and 'Close', along with a 'Help' link.

Figure 4: The Remedy Ticket Reporting System

Wiki

A web-based Wiki is a fairly new tool that is becoming popular in organizations for use as a central repository for network documentation. The advantage of using Wiki technology for documentation is that it is simple to access and search as well as helps to better get collaboration from multiple network administrators.

Added-Value Network Monitoring Tools

Sometimes basic IOS tools are not quite enough to help maintain and troubleshoot a network. There are several tools that add additional value by providing more robust and precise data that can help to trend and identify traffic problems that would otherwise be difficult to identify. We'll look at three of the most popular tools which are SNMP servers, NetFlow Servers and packet tracers.

SNMP Server

An SNMP server collects and stores SNMP data from Cisco (and other vendor) devices. On the Cisco side, you must configure an SNMP agent to run and point it to a SNMP server located on the network. The SNMP server then collects the sent network data such as utilization statistics or device configuration information. The server can then begin baselining normal network traffic so when a problem occurs, it can be identified by clues such as a spike in network traffic, increased CPU or memory utilization or increased interface errors. Here are a few of the benefits to monitoring a network using an SNMP Server:

- Baselining resource utilization on your network, which can help you recognize trends and determine when upgrades will be required.
- Troubleshooting performance issues.
- Verifying compliance with a service level agreement (SLA).

There are several SNMP tools available for free and for a fee. One popular free SNMP tool is called MRTG. Here is a screenshot of MRTG monitoring interface traffic utilization of LAN and WAN interfaces:

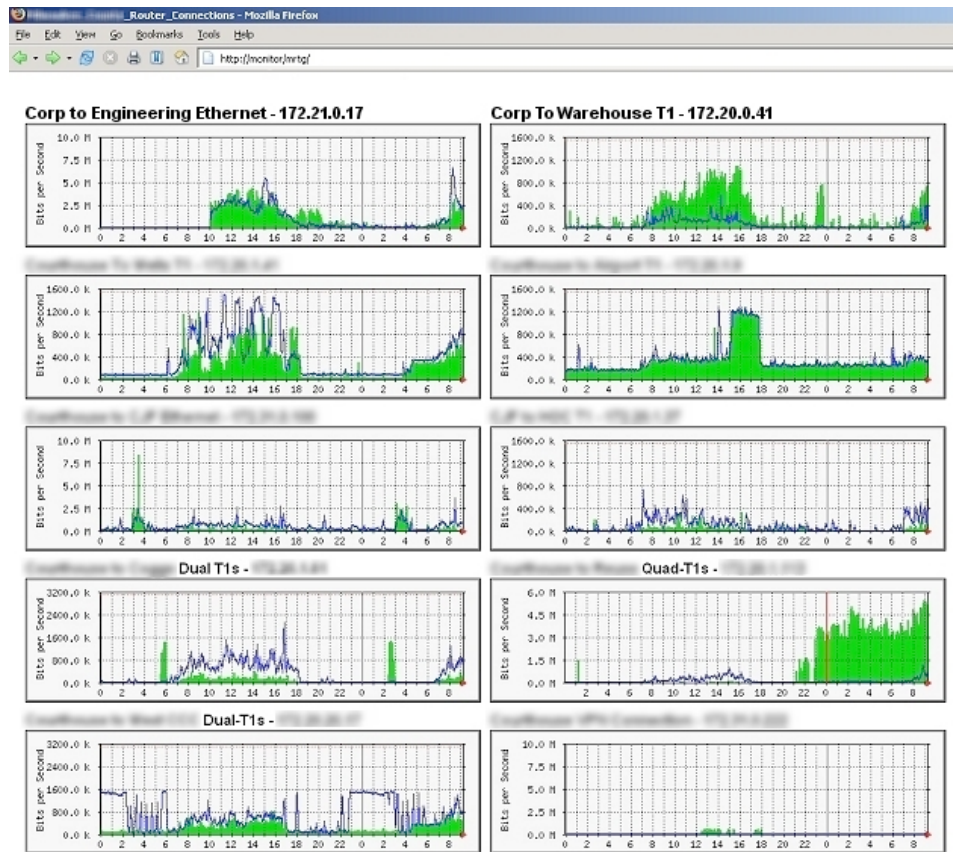


Figure 5: MRTG Monitoring Interface Traffic

A fairly new feature called **Embedded Event Manager (EEM)** works with SNMP to trigger customized actions based on events discovered by SNMP. EEM provides real-time network event detection and automation. The tool gives an administrator the ability to configure EEM to be triggered when criteria are met and adapt to the event without any human interaction. Events can be triggered by any SNMP event and actions can include SNMP traps, syslog messages, automated IOS commands, and email messages. EEM is used to either notify administrators of a problem or automatically attempt to correct the problem on the fly. The tool is basically a script editor. Here is a very basic example of a switch that monitors a specific the switch memory SNMP OID. If that OID counter goes over the configured threshold (5,120,000), a syslog message will be generated with a syslog level of 2 (critical). The EEM is checked every 30 seconds on the switch:

```
Switch(config)#event manager applet monitor_mem
Switch(config-applet)#event snmp oid 1.3.6.1.4.1.9.9.48.1.1.1.6.1 get-
type exact entry-op lt entry-val 5120000 poll-interval 30
Switch(config-applet)#action 1.0 syslog priority critical msg "Memory
threshold reached. Available memory is $_snmp_oid_val bytes"
Switch(config-applet)#end
Switch#
```

NetFlow Server

NetFlow is a network protocol developed by Cisco Systems to run on Cisco IOS-enabled equipment for collecting IP traffic flow information. An IP flow is defined as a unidirectional sequence of packets from a source IP to a destination IP. The following information is collected:

1. Source IP address
2. Destination IP address
3. Source port for UDP or TCP, 0 for other protocols
4. Destination port for UDP or TCP, type and code for ICMP, or 0 for other protocols
5. IP protocol
6. Ingress interface (SNMP ifIndex)
7. IP Type of Service

Here is a basic example on how to configure a Cisco router to send NetFlow version 5 information to a NetFlow server located at 192.168.17.101:

```
Router(config-if)#ip flow ingress
Router(config)#ip flow-export version 5
Router(config)#ip flow-export destination 192.168.17.101
```

An open-source NetFlow product called ntop collects and puts intelligence behind NetFlow data to troubleshoot networks. Here's a screenshot of what collected NetFlow data looks like:




Figure 6: NetFlow NTOP Data Monitoring

Notice how the graphs are broken out by TCP/UDP port. Only NetFlow dives deep into the packet to pull out detailed information like this and this is how it differentiates itself from SNMP monitoring.

Packet Tracers

Packet tracers (also called “sniffers”) enable an administrator to closely examine IP packets on a point of the network the administrator chooses. This information can be used to look at underlying protocol traffic and errors. A packet tracer is best for troubleshooting a single problem occurring on a network. This is because the tracer can examine only the traffic that passes through a specific port or trunk on the network. Packet tracers are often laptops or other portable devices with special packet tracer software installed on them. The laptops are plugged into an Ethernet switchport and the administrator uses Cisco SPAN (Switched Port Analyzer) commands to mirror traffic from one port to the port where the packet tracer is connected. To show you how this is done, let's assume that we have a problem with a server that is connected to port fa0/12 on a switch. We connect our laptop with the packet tracer software to port fa0/22. The following Cisco IOS switch commands can be used to mirror data coming and going on port 12 to port 22:

```
Switch(config)#Monitor session 1 source interface f0/12
Switch(config)#Monitor session 1 destination interface f0/22
```

The packet tracer will then begin receiving information on all the packets that go into and out of fa0/12. Here is a screenshot of a common packet tracer called Wireshark:

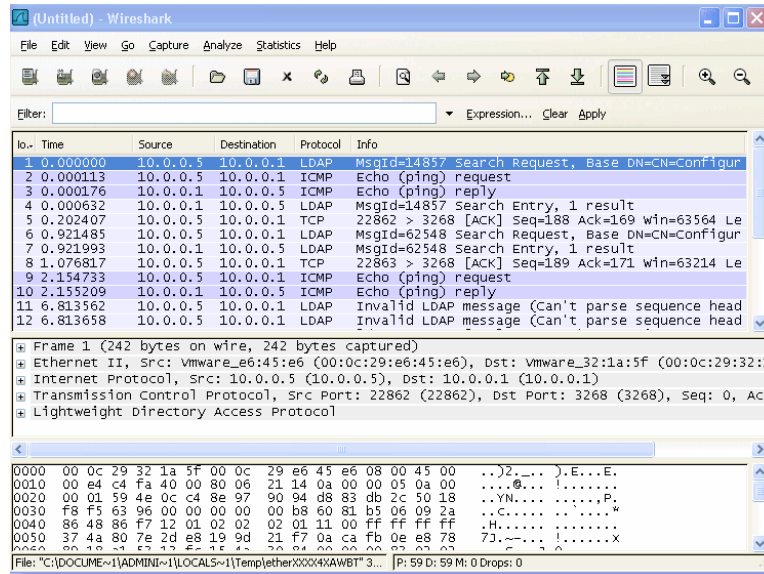


Figure 7: A Wireshark Trace

An alternative method to SPAN and a dedicated packet tracer is to use a Cisco router to capture interesting traffic using the Router IP Traffic Export (RITE) commands. RITE is configured on a router by defining a traffic export profile and then applying it to the appropriate interface. By default only outgoing traffic is captured on the specified interface. This can be manually changed using the bidirectional command. The following example shows a traffic export file “look_at_me” being applied to fa0/1. Both inbound and outbound data is collected and that information is then sent to a data collection device with the MAC address given:

```
Router(config)#ip traffic-export profile look_at_me
Router(config-rite)#interface FastEthernet1/0
Router(config-rite)#bidirectional
Router(config-rite)#mac-address 0111.2222.3334
```

Troubleshooting Methodology

The process of **Troubleshooting** can be broken down into the following steps:

1. Responding to a discovered problem.
2. Determining the root cause. AKA diagnosis.
3. Fixing the problem in the best method possible.

When a network grows, it becomes more and more complex from a troubleshooting standpoint. To help counter this, it is important that network administrators use tried-and-true troubleshooting methodologies to make it easier to diagnose problems. This section will explore several different techniques that can be used to ease the troubleshooting process.

Step 2 of the troubleshooting process is the most difficult to perform. This step requires a great deal of practice to successfully troubleshoot a wide array of possible issues. But if you follow these diagnosis steps, it can make the root cause discovery a much easier process:

1. **Collect information** – when a trouble ticket comes in, you will almost always have to get additional information. This step requires that the administrator either contact the customer to get the information or possibly the use of network administration tools.
2. **Examine the information** – once the administrator feels that enough information has been collected, they must examine the information to see if they can make correlations based on the findings. Baseline information can be very important in this step.
3. **Eliminate potential causes** – the easiest way to narrow down the possibilities of a failure is to first eliminate things that WOULD NOT cause the reported problem.
4. **Hypothesize underlying cause** – this step is where the administrator uses their skill to come up with a handful of possible causes. These hypotheses can be determined using hard facts or even gut feelings if the administrator is experienced enough.
5. **Verify hypothesis** – the administrator then tests their hypothesis, which will confirm or deny the theory of the root cause.

As stated earlier, if you are a fairly new network administrator, it would be wise to follow each of the five steps in a very structured manner. However, if you are a very experienced administrator and have seen similar problems like the one reported, you can use what's known as the "shoot from the hip" method. This method shortens the amount of time it takes from the time the problem is reported to the time it is resolved. Basically, once the information has been collected, the administrator skips the "Examine the information" and "Eliminate potential causes" steps and begins hypothesizing the underlying cause. The figure below helps to grasp this troubleshooting approach.

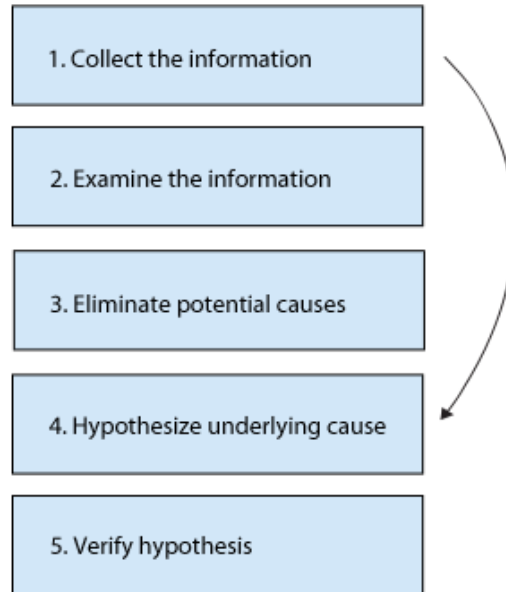


Figure 8: The Standard Troubleshooting Approach

There are additional structured methods to help find the root cause of a network problem. These methods include:

- **Top-Down** – the troubleshooter begins at the top layer of the OSI model (Layer 7). The administrator verifies the application first and then works down the OSI model as shown in this figure:




Figure 9: The Top-Down Approach

- **Bottom-Up** – the troubleshooter begins at the bottom layer of the OSI model (Layer 1). The administrator first checks the physical layer and then works up the OSI model as shown in this figure:




Figure 10: The Bottom-Up Approach

- **Divide and Conquer** – the troubleshooter begins in the middle of the OSI model. It often times is a more efficient method compared to the top-down and bottom-up approaches. Troubleshooting typically begins at Layer 3 with this approach. The administrator will often use a ping to verify IP connectivity. If the ping is successful, the administrator knows that layers 1 to 3 are working properly and can then begin troubleshooting upper-layer protocols. If the ping fails, the administrator can begin troubleshooting at this layer and work down the OSI model as shown in this figure:

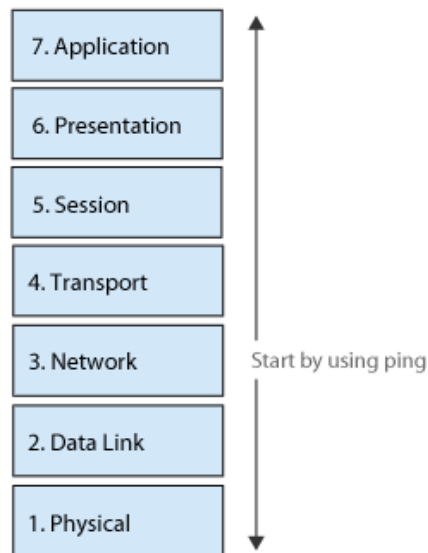


Figure 11: The Divide and Conquer Approach

- **Following Traffic Flows** – Method where the administrator knows the source and destinations of where the traffic is flowing. With this information, troubleshooting can begin by checking the links on the source and destination devices. If those check out, the next hop between the source and destination is checked until the problem is found along the path. A great troubleshooting tool for this approach is the traceroute command.
- **Comparing Configurations** – If you have network equipment that has very similar setups such as remote site routers, a good troubleshooting method to resolve a problem where one site functions properly while the other site does not is to compare the configurations between the working and non-working hardware. This approach is appropriate for new administrators but the problem is that while problems can often be resolved this way, sometimes the root cause is never discovered.
- **Component Swapping** – The final method is to physically change out hardware components (ports, cables, end devices, router/switch, etc) when a device is malfunctioning. This method is good for when a networking device that has been running with no problems and no changes have been made begins to malfunction. If replacing hardware fixes the problem using the exact same configuration, the administrator can assume that the root cause is faulty hardware.

As stated earlier, the majority of time spent troubleshooting occurs in the problem diagnosis phase. This is when the information is examined and hypotheses are made and eliminated. Examining these further, we know that a troubleshooter has two goals when examining information:

- Eliminating potential causes
- Looking for clues that point to the root cause

To meet these goals, two questions must be asked:

- What is happening on the network now?
- How does what is happening now compare to what should be happening?

By asking these questions, the troubleshooter can better begin to think about the true cause to the problem. Baseline data is extremely vital as you can see. Understanding how the network functioned prior to the problem can help to narrow down what component on the network is malfunctioning.

Root Cause Problem Resolution

After a problem has been discovered and resolved, the troubleshooter should document the reported problem, resolution and also document any steps that need to be completed to insure this problem will not occur again. This step is often called the Root Cause Analysis (RCA).

The Integration of Maintenance and Troubleshooting

So now that we've talked about network maintenance and network troubleshooting, how can we use what we know of these two topics to help make administration of a network easier?

When problems are discovered using troubleshooting methods, often times the problem is resolved but it may not be the optimal solution. Many times a band-aid solution is implemented but is not ideal and changes should be made to more permanently fix the problem. This is where routine network maintenance comes into play. Below is a network maintenance lifecycle that should be strictly followed to insure that solutions to problems are completely resolved:

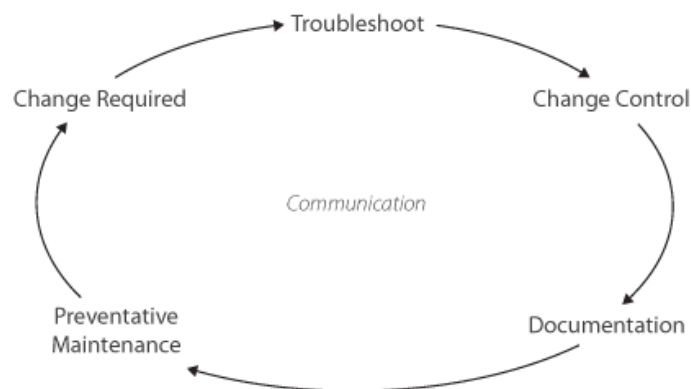


Figure 12: The Network Maintenance Lifecycle

Note that throughout the lifecycle, the central theme is communication. During each phase a network administrator must maintain contact with their customer. Also note that the customer often changes. For example, when troubleshooting the problem, communication often occurs with an end user. During the Change control, documentation and preventative maintenance phases, communication is commonly with other employees within the IT department such as server/applications administrators.

A Closer Look at IOS CLI Troubleshooting Tools

There are many CLI tools in your arsenal for network management and troubleshooting on a network. In this section, we'll break down these tools into the following sections:

- IOS Output Filtering Tools
- Output Redirection
- Troubleshooting Tools
- Hardware Diagnostics

IOS Output Filtering Tools

One problem that network administrators often have when using the CLI is the huge amount of information that they have to sift through in order to get the information they require. Fortunately, there are several IOS output filtering tools that can be used to help narrow down the information presented. This section will cover how to use several of these tools.

The first tool is the pipe (|) command. This command can be used to look for specific pieces of information. For example, let's say an administrator wants to see what processes are running on a router. They run the command as shown here:

```
router#show processes
CPU utilization for five seconds: 0%/0%; one minute: 0%; five minutes: 0%
PID Q  Ty  PC      Runtime(uS)  Invoked    uSecs   Stacks    TTY    Process
1   C  sp  602F3AF0  0             1627       0       2600/3000  0      Load Meter
2   L  we  60C5BE00  4             136        29      5572/6000  0      CEF Scanner
3   L  st  602D90F8  1676          837        2002    5740/6000  0      Check heaps
4   C  we  602D08F8  0             1           0       5568/6000  0      Chunk Manager
5   C  we  602DF0E8  0             1           0       5592/6000  0      Pool Manager
6   M  st  60251E38  0             2           0       5560/6000  0      Timers
7   M  we  600D4940  0             2           0       5568/6000  0      Serial Background
8   M  we  6034B718  0             1           0       2584/3000  0      OIR Handler
9   M  we  603FA3C8  0             1           0       5612/6000  0      IPC Zone Manager
10  M  we  603FA1A0  0            8124        0       5488/6000  0      IPC Periodic Tim
11  M  we  603FA220  0             9           0       4884/6000  0      IPC Seat Manager
12  L  we  60406818  124          2003        61      5300/6000  0      ARP Input
13  M  we  60581638  0             1           0       5760/6000  0      HC Counter Timer
14  M  we  605E3D00  0             2           0       5564/6000  0      DDR Timers
15  M  we  605FC6B8  0             2           0      11568/12000  0      Dialer event
```

Figure 13: Output from the show processes Command

There are many different processes running and our administrator is looking only to see the "Check heaps" process. Alternatively, the administrator could use the pipe command with the "include" statement as shown here:

```
router#show processes | include Check heaps
CPU utilization for five seconds: 0%/0%; one minute: 0%; five minutes: 0%
PID Q   Ty  PC           Runtime(uS)  Invoked uSecs  Stacks      TTY Process
3   L   st  602D90F8    1676         837          2002       5740/6000  0   Check heaps
```

Figure 14: Viewing Specific Processes

Just like the "include" statement, you can use the pipe command to filter using the "exclude" statement. The pipe command can also be used with a "begin" statement to find a specific tag on a show command. The output will find display the line that the tagged text is found and all other lines on the remainder of that show command. For example, a network engineer wants to look at the running-config of a router beginning with the static routes configured on it. To accomplish this, the administrator uses the following pipe show command:

```
show run | begin ip route
```

A third pipe statement TSHOOT candidates should be familiar with is the "section" statement. This is often used when checking the running configuration for a specific section of commands. For example, this command performs a show run and looks for any vty configurations. You can see that everything that falls within the vty statements are slightly indented. This is how you can identify "section" statements:

```
Router#show running-config | section vty
line vty 0 4
  exec-timeout 10 0
  password 7 045C021304288D4901
  logging synchronous
  transport input ssh
line vty 5 15
  exec-timeout 10 0
  password 7 14601BAE01042B2C28
  logging synchronous
  transport input ssh
```

The **show ip route** command has built-in features. This is because on very large routing networks with hundreds or even thousands of routes, sifting through all the information is very difficult. To help administrators looking for specific routes, the command lets you specify specific routes such as in the example here:

```
Router#show ip route 10.119.0.0

Routing entry for 10.119.0.0 (mask 255.255.0.0)
  Known via "igrp 109", distance 100, metric 10989
  Tag 0
  Redistributing via igmp 109
  Last update from 10.108.35.13 on TokenRing0, 0:00:58 ago
  Routing Descriptor Blocks:
  * 10.108.35.13, from 10.108.35.13, 0:00:58 ago, via TokenRing0
    Route metric is 10989, traffic share count is 1
    Total delay is 45130 microseconds, minimum bandwidth is 1544 Kbit
    Reliability 255/255, minimum MTU 1500 bytes

    Loading 2/255, Hops 4
```

An additional ip route show command includes the "longest-prefix" statement. When this statement is included, the address and subnet mask pair becomes the prefix, and any address that matches that prefix is displayed on the screen. The following is an example of this when running a show ip route bgp command:

```
Router#show ip bgp 10.92.0.0 255.255.0.0 longer-prefixes

BGP table version is 1738, local router ID is 192.168.72.24
Status codes: s suppressed, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
*> 10.92.0.0        10.92.72.30      8896           32768 ?
*                   10.92.72.30          0         109     108  ?
*> 10.92.1.0        10.92.72.30      8796           32768 ?
*                   10.92.72.30          0         109     108  ?
*> 10.92.11.0       10.92.72.30     42482           32768 ?
*                   10.92.72.30          0         109     108  ?
*> 10.92.14.0       10.92.72.30      8796           32768 ?
*                   10.92.72.30          0         109     108  ?
*> 10.92.15.0       10.92.72.30      8696           32768 ?
*                   10.92.72.30          0         109     108  ?
*> 10.92.16.0       10.92.72.30     1400           32768 ?
*                   10.92.72.30          0         109     108  ?
*> 10.92.17.0       10.92.72.30     1400           32768 ?
*                   10.92.72.30          0         109     108  ?
*> 10.92.18.0       10.92.72.30     8876           32768 ?
*                   10.92.72.30          0         109     108  ?
*> 10.92.19.0       10.92.72.30     8876           32768 ?
*                   10.92.72.30          0         109     108  ?
```

As you can see, only routes beginning with 10.92.0.0 having a longer subnet mask than a /16 is displayed. Keep in mind that this command can also be used with a "shorter-prefix" statement.

Output Redirection

If you've ever worked a troubleshooting case with Cisco TAC, you'll know that one of the first pieces of information they request is the output of a show tech-support command. You could run this command and copy and paste it into a text file to email it to them. Alternatively, you can automatically create a text file that is redirected to a tftp server as shown here:

```
Router11#show tech-support | redirect tftp://Router11_show-tech.txt
Translating "tftp"...domain server (10.10.100.30) [OK]
```

An administrator wants to not only write the command to a .txt file on a TFTP server, but they also want to view it themselves. To accomplish this, they can use the "tee" statement as opposed to the "redirect" statement, which both copies the output to a file on the TFTP server and also display it on your screen as shown here:

```
Router11#show tech-support | tee tftp://Router11_show-tech.txt
Translating "tftp"...domain server (10.10.100.30) [OK]
!
Building configuration...
Current configuration : 29295 bytes
version 12.0
no service udp-small-servers
no service tcp-small-servers
hostname Router11
<output omitted>
```

Troubleshooting Tools

There are several Cisco tools built-in to the IOS to assist with a multitude of troubleshooting tasks. We'll look at three of the most often used tools:

- Ping
- Traceroute
- Telnet

The **ping** command is great for troubleshooting layer 3 problems. The simplest way to use it is to issue the command and the IP address of a problematic device to see if it responds, as shown here:

```
Router#ping 192.168.99.99
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.99.99, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 8/9/12 ms
Router#
```

You can see that by default, the ping command sends 5 pings to the device with a datagram size of 100 bytes and a timeout of 2 seconds. If you want to change the defaults, you simply issue the ping command and you can then make modifications to the default settings. Here's an example where we use an **extended ping** that runs a 100 **ping sweep** between 100 and 1500 byte datagram sizes. In addition, we changed the timeout from 2 to 4 seconds and set the **don't fragment (DF) bit**:

```
Router#ping
Protocol [ip]:
Target IP address: 192.168.99.99
Repeat count [5]:
Datagram size [100]:
Timeout in seconds [2]: 4
Extended commands [n]: y
Source address or interface:
Type of service [0]:
Set DF bit in IP header? [no]: y
Validate reply data? [no]:
Data pattern [0xABCD]:
Loose, Strict, Record, Timestamp, Verbose [none]:
Sweep range of sizes [n]: y
Sweep min size [36]: 100
Sweep max size [18024]: 1500
Sweep interval [1]: 100
Type escape sequence to abort.
Sending 75, [100..1500]-byte ICMP Echos to 192.168.99.99, timeout is 4
seconds:
Packet sent with a source address of 192.168.1.1 Packet sent with the
DF bit set !!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!
Success rate is 100 percent (75/75), round-trip min/avg/max = 8/10/12 ms
```

Make sure you understand that extended ping command shows the default ping settings in the brackets [.]. Also, you can see that the exclamation mark (!) in the output means that the ping was successful. There are several other characters you might run across if the ping fails for some reason. The table here lists the characters and their meanings:

Ping Character	Description
!	Successful ping reply
.	Time-out while waiting for reply
U	Destination unreachable
Q	Destination too busy (quench)
M	Could not fragment
?	Unknown packet type
&	Packet lifetime exceeded

Figure 15: Ping Characters

The **tracert** command is used to check the layer 3 path between a source and destination. In the output, you can see the round-trip time (RTT) to a particular device along the path. Also remember that the Cisco **tracert** command uses **UDP** for transport. Just like the ping command, if you simply type **tracert** and hit enter, you can modify the default settings. Here's an example checking UDP port 2000 (what Cisco SCCP IPT phones use):

```
Router#tracert
Protocol [ip]:
Target IP address: 192.168.30.199
Source address:
Numeric display [n]:
Timeout in seconds [3]:
Probe count [3]:
Minimum Time to Live [1]:
Maximum Time to Live [30]:
Port Number [33434]: 2000
Loose, Strict, Record, Timestamp, Verbose [none]:
Type escape sequence to abort.
Tracing the route to 192.168.30.199
 0 10.0.0.2 4 msec 4 msec 4 msec
 1 172.16.0.3 20 msec 16 msec 16 msec
 2 192.168.30.199 msec * 16 msec
```

The last IOS troubleshooting tool we'll look at is the **telnet** command. You might think that telnet is more of a method of remotely connecting to other devices such as routers and switches. While, you may be right, experienced network administrators also use the telnet command to troubleshoot problems at the Session (layer 4) layer of the OSI model.

By default, when you issue the telnet command and a destination IP address, the default port it uses is TCP 23. One option that can be changed is the TCP port that is sent to open a session to the destination IP. You can change the destination port to check to see if devices have specific TCP ports open or closed. This gives the administrator an indication if an application is properly running on a server. For example, if you are having problems with an HTTP server, you can send a telnet request. If the response comes back as "open" then at least you know the server is running a web server (running at DNS www.test-webserver12.com). Here's an example of using this command for troubleshooting:

```
Router#telnet test-webserver12.com 80
Translating "test-webserver12.com"...domain server (10.100.10.4) [OK]
Trying test-webserver12.com (172.16.100.25, 25)... Open
```

Hardware Diagnostics

Many times, administrators suspect faulty hardware as the root cause of a networking problem. One sign is where a major outage occurs when no changes to the network and server/client processes were made. If the Cisco hardware is operational but possibly faulty, there are several hardware diagnostic show commands in the IOS that can be used to diagnose a problem. These tools include:

- show interfaces
- show inventory
- show environment
- show module
- show memory
- show platform

The **show interface** command is widely known and lets a network administrator view port statistics of any interface the hardware has installed. This example narrows the output to display only Serial interface 0 on the router:

```
Router# show interface serial 0

Serial 0 is up, line protocol is up
  Hardware is MCI Serial
  Internet address is 192.168.10.203, subnet mask is 255.255.255.0
  MTU 1500 bytes, BW 1544 Kbit, DLY 20000 usec, rely 255/255, load
  1/255
  Encapsulation HDLC, loopback not set, keepalive set (10 sec)
  Last input 0:00:07, output 0:00:00, output hang never
  Output queue 0/40, 0 drops; input queue 0/75, 0 drops
  Five minute input rate 0 bits/sec, 0 packets/sec
  Five minute output rate 0 bits/sec, 0 packets/sec
    16263 packets input, 1347238 bytes, 0 no buffer
    Received 13983 broadcasts, 0 runts, 0 giants
    2 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 2 abort
  1 carrier transitions

    22146 packets output, 2383680 bytes, 0 underruns
    0 output errors, 0 collisions, 2 interface resets, 0 restarts
```

Here you can check the input/output rates, number of packets transmitted and received and any errors that have occurred on the interface. These statistics can be reset by issuing the following command:

```
Router#clear counter interface serial 0
```


The **show inventory** command displays hardware information such as installed modules and serial numbers. This information can be useful to compare one device to another to detect differences in equipment as well as being very useful for Cisco TAC cases. Here's an example of this command:

```
Switch#show inventory
NAME: "Chassis", DESCR: "Cisco Systems WS-C6509 9 slot switch"
PID: WS-C6509          , VID:          , SN: SCA034401LQ

NAME: "Clock 1", DESCR: "Clock"
PID: WS-C6000-CL      , VID:          , SN: SMT03462479

NAME: "Clock 2", DESCR: "Clock"
PID: WS-C6000-CL      , VID:          , SN: SMT03462480

NAME: "VTT 1", DESCR: "VTT"
PID: WS-C6000-VTT     , VID:          , SN: SMT03460976

NAME: "VTT 2", DESCR: "VTT"
PID: WS-C6000-VTT     , VID:          , SN: SMT03460843

NAME: "VTT 3", DESCR: "VTT"
PID: WS-C6000-VTT     , VID:          , SN: SMT03461008

NAME: "2", DESCR: "1000BaseX Supervisor 2 port WS-X6K-SUP2-2GE Rev.
1.1"
PID: WS-X6K-SUP2-2GE  , VID:          , SN: SAD04450LF1

NAME: "submodule 2/1", DESCR: "L3 Switching Engine II"
PID: WS-F6K-PFC2      , VID:          , SN: SAD04440HVU

NAME: "3", DESCR: "10/100BaseTX Ethernet 48 port WS-X6248-RJ-45 Rev.
1.0"
PID: WS-X6248-RJ-45   , VID:          , SN: SAD03181468

NAME: "5", DESCR: "Switch Fabric Module 0 port WS-C6500-SFM Rev. 1.0"
PID: WS-C6500-SFM     , VID:          , SN: SAD04420JR5

NAME: "7", DESCR: "Network Analysis Module 2 port WS-X6380-NAM Rev.
0.201"
PID: WS-X6380-NAM     , VID:          , SN: JAB0343055Y

NAME: "8", DESCR: "1000BaseX Ethernet 8 port WS-X6408-GBIC Rev. 0.202"
PID: WS-X6408-GBIC    , VID:          , SN: SAD02430406

NAME: "PS 1", DESCR: "1300 watt supply AC"
PID: WS-CAC-1300W     , VID:          , SN: ACP03380477

NAME: "Fan 1", DESCR: "Fan 1"
PID: WS-C6K-9SLOT-FAN , VID:          , SN:
```

The **show environment** command displays system status information of the hardware including power-supplies, fans, EEPROM, clocking and VTT. These are shown in the following conditions:

- Pass
- Fail
- Unknown
- Not Present

Here's an example of this command output:

```
Switch#show environment
Environmental Status (. = Pass, F = Fail, U = Unknown, N = Not Present)
PS1:.      PS2:N      PS1 Fan:..  PS2 Fan:N
Chassis-Ser-EEPROM:..  Fan:..  Clock(A/B):A      Clock A:..  Clock B:..
VTT1:..   VTT2:..   VTT3:..
Switch#
```

The **show module** command is used to display module status and information on modular Cisco hardware such as the Catalyst 4500 and 6500 series switches. If a problem is detected on a specific hardware module, the status will be something other than "ok". These states include:

- Ok
- Disable
- Faulty
- Other
- Standby
- Error
- pwr-down
- pwr-deny

For supervisor engines, the command displays the supervisor engine number and appends the uplink daughter card's module type and information. These are shown in the "Sub-Type" and "Sub-Model" section shown here:

```
Switch#show module
Mod Slot Ports Module-Type           Model           Sub Status
-----
1   1   2   1000BaseX Supervisor   WS-X6K-SUP1A-2GE  yes ok
15  1   1   Multilayer Switch Feature WS-F6K-MSFC       no ok
8   8   48  10/100BaseTX Ethernet   WS-X6248-RJ-45   no ok
9   9   48  10/100BaseTX Ethernet   WS-X6348-RJ-45   yes ok

Mod Module-Name      Serial-Num
-----
1                    SAD03436055
15                   SAD03432597
9                    SAD03414268
```

```

Mod  MAC-Address (es)                               Hw      Fw      Sw
---  -
1    00-30-80-f7-a5-06 to 00-30-80-f7-a5-07  1.0     5.2(1)  6.1(0.12)
      00-30-80-f7-a5-04 to 00-30-80-f7-a5-05
      00-30-a3-4a-a0-00 to 00-30-a3-4a-a3-ff
15   00-d0-bc-ee-d0-dc to 00-d0-bc-ee-d1-1b  1.2     12.0(3)XE1 12.0(3)XE1
8    00-d0-c0-c8-83-ac to 00-d0-c0-c8-83-db  1.1     4.2(0.24)V 6.1(0.37)
FTL
9    00-50-3e-7c-43-00 to 00-50-3e-7c-43-2f  0.201   5.3(1)

Mod  Sub-Type                Sub-Model          Sub-Serial  Sub-Hw
---  -
1    L3 Switching Engine       WS-F6K-PFC        SAD03451187 1.0
9    Inline Power Module      WS-F6K-VPWR              1.0
Switch#

```

The **show memory** command displays a summary of memory usage on an IOS device including the size and number of blocks allocated for each address of the system call that allocated the block as shown here:

```

Router#show memory
Head      Total (b)  Used (b)  Free (b)  Lowest (b)  Largest (b)  Processor
60AB4ED0  5550384   2082996  3467388  3464996    3454608     I/O
40000000  16777216  1937280  14839936  14839936   14838908

Processor memory
Address  Bytes Prev.  Next  Ref  PrevF  NextF  Alloc PC  What
B0EE38  1056  0      B0F280  1    18F132  List Elements
B0F280  2656  B0EE38  B0FD08  1    18F132  List Headers
B0FD08  2520  B0F280  B10708  1    141384  TTY data
B10708  2000  B0FD08  B10F00  1    14353C  TTY Input Buf
B10F00  512   B10708  B11128  1    14356C  TTY Output Buf
B11128  2000  B10F00  B11920  1    1A110E  Interrupt Stack
B11920  44    B11128  B11974  1    970DE8  *Init*
B11974  1056  B11920  B11DBC  1    18F132  messages
B11DBC  84    B11974  B11E38  1    19ABCE  Watched Boolean
B11E38  84    B11DBC  B11EB4  1    19ABCE  Watched Boolean
B11EB4  84    B11E38  B11F30  1    19ABCE  Watched Boolean
B11F30  84    B11EB4  B11FAC  1    19ABCE  Watched Boolean
Router#

```

Figure 16: Output from the <show memory command> Command

The following table explains what each of these memory fields mean:

Field	Description
Total	Total amount of memory available after the system image loads and builds its data structures
Used	Amount of memory currently in use
Free	Amount of memory currently available for use
Lowest	Lowest amount of free memory recorded by the router since it was last booted
Largest	Largest free memory block currently available

Figure 17: Memory Fields

Domain 2: Troubleshoot Multi-Protocol System Networks

Potential Layer 2 Switch Problems

When a network administrator is troubleshooting a problem and has it narrowed down to the data-link (layer 2) layer of the OSI model or lower, there are three potential problems that they should investigate:

Hardware Problems

Hardware problems are considered to be a physical (layer 1) layer issue. Common things to check are the endpoint hardware network card, cabling and the Cisco switchport. Some ways to eliminate a physical layer problem is to swap out cabling and replace end-devices. You can also change switchports to a known working interface to verify it's not a problem.

VLAN Configuration Problems

A virtual LAN (VLAN) is a logical network that allows a group of devices to act as if they are on the same physical network. All devices within a VLAN share the same unicast, broadcast and multicast domain. If a device on one VLAN needs to communicate to a device on a different VLAN, that traffic must be routed. If you are troubleshooting a problem such as this, you will have to verify that the devices are either in the same VLAN or are properly routed.

Trunk Configuration Problems

When connecting multiple switches together, a trunk link can be used to transport traffic from multiple VLANs across a single point-to-point link. This is called a trunk. If you are troubleshooting a trunk link, you will need to verify that the configuration settings are the same on both switches. Settings like trunk encapsulation type (802.1q or ISL), native VLAN and trunk modes are configured for proper operation.

Layer 2 Troubleshooting Techniques

Network administrators have multiple show command options when troubleshooting physical, VLAN and trunk problems. The next section goes through some of the more important commands and instructs you when they should be used. The commands that you should familiarize yourself with include:

- show interfaces
- show interfaces counters
- show interfaces errors
- show mac-address-table (or "show mac address" depending on the IOS version/hardware being used)
- show mac-address-table <address>
- clear mac-address-table dynamic
- show vlan brief
- show interfaces switchport
- show interfaces trunk
- traceroute mac

Show interfaces

A network administrator can discover problems at the physical layer by using the show interfaces command. The following table lists some of these error counters and why they might be incrementing:

Error Counter	Common Problem
Align-Err	Cabling problems or duplex mismatch
FCS-Err	Cabling problems
Xmit-Err	Bottleneck on the local switch
Undersize Giants	Bad end-station network card
Single-Col	Duplex mismatch
Multi-Col	Duplex mismatch
Late-Col	Duplex mismatch
Excess-Col	Duplex mismatch

Figure 18: Error Counters

Below is an example of where these counters can be found when running the show interfaces command:

```
Switch#show interfaces fa0/10
FastEthernet0/10 is up, line protocol is up (connected)
  Hardware is Fast Ethernet, address is 000b.4660.840a (bia
000b.4660.840a)
  Description: SERVER-PANEL-1
  MTU 1500 bytes, BW 100000 Kbit, DLY 100 usec,
    reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  Keepalive set (10 sec)
  Full-duplex, 100Mb/s, media type is 100BaseTX
  input flow-control is off, output flow-control is unsupported
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input never, output 00:00:00, output hang never
  Last clearing of "show interface" counters never
  Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
  Queueing strategy: fifo
  Output queue: 0/40 (size/max)
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 1000 bits/sec, 3 packets/sec
    25694403 packets input, 274691981 bytes, 0 no buffer
    Received 34048 broadcasts (0 multicast)
    0 runts, 0 giants, 0 throttles
```

```

0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
0 watchdog, 441 multicast, 0 pause input
0 input packets with dribble condition detected
102642183 packets output, 4098595110 bytes, 0 underruns
0 output errors, 0 collisions, 1 interface resets
0 babbles, 0 late collision, 0 deferred
0 lost carrier, 0 no carrier, 0 PAUSE output
0 output buffer failures, 0 output buffers swapped out

```

Show interfaces counters

A different method of viewing interface stats is the `show interfaces <interface> counters` command. This command is useful for viewing traffic that is broken into broadcast, unicast and multicast counters. Here is an example of the command output:

```
Switch#show interfaces fa0/10 counters
```

Port	InOctets	InUcastPkts	InMcastPkts	InBcastPkts
Fa0/10	4569662887	25660389	441	33608

Port	OutOctets	OutUcastPkts	OutMcastPkts	OutBcastPkts
Fa0/10	51343291745	62123059	22030198	18489748

Show interfaces errors

Adding the keyword "errors" to the show interface counters command, displays detailed descriptions of the types of errors that may be occurring on your interface as shown here, where there are incrementing Transmit errors (Xmit-Err), which is likely due to a bottleneck when too much traffic from a larger bandwidth port (such as Gigabit Ethernet) is trying to send traffic to a lower bandwidth port and overloading it:

```
Switch#show interfaces gi0/1 counters errors
```

Port	Align-Err	FCS-Err	Xmit-Err	Rcv-Err	UnderSize
gi0/1	0	154423	0	0	0

Port	Single-Col	Multi-Col	Late-Col	Excess-Col	Carri-Sen	Runts	Giants
gi0/1	0	0	0	0	0	0	0

Show mac-address-table

The content addressable memory (CAM) table holds the port to MAC address mappings on a switch. This is useful when troubleshooting layer 2 problems or when tracking bad network card that is malfunctioning. There are basically two different types of CAM table entries. Static mappings are manually configured on the switch and dynamic mappings are learned by the switch when a network card begins transmitting data. By default, dynamic mappings are stored in the CAM table for 4 hours. If no traffic is seen by the switch after the timer expires, the entry is removed from the table.

Here is an example of how to view the entire CAM table using the show mac-address-table command:

```
Switch#show mac-address-table
      Mac Address Table
-----
Vlan    Mac Address      Type      Ports
----    -
All     000b.43d0.cb30   STATIC    CPU
All     000b.46d1.cb21   STATIC    CPU
30      0000.0c07.a20d   DYNAMIC   Gi0/1
30      0001.2900.27ea   DYNAMIC   Gi0/1
30      0001.e698.98fb   DYNAMIC   Fa0/34
30      0001.e6a0.bb9a   DYNAMIC   Fa0/39
```

Show mac-address-table address

If a network administrator is looking for a specific MAC address on a switch they can add the keyword "address" at the end of the command and enter the specific MAC as shown here:

```
Router# show mac-address-table address 001.6441.60ca
Codes: * - primary entry

  vlan  mac address      type  learn qos      ports
-----+-----+-----+-----+-----+-----
Supervisor:
* --- 0001.6441.60ca   static No  -- Router
```

Clear mac-address-table dynamic

An administrator will come across times when they want to quickly purge dynamic MAC address from the CAM table. This is just a simple matter of entering the command clear mac-address-table dynamic. This will clear out all of the dynamic mappings on that particular switch. The table will then be rebuilt when the connected devices begin attempting to communicate through the switch port.

Show vlan brief

When an administrator is troubleshooting switch problems that might be VLAN related, the most obvious piece of information that should be obtained on the switch is what VLANs are actually configured on it. Using the `show vlan brief` command displays all of the configured VLANs with their names, the current status and what ports are attached to that VLAN as seen in this example output:

```
Router# show vlan brief
VLAN Name                Status    Ports
-----
1    default                 active    Fa0/1, Fa0/2
2    VLAN0002                active    Fa0/3, Fa0/4
3    VLAN0003                active    Fa0/5, Fa0/6
4    VLAN0004                active    Fa0/7, Fa0/8
5    VLAN0005                active    Fa0/9
10   VLAN0010                active    Fa0/10
.
.
.
999  VLAN0999                active    Fa0/11, Fa0/12
1002 fddi-default            active
1003 trcrf-default        active
1004 fddinet-default       active
1005 trbrf-default        active
Router#
```

Show interfaces switchport

If you want to get detail into both the administrative and operational status of a switchport, use the `show interfaces switchport` command. This command is very useful to provide quick insight into the administration setup of trunk ports that carry multiple VLANs to other switches. Looking at this example here, we can quickly see the configured trunking encapsulation, negotiation, and encapsulation types. Also listed are the VLANs that are allowed to propagate the trunk port:

```
Switch#show interfaces switchport
Name: Fa0/1
Switchport: Enabled
Administrative Mode: static access
Operational Mode: down
Administrative Trunking Encapsulation: negotiate
Negotiation of Trunking: Off
Access Mode VLAN: 1 (default)
Trunking Native Mode VLAN: 1 (default)
Voice VLAN: none
Administrative private-vlan host-association: none
Administrative private-vlan mapping: none
Administrative private-vlan trunk native VLAN: none
Administrative private-vlan trunk encapsulation: dot1q
Administrative private-vlan trunk normal VLANs: none
Administrative private-vlan trunk private VLANs: none
```



```

Operational private-vlan: none
Trunking VLANs Enabled: ALL
Pruning VLANs Enabled: 2-1001
Capture Mode Disabled
Capture VLANs Allowed: ALL
Protected: false
Unknown unicast blocked: disabled
Unknown multicast blocked: disabled
Appliance trust: none
<output omitted>

```

Show interfaces trunk

If the network engineer logs into a switch and wants to quickly view what port(s) are configured as trunk ports, the show interfaces trunk command is a great tool. This command displays VLANs that are allowed on specific trunk ports and which VLANs are being forwarded as seen in this output:

```

Switch#show interfaces trunk

Port      Mode           Encapsulation  Status        Native vlan
Fa0/1     desirable     802.1q         trunking     1

Port      Vlans allowed on trunk
Fa0/1     1-1005

Port      Vlans allowed and active in management domain
Fa0/1     1-6,10,20,50,100,152,200,300,303-305,349-
351,400,500,521,524,570,801-8
02,850,917,999,1002-1005

Port      Vlans in spanning tree forwarding state and not pruned
Fa0/1     1-6,10,20,50,100,152,200,300,303-305,349-
351,400,500,521,524,570,801-8
02,850,917,999,1002-1005

```

Traceroute mac

We've already covered the traceroute command at layer 3. You can also use this command on a Cisco switch to perform a layer 2 trace. The traceroute mac <source-MAC> <destination-MAC> can be used to see the layer 2 hop-by-hop path a datagram takes. The hops are displayed as physical switchports as seen here:

```

Switch# traceroute mac 0000.0201.0601 0000.0201.0201

Source 0000.0201.0601 found on con6 [WS-C3750-12T] (2.2.6.6)
con6 (2.2.6.6) :Gi0/0/1 => Gi0/0/3
con5          (2.2.5.5       ) :    Gi0/0/3 => Gi0/0/1
con1          (2.2.1.1       ) :    Gi0/0/1 => Gi0/0/2
con2          (2.2.2.2       ) :    Gi0/0/2 => Gi0/0/1
Destination 0000.0201.0201 found on con2 [WS-C3550-24] (2.2.2.2)
Layer 2 trace completed

```

Spanning Tree Troubleshooting Techniques

Before we cover some different troubleshooting commands, let's review the major points of STP that a TSHOOT candidate should fully understand.

The Spanning Tree Protocol (STP) is an IEEE standard based on the 802.1D algorithm. The purpose of STP is to prevent layer 2 loops on a network. It uses bridge protocol data units (BPDU) to talk to other switches to remove loops by putting them into a blocking state. The exchanges of BPDUs also can turn blocked ports back into speaking ports when a link failure is detected. This allows engineers to configure redundant connections between switches while guaranteeing that a loop will not be formed.

Every port on a switch has a unique identifier, which is its MAC address. When two switches are connected, the STP algorithm assigns a default path cost to the link based on the type and speed of the link. This path cost can be manipulated as needed.

The Spanning Tree Algorithm (STA) is recalculated each time a network topology change occurs such as an adding or removal of a new link. This is determined when a switch port begins or stops receiving BPDUs. BPDUs contain a wealth of information that assist the switch in figuring out the optimal path to forward and which path to block. The default STP information that is contained in the BPDU is as follows:

Setting:	Default:
Switch Priority	32768
Port Cost	1000 Mbps: 4
	100 Mbps: 19
	10 Mbps: 100
Port Priority	128
Hello interval	2 seconds
Forward delay	15 seconds
Max age	20 seconds

Figure 19: BPDU Information

The switch priority plays a determining role in which switch becomes the root bridge. The switch with the lowest priority becomes the root. The other factor that determines the root bridge is the port MAC address. The lower MAC address becomes the root if all the switch priority is the same on two or more switches. The path costs and port priority help to compute which links are placed into forwarding or blocking state. The hello interval is the time between BPDUs being sent. The forward delay timer determines the amount of time the listening and learning states last prior to being moved to a forwarding state. The Max age timer determines the amount of time the switch stores STP information received on a port.

There are several types of STP that can be configured on Cisco devices:

- **Common Spanning Tree (CST)** – this STP protocol provides a single instance for the entire layer 2 network. All VLANs within the switched network share the same instance.
- **Per VLAN Spanning Tree (PVST)** – this protocol maintains a separate STP instance for every VLAN configured on the switched network. One of the advantages of PVST is that it can load balance and separate traffic across trunk links because one instance can block a port on a trunk inside a particular VLAN while the other instance will forward their traffic out the same trunk. This is the default mode on Cisco switches.

- **Multiple Instance Spanning Tree (MISTP)** – this is an IEEE 802.1S protocol that allows engineers to map multiple VLANs to an instance of STP.
- **Rapid Spanning Tree (RSTP)** – RSTP is an IEEE 802.1w standard that is the evolution of STP over time. The protocol implements many new features to allow for faster convergence after a topology change has occurred. The features are the equivalent to PortFast, BackboneFast and UplinkFast which can be configured with regular STP but are Cisco proprietary. RSTP can converge in less than a second compared to up to 50 seconds with standard STP.
- **Per VLAN Rapid Spanning Tree (PVRST)** – a combination of PVST+ and Rapid Spanning Tree that allows for use of the Rapid Spanning Tree features on a per VLAN basis.

As stated before, STP detects potential loops and blocks them by doing the following:

- Each switch advertises Bridge Protocol Data Units (BPDU) that announces its bridge-ID, current root bridge, and path cost to the root to all other connected switches. When the STP process first begins, every switch believes it is the root!
- When a switch receives a BPDU from a neighbor that contains a different root, it compares its own root with the new root suggestion. If the received BPDU has a lower root cost, the switch changes its mind and recalculates the cost to the root bridge. The port that received the superior BPDU is the **root port** because it is the port with the lowest cost to the root bridge. All other ports on the switch are then defined as **designated ports**. Every layer 2 segment has a single designated port, which is based on the lowest cost to the root bridge. That means that all ports on a root bridge switch are designated ports. **Nondesignated ports** are the only ports that block traffic which creates a loop-free layer 2 segment.

Spanning Tree Stages

A network administrator must also understand the five possible STP stages:

- **Disabled** – the disabled state means that the switch port is administratively shutdown.
- **Blocking** – does not forward any frames but still listens to BPDUs from other switches. The default timer is 20 seconds.
- **Listening** – this is the next stage after blocking. In the Listening mode, the port will begin sending BPDUs out the port. The default timer is 15 seconds.
- **Learning** – this stage is where the switch receives MAC address information from connected devices on the switchport. No data traffic is passed at this point of the process. The default timer is 15 seconds.
- **Forwarding** – this stage is where the port actually begins the transmission and reception of data frames.

Now that we've covered what a TSHOOT candidate is required to know for the exam, we'll focus our attention on commands used to help fix any problems with STP.

Show spanning tree vlan

The first command is `show spanning-tree vlan <vlan-id>`. This command displays information about the STP state of the local switch. This example shows how to display spanning tree information for a specific VLAN. Please note that because the Root ID and Bridge ID MAC addresses do not match, we know that this switch is not the root bridge. Also note that port `gi0/1` is currently the root port and therefore forwarding while `gi0/2` is the non-designated port and is blocking:

```
Switch#show spanning-tree vlan 30
VLAN0030
  Spanning tree enabled protocol ieee
  Root ID    Priority    8222
             Address    00d0.02e5.c800
             Cost      3008
             Port      49 (GigabitEthernet0/1)
             Hello Time 2 sec  Max Age 20 sec  Forward Delay 15 sec
  Bridge ID  Priority    49182 (priority 49152 sys-id-ext 30)
             Address    000b.46d0.cb80
             Hello Time 2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time 300
  Uplinkfast enabled
Interface          Role Sts Cost          Prio.Nbr Type
-----
Gi0/1              Root FWD 3004          128.49  P2p
Gi0/2              Altn BLK 3004          128.50  P2p
```

Show spanning tree <interface> detail

The `show spanning-tree <interface> detail` command sheds light into what information is contained within BPDUs. This is the information that this particular switchport is sending out. It also has counters as to the number of BPDUs sent and received:

```
Switch#show spanning-tree interface gi0/1 detail
Port 49 (GigabitEthernet0/1) of VLAN0001 is forwarding
  Port path cost 3004, Port priority 128, Port Identifier 128.49.
  Designated root has priority 8300, address 000b.450d.d000
  Designated bridge has priority 16492, address 000b.450d.cc00
  Designated port id is 128.400, designated path cost 4
  Timers: message age 2, forward delay 0, hold 0
  Number of transitions to forwarding state: 1
  Link type is point-to-point by default
  BPDU: sent 6, received 700
```

Common Spanning Tree Problems

In this section, we'll explore how to identify two of the more common problems a network engineer would see on a switched network.

MAC Address Table Corruption

As we know, the CAM table is responsible for keeping track of MAC address to switchport mappings on a local switch. Sometimes things go wrong and for whatever reason, the CAM table becomes corrupted. What usually happens when the corruption occurs is that two copies of the same Ethernet frame will be sent out onto the network. Depending on the configuration of your network, downstream switches may incorrectly modify their CAM table thinking the source device is located upstream from them. Bottom line is that a MAC address table corruption of one switch can cause catastrophic problems on the entire network.

The key to identifying this problem is by tracking the MAC address from switch to switch, continuing to get closer to the source PC and ultimately finding the problematic switch.

Broadcast Storm

Switches flood broadcasts out all interfaces in the same VLAN, except the interface in which the frame arrived. If your layer 2 segment has a loop, it will continuously forward datagrams from one switch to another in a never-ending cycle. Remember that there is no TTL in an Ethernet frame so it can go on forever! Many times, a broadcast storm is not due to a failure within STP but rather, a configuration error or adding a new switch/hub on the network can cause the LAN segment to loop.

Symptoms of a broadcast storm include very slow network response times and network operations ceasing to function. Eventually, a broadcast storm can bring down a large portion of your network if you are spanning VLANs across distribution blocks.

EtherChannel Troubleshooting Techniques

EtherChannel provides a way to aggregate multiple Ethernet connections into one logical link. It's a simple way to increase bandwidth between your switches where you find bottlenecks. Once a group of links are bonded with EtherChannel, it has the ability to recover when a link failure occurs within the aggregate channel. Up to 8 fast Ethernet or gigabit Ethernet links can be channeled to provide for up to 800 Mbps, 8 Gbps and with new 10 Gigabit interfaces, up to 80 Gbps full duplex bandwidth! All the links in the EtherChannel must be of the same speed and they need to be configured as layer 2 or layer 3 interfaces. Also remember that when running STP across an EtherChannel link, STP treats the bonded links as one physical link.




Figure 20: Bonded EtherChannel Links

EtherChannel is configured in one of three ways on a Cisco switch. EtherChannel On mode statically sets up the links into a channel. **Link Aggregation Control Protocol (LACP)** and **Port Aggregation Control Protocol (PAgP)** negotiates channelized ports with the other side. No matter which method you choose, both sides of the channel link must be configured the same.

Load balancing on EtherChannels happens at the MAC address level. By default, source-address load balancing is configured. What this means is that when a frame enters an EtherChannel, it gets randomly assigned a link to use based on the source MAC address. This works great if you have multiple hosts attached to a switch that is configured for EtherChannel. Each PC or end device has a unique MAC address and load-balancing will occur across the links. When traffic has to go through a layer 3 device that connects to an EtherChannel link, source based MAC address balancing does not work because all traffic will come from the MAC address assigned to the layer 3 interface. In this case, it is better to use destination-based Load balancing which randomly assigns frames to ports within the EtherChannel based on the destination MAC address.

There are several things that a network administrator must verify on a problematic EtherChannel link:

- All ports must be of the same type including
 - ▶ Speed
 - ▶ Duplex
 - ▶ Access port or Trunk
 - ▶ VLAN settings
- Both ends of the channel must be configured for the same EtherChannel type being LACP, PAgP or "on" with no negotiation. If only one side is configured as an EtherChannel, the ports on the remote switch will be placed into Err-disable mode and cannot function.
- The channel might form, but traffic might still be traveling over primarily one link. This is due to the distribution algorithm being used for load balancing. If your traffic is primarily used by a small number of devices, changing the load-balancing algorithm from source-based to destination-based MAC addressing can help eliminate this problem. A great show command to see how well your distribution algorithm is balancing traffic is the show interface counters etherchannel command as shown in this example where we are looking at input and output traffic of the individual links configured as a single EtherChannel interface (po1).

```
sw4#sh interface po1 counters etherchannel
Port          InOctets    InUcastPkts  InMcastPkts  InBcastPkts
Po1           23863       15           229           0
Fa0/19        1899894     835          18864         0
Fa0/20        1903864     835          18862         5
Fa0/21        1945733     835          19307         5
Port          OutOctets    OutUcastPkts OutMcastPkts  OutBcastPkts
Po1           3653        15           14            0
Fa0/19        181124     834           533           5
Fa0/20        171050     834           494           0
Fa0/21        169510     834           469           0
```

Troubleshooting InterVLAN Routing (IVR)

When two devices on different VLANs need to talk to each other, they require assistance from a routing device to communicate. VLANs are logical broadcast segmentations of a network. In order to send traffic from a device on VLAN A to a device on VLAN B, the traffic must go through a layer 3 routed interface. Previously, routers handled this task either using individual connections for each broadcast segment or a trunked connection from a switch to a router which is often referred to as a router-on-a-stick. Here is a diagram depicting a network utilizing single routed links:




Figure 21: A Network Using Single Routed Links

And here is a diagram showing the router-on-a-stick technique that uses a single trunked link carrying multiple VLANs:




Figure 22: A VLAN using the Router-On-A-Stick Method

More recently, switches themselves have incorporated the ability to both switch and route on a single device. **Inter-VLAN routing (IVR)** provides the mechanism to enable routing of traffic between VLANs. Switches capable of handling both switching and routing tasks are referred to as **multilayer switches (MLS)**. There are a few advantages to using the old method of using a standard router that handles the layer 3 work as opposed to using a multilayer switch. These are:

- Routers can handle a variety of different interface types that multilayer switches do not have including T1s, DS3s, HSSI and other serial interface types.
- The physical separation of layer 2 and layer 3 devices can make troubleshooting of a problem easier for less-experienced network administrators.

On the other hand, multilayer switches have these benefits:

- MLS devices use hardware ASICs that can route traffic at wire speed. This makes routing tasks much faster compared to routed methods.
- The backplane of a MLS device is much larger than most any interface port. This helps to eliminate any bottlenecks on the network when routing traffic from one VLAN to another.
- MLS devices provide many more software features to enhance the performance and security of routed traffic that you cannot achieve using standard routing methods.

The TSHOOT exam focuses mainly on troubleshooting IVR on MLS switches. The two most common commands used when troubleshooting IVR on any layer 3 device is the show ip route and show ip arp commands.

Show ip route

Everyone should be familiar with the show ip route command that builds a layer 3 table of IP subnets and their next hop address as shown here:

```
Router#show ip route

Codes: I - IGRP derived, R - RIP derived, O - OSPF derived,
       C - connected, S - static, E - EGP derived, B - BGP derived,
       * - candidate default route, IA - OSPF inter area route,
       i - IS-IS derived, ia - IS-IS, U - per-user static route,
       o - on-demand routing, M - mobile, P - periodic downloaded
       static route, D - EIGRP, EX - EIGRP external, E1 - OSPF
       external type 1 route, E2 - OSPF external type 2 route, N1 -
       OSPF NSSA external type 1 route, N2 - OSPF NSSA external type
       2 route

Gateway of last resort is 10.119.254.240 to network 10.140.0.0

O E2 172.150.0.0 [160/5] via 10.119.254.6, 0:01:00, Ethernet2
E    172.17.10.0 [200/128] via 10.119.254.244, 0:02:22, Ethernet2
O E2 172.70.132.0 [160/5] via 10.119.254.6, 0:00:59, Ethernet2
O E2 10.130.0.0 [160/5] via 10.119.254.6, 0:00:59, Ethernet2
E    172.30.0.0 [200/128] via 10.119.254.244, 0:02:22, Ethernet2
E    10.129.0.0 [200/129] via 10.119.254.240, 0:02:22, Ethernet2
```

Show ip arp

The ARP table maps layer 3 IP addresses to MAC addresses as seen in the output of this show command:

```
Router#show ip arp

Protocol  Address          Age (min)  Hardware Addr  Type   Interface
Internet  172.16.233.22    9          0000.0c59.f892  ARPA   Ethernet0/0
Internet  172.16.233.21    8          0000.0c07.ac00  ARPA   Ethernet0/0
Internet  172.16.233.19    -          0000.0c63.1300  ARPA   Ethernet0/0
Internet  172.16.233.30    9          0000.0c36.6965  ARPA   Ethernet0/0
Internet  172.16.168.11    -          0000.0c63.1300  ARPA   Ethernet0/0
```


The previous two commands should be known by CCNA level students. However, understanding ARP and IP routing is important to the next discussion of how MLS switches actually work and the new tables that are created.

How MLS works from a Layer 2 and Layer 3 point of View:

From a logical standpoint, you can think of a multilayer switch as two separate entities that work together. The **data plane** of the switch is responsible for forwarding data frames at layer 2 while the **control plane** routes packets between VLANs at layer 3. When a frame from one VLAN needs to get to a device on a different VLAN, it is passed from the data plane to the control plane across the switch **backplane**. This is when the frame is encapsulated in an IP packet and routed to the destination VLAN where it is then de-encapsulated, and the frame is then switched out the destination port on the data plane. This process uses either route caching or the more advanced Cisco Express Forwarding (CEF).

Route caching is a technique that uses a **route-switch processor (RSP)** to forward IP packets. This technology uses a route-one, switch many technique that identifies a traffic flow, spends time looking up routing information for the first packet and routes it accordingly. All other packets in the flow then do not have to go through the trouble of having the RSP look into the routing table to find the destination route. Instead, that information is cached in an ASIC and the packet is switched at wire-speed.

Cisco Express Forwarding (CEF) is a newer layer 3 switching technology that offers the following benefits compared to the RSP method:

- **Improved performance** – CEF is less CPU-intensive than route caching.
- **More scalable.**
- **Resilience** – CEF works better in inconsistent routing environments where there are many different data stream types and when constant changes to the routed network occur. The goal of CEF is to eliminate route-table lookups and the technology has what's known as a **Forwarding Information Base (FIB)** that keeps track of all known routes in the routing table. Accessing the FIB rather than the routing table helps to speed up the packet switching process.

Show ip cef

One key command with troubleshooting IVR when CEF is being used is the show ip cef command. This command shows routes pulled from the routing table and put into the FIB. The table also has the next-hop IP address and interface as seen here:

```
Switch#show ip cef
Prefix                Next Hop                Interface
0.0.0.0/0             192.168.1.5            FastEthernet0/0
0.0.0.0/32            receive
192.168.0.0/24        192.168.10.1           Serial0/2/0
192.168.2.0/30        192.168.10.1           Serial0/2/0
192.168.3.0/30        192.168.10.1           Serial0/2/0
192.168.4.0/24        192.168.10.1           Serial0/2/0
192.168.5.0/30        192.168.10.1           Serial0/2/0
192.168.6.0/30        192.168.11.1           FastEthernet0/1
192.168.7.0/30        192.168.11.1           FastEthernet0/1
192.168.8.0/24        192.168.11.1           FastEthernet0/1
192.168.9.0/30        192.168.11.1           FastEthernet0/1
```

When troubleshooting CEF problems, a good place to start is to make sure your FIB database (viewed by using the command above) matches your IP routing table (found using the show ip route command).

Show adjacency

Another useful command when troubleshooting multilayer switches running CEF is the show adjacency command. This command provides useful information for verifying an adjacency and interface for routes :

```
Switch#show adjacency

Protocol Interface          Address
IP        FastEthernet2/3          172.20.52.1 (3045)
IP        FastEthernet2/3          172.20.52.22 (11)
```

Troubleshooting SVIs

When you are using a multi-layer switch (MLS), instead of configuring physical interfaces with IP addresses, you configure switch virtual interfaces (SVI). SVIs are virtual layer 3 interfaces that have the ability to route between VLANs. There are a few distinct differences between routed interfaces and SVIs that you need to understand when troubleshooting:

- A routed port is shown as down when performing a show interface command if the port is not operational at both Layer 1 and 2.
- An SVI is shown as down when there are no active ports on that particular VLAN. As soon as an access or trunk port configured in the VLAN comes up, the VLAN will also transition to an up state.
- A routed port does not run layer 2 protocols such as any flavor of STP or any dynamic trunking protocols.

You can verify the state of an SVI by issuing the show interface vlan <vlan> command as seen here when viewing VLAN 11:

```
Switch#show interface vlan11
Vlan11 is up, line protocol is up
Hardware is EtherSVI, address is 0012.7f02.4b41 (bia 0012.7f02.4b41)
Internet address is 10.1.1.11/24
```

Troubleshooting Gateway Redundancy Protocols

In this section, we'll briefly review gateway redundancy protocols and then how network administrators can troubleshoot them.

Gateway redundancy is also known as first hop redundancy due to the fact that it provides end-device redundancy from a default gateway standpoint. The three main redundancy protocols that Cisco layer 3 devices utilize are Hot Standby Routing Protocol (HSRP), Virtual Router Redundancy Protocol (VRRP) and the Gateway Load Balancing Protocol (GLBP).

Hot Standby Routing Protocol (HSRP)

HSRP is a Cisco proprietary protocol that performs automatic router backup when you configure it on Cisco capable devices. There are three interface types that can be configured for HSRP:

1. Standard Routed Ports
2. MLS switched virtual interfaces (SVI)
3. Layer 3 EtherChannels

HSRP takes 2 redundant layer 3 gateways and groups them together. Once they are inside an HSRP group, they are configured to share a Virtual MAC and IP address. That virtual IP address is what the end stations use as the gateway address. One of the HSRP interfaces is the primary and the other is the standby interface. If for some reason, the primary fails, the standby interface detects the failure and takes over the gateway duties. HSRP uses a priority configuration to determine which layer 3 interface is primary. When HSRP is initially configured, the device has a default priority of 100. The highest priority becomes the primary interface. HSRP communicates to its peers by sending multicast messages.

These messages include:

- Hello – verifies that the routers are still functioning properly.
- Coup – when a standby router becomes the active router.
- Resign – when a router that is the active router sends this message when it is about to give up being the active due to another router with a higher priority coming online.

show standby <group>

The show standby <group> command displays real-time information about a configured standby group. This command displays detailed information which is valuable for troubleshooting such as:

- Last state change
- Virtual IP address
- Hello timer settings
- Hold timer settings
- Preemption enabled/disabled
- Priority settings

In addition, the “State” portion of the command lets you know what state the local MLS switch is currently in. A standby group can be in one of these five possible states:

- **Active** – the current HSRP MLS.
- **Standby** – the MLS/router is next in line to be the active HSRP switch.
- **Speak** – the MLS/router is sending HSRP hellos to neighbor MLS/routers in the same group to determine which switch will be active.
- **Listen** – the MLS/router is listening for HSRP messages from other MLS/routers in the group.
- **Init or Disabled** – the MLS/router is not yet ready or able to participate in HSRP, possibly because the associated interface (either physical or SVI) is not up. The state is listed as “Disabled” if the standby ip command has not been configured on the MLS/router.

Here is an example of the show standby <group> command:

```
Router#show standby 1
Ethernet0/1 - Group 1
  State is Active
    2 state changes, last state change 00:30:59
  Virtual IP address is 10.1.0.20
    Secondary virtual IP address 10.1.0.21
  Active virtual MAC address is 0004.4d82.7981
    Local virtual MAC address is 0004.4d82.7981 (bia)
  Hello time 4 sec, hold time 12 sec
    Next hello sent in 1.412 secs
  Preemption enabled, min delay 50 sec, sync delay 40 sec
  Active router is local
  Standby router is 10.1.0.6, priority 75 (expires in 9.184 sec)
  Priority 95 (configured 120)
  Tracking 2 objects, 0 up
    Down Interface Ethernet0/2, pri 15
    Down Interface Ethernet0/3
  IP redundancy name is "HSRP1", advertisement interval is 34 sec
```

show standby brief

Adding the keyword "brief" to the show standby command displays basic information about all of the HSRP interfaces configured on your Router or MLS as shown here:

```
Router#show standby brief

Interface  Grp Prio P State   Active addr   Standby addr   Group addr
Et0        0   120  Init   10.0.0.1      unknown        10.0.0.12
```

debug standby

The final HSRP troubleshooting command a TSHOOT candidate should be familiar with is the debug standby command. This command outputs HSRP hello and State information to the command line. This can help a network administrator to see if hello's are being properly sent and received as well as possibly identifying what state in the HSRP process the MLS/router is failing at:

```
Switch#debug standby

*Mar 1 02:55:56: SB0:FastEthernet3/0 Hello out 10.144.220.2 Active pri
110 hel 3 hol 10 ip 10.144.220.1
*Mar 1 02:56:08: SB0:FastEthernet3/0 Hello in 10.144.220.3 Active pri
120 hel 3 hol 10 ip 10.144.220.1
*Mar 1 02:56:08: SB0: FastEthernet3/0 state Active -> Speak
*Mar 1 02:56:08: SB0:FastEthernet3/0 Resign out 10.144.220.2 Speak pri
110 hel 3 hol 10 ip 10.144.220.1
*Mar 1 02:56:08: SB0:FastEthernet3/0 Hello out 10.144.220.2 Speak pri
110 hel 3 hol 10 ip 10.144.220.1
```

```
*Mar 1 02:56:09: %LINEPROTO-5-UPDOWN: Line protocol on Interface
FastEthernet3/0, changed state to down
*Mar 1 02:56:11: SB0: FastEthernet3/0 state Speak -> Init
*Mar 1 02:56:13: %LINEPROTO-5-UPDOWN: Line protocol on Interface
FastEthernet3/0, changed state to up
*Mar 1 02:56:13: SB0: FastEthernet3/0 state Init -> Listen
*Mar 1 02:56:14: %LINEPROTO-5-UPDOWN: Line protocol on Interface
FastEthernet3/0, changed state to down *Mar 1 02:56:14: SB0:
FastEthernet3/0 state Listen -> Init
*Mar 1 02:56:20: %LINEPROTO-5-UPDOWN: Line protocol on Interface
FastEthernet3/0, changed state to up
*Mar 1 02:56:20: SB0: FastEthernet3/0 state Init -> Listen
*Mar 1 02:56:30: SB0: FastEthernet3/0 state Listen -> Speak
*Mar 1 02:56:40: SB0: FastEthernet3/0 state Speak -> Standby
*Mar 1 02:56:41: SB0: FastEthernet3/0 state Standby -> Active
*Mar 1 02:56:41: SB: FastEthernet3/0 Adding 0000.0c07.ac00 to address
filter
*Mar 1 02:56:41: SB0:FastEthernet3/0 Hello out 10.144.220.2 Active pri
110 hel 3 hol 10 ip 10.144.220.1
*Mar 1 02:56:44: SB0:FastEthernet3/0 Hello in 10.144.220.3 Active pri
120 hel 3 hol 10 ip 10.144.220.1
*Mar 1 02:56:44: SB0: FastEthernet3/0 state Active -> Speak
```

Virtual Router Redundancy Protocol (VRRP)

VRRP is an IETF standard gateway redundancy protocol. Because it's a standard, it can be used in multi-vendor environments. VRRP configuration enables a group of layer 3 devices to form a single virtual router. The end devices then use the virtual router as their default gateway address. VRRP works very much like HSRP but instead of creating a Virtual IP address, VRRP uses the master interface physical IP address. If for some reason, the master VRRP router fails, the backup VRRP router takes over the IP address and applies it to its physical interface.

VRRP uses a priority feature to determine which device is the master and which ones are the backups. The highest priority is the master. By default, the master will always preempt and become the master when it comes online. The priority can be configured between 1 and 254. If two routers have the same priority, the router with the highest IP address becomes the master. The default priority is 100.

show vrrp brief

To view VRRP information, issue the show vrrp brief command. This command displays the information found in the following table:

Field	Description
Interface	Interface or SVI used
Group	VRRP group
Prio	VRRP priority (higher is preferred)
Time	Amount of time before backup takes over
Own	Owner of the IP
Pre	If preemption (P) is configured
State	Role the interface or SVI is currently in
Master addr	IP of the master virtual router
Group addr	IP of the virtual router

Figure 23: VRRP Field Information

Here is an example of the command output on a router:

```
Router#show vrrp brief
Interface  Grp Prio  Time  Own Pre  State  Master addr  Group addr
Ethernet1/0  1  100  3609      P  Master  1.0.0.4  1.0.0.10
Ethernet1/0  2  105  3589      P  Master  1.0.0.4  1.0.0.20
```

Gateway Load Balancing Protocol (GLBP)

The main difference between GLBP and HSRP/VRRP is in its ability to load balance by default. With HSRP and VRRP, there are ways to design your network to load balance at the VLAN level, but this can lead to an imbalance if one VLAN utilizes the vast majority of traffic. GLBP is a per-MAC address way to load balance while still providing gateway redundancy. GLBP routers configured in the same group communicate between each other using multicast hello packets.

With GLBP, a router is elected to be the **active virtual gateway (AVG)**. The AVG's responsibility is to hand out virtual MAC addresses to the other routers in the GLBP group. All the other routers in the group are called active virtual forwarders (AVF). They use the same gateway IP address as the rest but with a different MAC address. End devices are all configured with the same gateway address. The difference is that when the end device does an ARP lookup to get the MAC address of its gateway IP, the AVG hands out the different MAC addresses of itself and the other AVFs that are configured. This is how load balancing is achieved. A single GLBP group can have up to four AVFs. If one of the GLBP routers becomes unavailable, the other router assumes responsibility to handle traffic for the downed MAC address and its own MAC address. The end devices never notice because they're still sending their traffic to the same layer 2 MAC and layer 3 IP address.

show glbp brief

Similar to both HSRP and VRRP, the show glbp brief command displays information about the currently active glbp virtual gateways. Similar to the others, the command output displays real-time information including:

- Interface
- Group
- Priority
- Address
- Active Router
- Standby Router

At this point, you begin to see that the three gateway load balancing options have fairly similar troubleshooting outputs. The key is to know the unique differences to be able to differentiate between them. Here is the output of the command:

```
Router#show glbp brief
```

Interface	Grp	Fwd	Pri	State	Address	Active router	Standby router
Fa0/0	10	-	254	Active	10.21.8.10	local	unknown
Fa0/0	10	1	7	Active	0007.b400.0101	local	-

Troubleshooting Network Derogation Problems

Troubleshooting issues that are not configuration problems can often times be the most difficult. A decent amount of trial-and-error must go into the troubleshooting process if you are fairly new network administrator. After time, you will begin to have a “sixth-sense” about many problems however and will have the ability to resolve them quickly. This section covers a few of the more common non-configuration related problems. First we’ll look at two common interface related problems mainly found on access port interfaces that connect end devices. Then we’ll look at problems with the switch forwarding process of a multilayer switch and why the TCAM might “punt” packets up to the main CPU causing derogation problems.

Straight-through vs. crossover

There are two basic standards for Ethernet pinouts. The first is the straight-through cable. This cable can be used to connect non-like devices together such as a PC to a switch or a switch to a router. If you needed to connect two of the same device together, you needed to use what’s known as a crossover cable which literally flips the Rx and Tx wires on each end of the cable.

More recently, Cisco has implemented what’s known as automatic medium dependent interface crossover (auto-MDIX). This switchport feature can be enabled so the switchport can detect which pins to transmit and receive based on learning if the connected device is another switch or a different device. The command to configure this feature on a switchport looks like this:

```
Switch#configure terminal
Switch(config)#interface fa0/10
Switch(config-if)#mdix auto
```

Speed/Duplex Mismatch

Network derogation on an Ethernet LAN is often times a simple speed and duplex mismatch. A duplex mismatch results in slow performance, dropped frames, data link errors, and other issues:

One side of the connection is hard coded and the other is set to auto-negotiate. In this example, our FastEthernet 0/5 switchport is set to auto-negotiate. When the port comes online it looks like this:

```
Switch#sh int fa0/5
FastEthernet0/7 is up, line protocol is up (connected)
  Hardware is Fast Ethernet, address is 000b.4660.8407 (bia
000b.4660.8407)
  Description:
  MTU 1500 bytes, BW 100000 Kbit, DLY 100 usec,
    reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  Keepalive set (10 sec)
  Half-duplex, 100Mb/s, media type is 100BaseTX
  input flow-control is off, output flow-control is unsupported
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input never, output 00:00:00, output hang never
  Last clearing of "show interface" counters never
  Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
  Queueing strategy: fifo
  Output queue: 0/40 (size/max)
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 1000 bits/sec, 2 packets/sec
    45462786 packets input, 710370902 bytes, 0 no buffer
    Received 88119 broadcasts (0 multicast)
    0 runts, 0 giants, 0 throttles
    87342 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
    0 watchdog, 0 multicast, 0 pause input
    0 input packets with dribble condition detected
  521051404 packets output, 3353431321 bytes, 0 underruns
  0 output errors, 0 collisions, 6 interface resets
  0 babbles, 0 late collision, 0 deferred
  0 lost carrier, 0 no carrier, 0 PAUSE output
  0 output buffer failures, 0 output buffers swapped out
```

You'll notice two things about the command output. First, the port came up as a half-duplex link. Secondly there are a high amount of "input errors" on the interface. This is a second sign that the device that is connected to this port is hard-coded as a 100/full link. To remedy this situation, it is recommended that you change the end-device to auto-negotiate speed/duplex settings.

Previously, it was widely practiced to manually set speed and duplex settings on both sides of the connection. This was because the auto-negotiation feature had some bugs to work out. Now, just the opposite is recommended as a best-practice by Cisco. Network Administrators should let auto-negotiation set the speed/duplex settings on devices.

Switch Forwarding Processes

A second network derogation problem TSHOOT candidates should familiarize themselves with is how to understand and troubleshoot the multilayer switch packet forwarding processes when using TCAM memory.

A multilayer switch's forwarding logic is compiled into a special type of memory called ternary content addressable memory (TCAM). TCAM works with Cisco Express Forward (CEF) feature to provide extremely fast forwarding decisions. Sometimes there are circumstances where packets cannot flow from the CEF to the TCAM and instead traffic is forwarded to the multilayer switch's main CPU. This CPU is not optimized for forwarding packets like the TCAM and therefore increases overall CPU utilization and slows down the maximum number of packets the MLS can process. The process of sending packets to the main CPU for processing is called "punting".

There are several reasons why a packet might be punted from the TCAM to the main CPU.

These reasons include:

- If a switch's TCAM has reached capacity and cannot process any more packets, all packets that go over the TCAM limit will be punted to the main CPU. An overtaxed MLS switch with too many routes or ACLs to process can cause the TCAM to reach capacity. This is the most common occurrence of punting on most networks.
- Routing protocols and data control plane protocols (Like STP) communicate using either broadcast or multicast packets. This traffic will be sent to the main CPU.
- An administrator that remotely connects to the local switch using Telnet/SSH or the web GUI will have their packets sent to the main CPU.
- Packets using a feature that not supported in by hardware ASICs such as traffic traversing a GRE tunnel are sent to the CPU.

Troubleshooting Switch Supervisor Redundancy

Modular switches such as the Catalyst 6500 series platform support supervisor redundancy. Two supervisor cards can be placed into a switch. Only one supervisor module will function at a time. If for some reason, the active supervisor module were to fail, the second module would assume routing and switching responsibilities. There are two types of supervisor redundancy currently available on Cisco switches, they are:

- Route Processor Redundancy (RPR)
 - ▶ Secondary supervisor module waits for a failure and then boots and initializes. The process of taking over routing and switching is between 2 and 4 minutes
- Route Processor Redundancy Plus (RPR+)
 - ▶ A newer supervisor redundancy method where the secondary processor is online and fully initialized. The process of taking over routing and switching is between 30 and 60 seconds
 - ▶ The IOS versions must be identical when using RPR+ otherwise an error will occur and RPR+ will fall back to using standard RPR

A TSHOOT candidate should know how to verify the proper operation of RPR and RPR+. To verify this, use the `show redundancy states` command as shown below. Notice you can see the primary supervisor is ACTIVE and the peer state is STANDBY HOT:

```
Router#show redundancy states
my state = 13 -ACTIVE
  peer state = 8 -STANDBY HOT
    Mode = Duplex
    Unit = Primary
    Unit ID = 1

Redundancy Mode (Operational) = Route Processor Redundancy Plus
Redundancy Mode (Configured) = Route Processor Redundancy Plus
  Split Mode = Disabled
  Manual Swact = Enabled
  Communications = Up

  client count = 11
  client_notification_TMR = 30000 milliseconds
  keep_alive TMR = 9000 milliseconds
  keep_alive count = 0
  keep_alive threshold = 18
  RF debug mask = 0x0
```

Troubleshooting IP Routing Protocols

When troubleshooting any type of connectivity problems, many network administrators choose to use the “divide-and-conquer” troubleshooting method and begin working a problem at the network layer of the OSI model. Typically the first step is to check IP connectivity by pinging from the source to the destination to see if it is successful. If the problem is related to a routing protocol there are some general troubleshooting tips that work for troubleshooting all IP routing protocols.

Basic Routing Concepts

Remember that when a PC residing on one subnet needs to communicate to a PC residing on a different subnet, the PC must send the data to its default gateway IP address. The default gateway is going to be either a physical interface on a router or an SVI on a multilayer switch. If the router or MLS knows that the destination subnet of the PC resides off one of the network it is configured for locally, the router does a routing table lookup and sees the networks listed as “C” for connected. The router then forwards traffic out that specific interface or SVI that is listed in the routing table.

However, if the subnet that the destination PC resides on is not directly connected to the router, the router must rely on some type of routing to know where the remote network is located. For the sake of the TSHOOT exam, these routing methods include the following:

- Static routes
- EIGRP
- OSPF
- BGP

In addition to routing, a TSHOOT candidate must be familiar with the following routing terms:

- **IP routing table** – a database that stores the routes and metrics to remote networks. This information contains the topology of all connected networks as well as any statically configured routes and routes learned through dynamic routing protocols.
- **ARP table** – a table of MAC address to IP address mappings that are either dynamically or statically learned by the local router.
- **Forwarding Information Base (FIB)** – similar to a routing table, the FIB contains information regarding how to reach remote networks. The difference between the FIB and a standard routing table is that the FIB is optimized for very fast routing table lookups of destination addresses. The FIB is used on CEF enabled routers.
- **Adjacency Table** – another table used by CEF enabled devices. This table maintains layer 2 next-hop addresses for layer 3 FIB entries.

Administrative Distance

When there are multiple routes to the same network found by different static or dynamic routing protocols, the method that has the lowest administrative distance will be placed into the routing table. Administrative distance is considered to be how “believable” the method is. For example, connected routes are the most believable, because they are directly connected to the local router. Static routes are the next most believable because the thinking is, since an administrator had to manually enter the route, it must be believable. Lastly, all dynamic routing protocols have administrative distances. Some protocols are considered to be more believable than others. Also, some protocols have different distances depending on how the route was learned (redistributed routes are less believable). The following shows the Administrative distances for various protocols:

Method	Distance
Connected	0
Static	1
EIGRP Summary	5
External BGP	20
Internal EIGRP	90
OSPF	110
RIP	120
External EIGRP	170
Internal BGP	200
Unknown	255

Figure 24: Administrative Distances

Basic Routing Concepts

When troubleshooting basic routing problems, don't forget to investigate both the control AND data plane portions of the router/MLS. Many people often check the routing table by issuing the show ip route command. This is part of the control plane. However, don't forget that the FIB is part of the data plane and also must be checked for proper operation! Below are some of the best-practice troubleshooting commands used to troubleshoot both the control and data planes:

Show ip route <network> <mask>

The show ip route <network> <mask> command details the current state of the router's IP routing table for a specific network. You can see that from the example below, even though we've specified a /24 network, the command output displays the /16 supernet network:

```
Router#show ip route 192.168.10.0 255.255.255.0
Routing entry for 192.168.0.0/16, supernet
  Known via "eigrp 100", distance 170, metric 3072, type external
  Redistributing via eigrp 100
  Last update from 192.168.20.17 on GigabitEthernet1/2, 7w0d ago
  Routing Descriptor Blocks:
  * 192.168.20.17, from 192.168.20.17, 7w0d ago, via GigabitEthernet1/2
    Route metric is 3072, traffic share count is 1
    Total delay is 20 microseconds, minimum bandwidth is 1000000 Kbit
    Reliability 255/255, minimum MTU 1500 bytes
    Loading 1/255, Hops 2
```

Show ip route <network> <mask> longer-prefixes

If we add the keyword longer-prefixes for our supernet network, the command output displays that only routes matching the ip-address and mask pair or any higher mask should be displayed. In our example, since a /24 mask is specified, we see all routes that are /24 and higher:

```
Router#show ip route 192.168.10.0 255.255.255.0 longer-prefixes
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS
level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static
route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is 192.168.20.17 to network 0.0.0.0

 192.168.10.0/32 is subnetted, 7 subnets
 D       192.168.10.32
         [90/179712] via 192.168.20.17, 7w0d, GigabitEthernet1/2
```

```

D      192.168.10.40
        [90/131072] via 192.168.20.17, 7w0d, GigabitEthernet1/2
D      192.168.10.5
        [90/130816] via 192.168.20.17, 7w0d, GigabitEthernet1/2
D      192.168.10.10
        [90/131072] via 192.168.20.17, 7w0d, GigabitEthernet1/2
D      192.168.10.21
        [90/131328] via 192.168.69.2, 7w0d, GigabitEthernet1/1
        [90/131328] via 192.168.20.17, 7w0d, GigabitEthernet1/2
D      192.168.10.20
        [90/131072] via 192.168.20.17, 7w0d, GigabitEthernet1/2
D      192.168.10.30
        [90/130816] via 192.168.20.17, 7w0d, GigabitEthernet1/2

```

Show ip cef <network> <mask>

If you are having routing difficulties but your routing table looks fine, the next place a network administrator should look is the FIB table. To display current FIB table entries, use the show ip cef <network> <mask> command shown here:

```

Router#show ip cef 192.168.60.0 255.255.255.0
192.168.60.0/24
    nexthop 192.168.60.2 GigabitEthernet2/1

```

Show ip cef exact-route <source_ip> <destination_ip>

An enhancement to the standard show ip cef command is to use the exact-route keyword. The network administrator then specifies a specific source and destination IP address to see the specific route a packet would take given the source and destination IPs as shown in this example where we want to see the next-hop path a packet sourced from 10.10.10.1 to the destination address of 192.168.20.10:

```

Router#show ip cef exact-route 10.10.10.1 192.168.20.10
10.10.10.1      -> 192.168.20.10  :Ethernet2/0/0 (next hop
172.1.1.1)

```

Show ip arp

The show ip arp command is one of the most frequently used commands in a network administrator's arsenal. It is used to display IP address to MAC address mappings as shown here:

```
Router#show ip arp
```

Protocol	Address	Age (min)	Hardware Addr	Type	Interface
Internet	172.16.233.22	9	0000.0c59.f892	ARPA	Ethernet0/0
Internet	172.16.233.21	8	0000.0c07.ac00	ARPA	Ethernet0/0
Internet	172.16.233.19	-	0000.0c63.1300	ARPA	Ethernet0/0
Internet	172.16.233.30	9	0000.0c36.6965	ARPA	Ethernet0/0
Internet	172.16.168.11	-	0000.0c63.1300	ARPA	Ethernet0/0
Internet	172.16.168.254	9	0000.0c36.6965	ARPA	Ethernet0/0

Show frame-relay map

If you have remote sites that utilize frame-relay circuits, a very common troubleshooting command is show frame-relay map. This command displays the following information useful for troubleshooting:

- Link status (up or down)
- Layer 2 DLCI to IP address mappings
- Dynamically or statically configured mapping
- Frame-relay encapsulation type
- If TCP/IP header compression is used

Here is an example of the output found using this show command:

```
Router#show frame-relay map
Serial1/2 (up): ip 172.16.1.4 dlci 401(0x191,0x6410), dynamic,
broadcast,, status defined, active Serial1/2 (up): ip 172.16.1.5
dlci 501(0x1F5,0x7C50), dynamic, broadcast,, status defined,
active Serial1/2 (up): ip 172.16.1.2 dlci 301(0x12D,0x48D0), dynamic,
broadcast,, status defined, active
TCP/IP Header Compression (inherited), passive (inherited)
```

Show adjacency detail

The show adjacency detail command lists all of the local router's routing protocol information and adjacencies with detailed timer information as shown here:

```
Router#show adjacency detail
Protocol Interface Address
IP EOBC0/0 127.0.0.51(4)
0 packets, 0 bytes
epoch 0
sourced in sev-epoch 11
Encap length 14
00001A0000000000150000000700
ARP
IP Vlan99 192.168.99.2(7)
4624 packets, 402636 bytes
epoch 0
sourced in sev-epoch 11
Encap length 14
00602B0021FE000E392C78000800
ARP
```

Debug ip routing

The debug ip routing command shows debug level routing information for all IP routing protocols. It details information such as routing table and route cache updates as shown in the output here of a router running the RIP protocol:

```
Router#debug ip routing

RT: add 172.25.168.0 255.255.255.0 via 172.24.76.30, igrp metric
[100/3020]
RT: metric change to 172.25.168.0 via 172.24.76.30, igrp metric
[100/3020]
      new metric [100/2930]
IP: cache invalidation from 0x115248 0x1378A, new version 5736
RT: add 172.26.219.0 255.255.255.0 via 172.24.76.30, igrp metric
[100/16200]
RT: metric change to 172.26.219.0 via 172.24.76.30, igrp metric
[100/16200]
      new metric [100/10816]
RT: delete route to 172.26.219.0 via 172.24.76.30, igrp metric
[100/10816]
RT: no routes to 172.26.219.0, entering holddown
IP: cache invalidation from 0x115248 0x1378A, new version 5737
RT: 172.26.219.0 came out of holddown
RT: garbage collecting entry for 172.26.219.0
IP: cache invalidation from 0x115248 0x1378A, new version 5738
RT: add 172.26.219.0 255.255.255.0 via 172.24.76.30, igrp metric
[100/10816]
RT: delete route to 172.26.219.0 via 172.24.76.30, igrp metric
[100/10816]
RT: no routes to 172.26.219.0, entering holddown
IP: cache invalidation from 0x115248 0x1378A, new version 5739
RT: 172.26.219.0 came out of holddown
RT: garbage collecting entry for 172.26.219.0
IP: cache invalidation from 0x115248 0x1378A, new version 5740
RT: add 172.26.219.0 255.255.255.0 via 172.24.76.30, igrp metric
[100/16200]
RT: metric change to 172.26.219.0 via 172.24.76.30, igrp metric
[100/16200]
      new metric [100/10816]
RT: delete route to 172.26.219.0 via 172.24.76.30, igrp metric
[100/10816]
RT: no routes to 172.26.219.0, entering holddown
IP: cache invalidation from 0x115248 0x1378A, new version 5741
```

Troubleshooting EIGRP

The Enhanced Interior Gateway Routing Protocol (EIGRP) is an advanced distance vector protocol routing protocol. This protocol was developed by Cisco Systems and is considered to be proprietary. The CCNP route goes into great detail describing this protocol and how to configure it. The TSHOOT exam covers what you need to know about the protocol to troubleshoot it as well as understanding different show and debug commands for troubleshooting purposes.

EIGRP Neighbors

Once a routing protocol is enabled on a router and networks are defined, the EIGRP process begins sending out EIGRP hello multicast packets out interfaces defined in the eigrp network statements. Other routers that are also configured with EIGRP on these networks will listen for and return hello packets. Once both routers become aware of each other and establish a neighbor relationship, they proceed to exchange their full routing table. After this single occurrence of a full table exchange, only routing table updates are sent between neighbor routers to reduce the amount of bandwidth the routing protocol consumes.

EIGRP Tables

Routing with EIGRP requires for different tables as described here:

- **Interface Table** – table listing all interfaces that EIGRP is running on.
- **Neighbor Table** – table listing all known EIGRP neighbors.
- **Topology Table** – table listing all EIGRP routes. This includes successor and feasible successor routes. The EIGRP DUAL algorithm uses this table to select the best option to be placed into the routing table. The best routes are defined as the lowest cost to a destination.
- **Routing Table** – the table used by the router for path selection to remote networks.

EIGRP also supports both equal and unequal cost load balancing. By default, if there are multiple paths to a network with an equal cost, the router will load balance traffic for up to four paths. If an administrator wants to load balance using unequal cost paths, the variance command can be used. This command defines a range of costs. If the unequal costs fall within the specified range, the router will load balance traffic between the multiple paths.

Routing information learned from adjacent neighbors is inserted into the EIGRP topology table. The best route for a specific network in the IP EIGRP topology table is called the successor route and is placed into the routing table. Other routes to the same network that have a higher cost are considered to be feasible successor routes.

Here is how EIGRP chooses the best routes to be placed into the routing table:

- **Advertised Distance** – The metric a neighbor router is advertising to a remote network.
- **Feasible Distance** – The metric a neighbor router is advertising to a remote network plus the additional metric cost of the link to that neighbor.

EIGRP Metric

So as you can see, the metric calculation plays a vital role in the selection of successor routes which are ultimately placed into the routing table. There is a formula used to calculate the metrics for each remote route. The full EIGRP calculation is as follows:

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

This calculation may be a bit intimidating at first, but once you understand the K values, it becomes much easier to understand how to calculate metrics.

EIGRP default K values are:

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

So you can see that by default, EIGRP only uses bandwidth and delay values to calculate the metrics. All other K values are 0 and when 0 is multiplied by anything, the result is obviously, 0.

Essential EIGRP Troubleshooting Commands

This next section describes the most important show and debug commands that can be used to troubleshoot EIGRP networks.

Show ip eigrp interfaces

If an administrator wants to see which EIGRP interfaces are currently active, a great command to use is the show ip eigrp interfaces command. In addition, the command displays the following information regarding the individual EIGRP interfaces:

Field	Description
Interface	Interface that EIGRP is active on
Peers	Number of EIGRP peers on that interface
Xmit Queue Un/Reliable	Number of packets in queues
Mean SRTT	Mean smooth round-trip time (ms)
Pacing Time Un/Reliable	Time used to pace EIGRP protocol packets
Multicast Flow Timer	Maximum number of seconds the interface will send multicast EIGRP packets
Pending Routes	Number of routes in queued packets

Figure 25: Information from the show ip eigrp interfaces Command

Here is an example of the command output where you can gather all of this information:

```
Router#show ip eigrp interfaces
IP EIGRP interfaces for process 100
```

Interface	Peers	Xmit Queue Un/Reliable	Mean SRTT	Pacing Time Un/Reliable	Multicast Flow Timer	Pending Routes
Di0	0	0/0	0	11/434	0	0
Et0	1	0/0	337	0/10	0	0
SE0:1.16	1	0/0	10	1/63	103	0
Tu0	1	0/0	330	0/16	0	0

Show ip eigrp neighbors

If an EIGRP network is suffering from stability issues, use the show ip eigrp neighbors command to not only see what EIGRP neighbors the local router has discovered but also to view the uptime of the neighbor connection as shown here:

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

Address	Interface	Holdtime (secs)	Uptime (h:m:s)	Q Count	Seq Num	SRTT (ms)	RTO (ms)
192.168.81.28	Ethernet1	13	0:48:41	0	11	4	20
192.168.80.28	Ethernet0	14	0:48:41	0	10	12	24
192.168.80.31	Ethernet0	12	0:00:02	0	1	5	20

Just looking at the uptime for these three neighbors, one can see that two neighbors have an uptime of over 48 minutes while the 192.168.80.31 neighbor has an uptime of only two seconds. This may be the source of a stability problem.

Show ip eigrp topology

The show ip route command will show which EIGRP routes are placed into the routing table. However, if you are troubleshooting EIGRP and need to examine all the routes (successor and feasible successor) that the DUAL algorithm has calculated, the show ip eigrp topology command will give you the detail that you need. The output lists all the learned EIGRP routes, lists the successors, next hop routes and interfaces as well as the metrics (cost to destination and advertised cost):

```
Router#show ip eigrp topology

IP-EIGRP Topology Table for process 100
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - Reply status
P 172.16.90.0 255.255.255.0, 2 successors, FD is 0
   via 172.16.80.28 (46251776/46226176), Ethernet0
   via 172.16.81.28 (46251776/46226176), Ethernet1
   via 172.16.80.31 (46277376/46251776), Serial0
P 172.16.81.0 255.255.255.0, 1 successors, FD is 307200
   via Connected, Ethernet1
   via 172.16.81.28 (307200/281600), Ethernet1
   via 172.16.80.28 (307200/281600), Ethernet0
   via 172.16.80.31 (332800/307200), Serial0
```

Show ip eigrp traffic

One often overlooked eigrp troubleshooting command that nicely displays the number of EIGRP messages sent and received on an EIGRP enabled router is the show ip eigrp traffic command. This can help to identify specific EIGRP messages that are not getting through. Here is an example of the output:

```
Router#show ip eigrp traffic
IP-EIGRP Traffic Statistics for process 100
  Hellos sent/received: 218/205
  Updates sent/received: 7/23
  Queries sent/received: 2/0
  Replies sent/received: 0/2
  Acks sent/received: 21/14
```

Debug ip eigrp

To display real-time EIGRP packet processing on a local switch, enable the debug ip eigrp command. The output can be useful to analyze the EIGRP packets being sent and received. Here is an example of this command output:

```
Router#debug ip eigrp
IP-EIGRP: Processing incoming UPDATE packet
IP-EIGRP: Ext 192.168.3.0 255.255.255.0 M 386560 - 256000 130560 SM
360960 - 256000 104960
IP-EIGRP: Ext 192.168.0.0 255.255.255.0 M 386560 - 256000 130560 SM
360960 - 256000 104960
IP-EIGRP: Ext 192.168.3.0 255.255.255.0 M 386560 - 256000 130560 SM
360960 - 256000 104960
IP-EIGRP: 172.69.43.0 255.255.255.0, - do advertise out Ethernet0/1
IP-EIGRP: Ext 172.69.43.0 255.255.255.0 metric 371200 - 256000 115200
IP-EIGRP: 192.135.246.0 255.255.255.0, - do advertise out Ethernet0/1
IP-EIGRP: Ext 192.135.246.0 255.255.255.0 metric 46310656 - 45714176
596480
IP-EIGRP: 172.69.40.0 255.255.255.0, - do advertise out Ethernet0/1
IP-EIGRP: Ext 172.69.40.0 255.255.255.0 metric 2272256 - 1657856 614400
IP-EIGRP: 192.135.245.0 255.255.255.0, - do advertise out Ethernet0/1
IP-EIGRP: Ext 192.135.245.0 255.255.255.0 metric 40622080 - 40000000
622080
IP-EIGRP: 192.135.244.0 255.255.255.0, - do advertise out Ethernet0/1
```

Debug eigrp packet

The debug ip eigrp command is great but when you want to see the request and acknowledgement of EIGRP packets, the debug eigrp packet command is a much better option. This command displays useful information such as:

- Sending or Receiving Packet
- EIGRP Packet type (HELLO, UPDATE, REQUEST, QUERY or REPLY)
- Interface number
- IP address (where applicable)
- AS number

Here is an example of this command output:

```
Router#debug eigrp packet
EIGRP: Sending HELLO on Ethernet0/1
      AS 100, Flags 0x0, Seq 0, Ack 0
EIGRP: Sending HELLO on Ethernet0/1
      AS 100, Flags 0x0, Seq 0, Ack 0
EIGRP: Sending HELLO on Ethernet0/1
      AS 100, Flags 0x0, Seq 0, Ack 0
EIGRP: Received UPDATE on Ethernet0/1 from 192.168.78.4,
      AS 100, Flags 0x1, Seq 1, Ack 0
EIGRP: Sending HELLO/ACK on Ethernet0/1 to 192.168.78.4,
      AS 100, Flags 0x0, Seq 0, Ack 1
EIGRP: Sending HELLO/ACK on Ethernet0/1 to 192.168.78.4,
      AS 100, Flags 0x0, Seq 0, Ack 1
EIGRP: Received UPDATE on Ethernet0/1 from 192.168.78.4,
      AS 100, Flags 0x0, Seq 2, Ack 0
```

Troubleshooting OSPF

OSPF is an open standard routing protocol that has very fast convergence similar to EIGRP. OSPF and EIGRP are the two main routing protocols found in large Enterprise networks for routing within an organization. OSPF uses a link-state method of determining the best routes to a remote network. This section will cover the basics needed to troubleshoot OSPF and some of the show and debug commands commonly used to identify OSPF related problems.

OSPF Tables

OSPF deals with the following five tables:

- **Interface Table** – table listing all of the local interfaces participating in OSPF. The participating interfaces are based on the network command statements listing the IP networks that OSPF should advertise. The interfaces that are configured on those defined OSPF network statements will be listed here.
- **Neighbor Table** – table listing all of the OSPF neighbors that the local router had formed adjacencies with.
- **Link-state Database** – table containing all of the OSPF learned routes and how to reach them.
- **Local Routing Information Base (RIB)** – table listing all of the OSPF learned routes and the costs after the SPF (Dijkstra) algorithm has been calculated. The best path metric discovered using the SPF algorithm will be placed into the routing table.
- **Routing Table** – the table used by the router for path selection to remote networks.

OSPF Metric

The OSPF Dijkstra shortest path first (SPF) algorithm calculates the cost to a remote network, lower the cost, the better the path. OSPF uses a simple calculation to determine this metric:

Cost = 100,000,000 / bandwidth (kbps)

Purpose of the Designated Router (DR)

A multi-access network can have multiple routers residing on a common network segment. Instead of all routers creating a full-mesh of adjacencies with one another, a single designated router (DR) will be elected, and all other routers on the segment can form an adjacency with the DR. This method helps to eliminate a great deal of routing protocol overhead in the form of multicast packet communication between neighbors. A backup designated router (BDR) is also chosen to take over DR responsibilities if the DR were to be unreachable due to some type of failure.

The most popular example of a multi-access network that would use DRs and BDRs is an Ethernet segment with several routers within the same IP subnet.

Purpose of OSPF Areas

OSPF routers are grouped logically into areas that are identified by a unique area number. All OSPF networks must have an area 0 called the backbone area. Any other routers configured in a different area must connect directly to area 0 or at least tunneled using OSPF virtual link techniques. The purpose of breaking an OSPF network into multiple areas is to reduce the size and frequency of routing tables and their updates. A router defined in one area knows about routes to other routes in their area only. Any other networks outside of the area are summarized into LSA type 3 routes which are propagated to routers for distribution. Depending on the location of the router within an OSPF area, a router can be defined as being one of the four following types:

- Internal Router
- Area Border Router (ABR)
- Backbone Router
- Autonomous System Boundary Router (ASBR)

The following diagram depicts each of these four OSPF router types:




Figure 26: OSPF Router Types

Link State Algorithm Types

In addition to the OSPF router types, each router will send different link state algorithm (LSA) messages. The following table defines each LSA type:

LSA Type	Description
Type 1	Router LSAs generated by each router and sent to all other routers within an area.
Type 2	Network LSAs generated by the DR and sent to all other routers within an area.
Type 3/4	Summary LSAs generated by ABRs that contain information about inter-area routes. Type 3s are routes to networks and Type 4s are routes to ASBRs.
Type 5	Network LSAs generated by the ASBRs that contains information about external links. These are sent to all other areas except for stub areas.
Type 6	Group membership LSAs generated by multicast OSPF routers. These are not supported by Cisco routers.
Type 7	NSSA route LSAs generated by the ASBR and sent to NSSA routers. The ABR converts type 7 LSAs to Type 5 before flooding them to the backbone.

Figure 27: LSA Types

OSPF Network Types

OSPF treats different network types differently based on the technology. The reason it does this is so OSPF can be run as optimally as possible depending on the type of LAN or WAN connection being used. Below are the OSPF Network types available:

- **Broadcast** – a network segment that connects a large number of broadcast-capable devices. Typically found on Ethernet connections.
- **Non-broadcast** – a network segment that connects a large number of devices that are not broadcast capable. Typically found on ATM and frame relay connection.
- **Point-to-Point** – a link between exactly two points. Typically found on WAN circuits such as T1 data serial connection.
- **Point-to-Multipoint** – organizes PVCs into a group of logical point-to-point connections. Typically found on frame relay connections.

OSPF Hellos

Routers running OSPF must establish adjacencies with neighbors. We know that the routers send Hello packets that contain information about the router and other details that help the protocol to operate properly. There are specific settings that must exactly match on any router that wishes to form an adjacency with a neighbor. These settings include:

- Hello Timer (default 10 seconds)
- Dead Timer (default 40 seconds)
- Area Number
- Area Type
- Subnet
- Authentication details (if configured)

Adjacency Status States

OSPF routers go through a process of transitioning between different adjacency status states before becoming fully adjacent with a neighbor. The following is the correct order of adjacency states that all OSPF neighbors must go through.

1. Down
2. Init
3. 2-Way
4. ExStart
5. Exchange
6. Loading
7. Full

As soon as the neighbors reach a “Full” adjacency state, the routing tables between neighbors have been fully propagated. Once neighbors reach this point, only incremental routing table updates will be sent.

Essential OSPF Troubleshooting Commands

The following commands in this section are common show and debug tools for troubleshooting OSPF network problems. Please make sure that you fully understand the command and what information it can provide.

Show ip route ospf

The show ip route ospf command is used to display the routing table for only OSPF routes. If you have multiple routing protocols, this command helps to filter out everything except for OSPF learned routes as shown here:

```
RouterB#show ip route ospf
      172.16.0.0/24 is variably subnetted, 3 subnets, 2 masks
O       172.16.0.0/24 is a summary, 00:00:29, Null0
      172.16.0.0/24 is variably subnetted, 3 subnets, 2 masks
O       201.0.0.0/24 is a summary, 00:00:17, Null0
```

Show ip ospf interface

The show ip ospf interface command details OSPF enabled interfaces. This command specifies important information such as:

- Area
- Router ID
- Network type
- DR IP address
- BDR IP address
- Neighbor count
- Timer settings

Here is the show command output:

```
Router#show ip ospf interface FastEthernet 0/1
FastEthernet 0/1 is up, line protocol is up
Internet Address 192.168.250.202, Mask 255.255.255.0, Area 0.0.0.0
AS 201, Router ID 192.77.99.1, Network Type BROADCAST, Cost: 10
Transmit Delay is 1 sec, State OTHER, Priority 1
Designated Router id 192.168.250.10, Interface address 192.168.250.10
Backup Designated router id 192.168.254.28, Interface addr
192.168.250.28
Timer intervals configured, Hello 10, Dead 60, Wait 40, Retransmit 5
Hello due in 0:00:03
Neighbor Count is 8, Adjacent neighbor count is 2
  Adjacent with neighbor 192.168.250.28 (Backup Designated Router)
  Adjacent with neighbor 192.168.250.10 (Designated Router)
```


Show ip ospf neighbor

The show ip ospf neighbor command displays OSPF neighbor information on an interface-by-interface basis. This is a quick way to check your neighbors to see which ones are adjacent and which are DR/BDR if the interfaces are of OSPF type Broadcast such as Ethernet connections as shown here:

```
Router#show ip ospf neighbor

Neighbor ID      Pri   State           Dead Time   Address      Interface
192.168.10.1    1     FULL/DR         0:00:36    10.0.0.1    FastEthernet0/1
172.1.4.1       1     FULL/DROTHER    0:00:33    172.1.4.1   FastEthernet0/2
```

Show ip ospf database

When a network has multiple areas configured on it, an excellent way to view the router IDs configured in each area is to use the show ip ospf database command. In addition, you can see the following information:

Field	Description
Link ID	Router ID IP address
ADV Router	Advertising Router IP address
Age	Age of the neighbor connection
Link count	Number of interfaces detected

Figure 28: Information from the show ip ospf database Command

Here is an example of the output of this show command:

```
Router#show ip ospf database
      OSPF Router with ID (203.250.15.67) (Process ID 10)
          Router Link States (Area 1)
Link ID      ADV Router    Age      Seq#          Checksum Link count
203.250.15.67 203.250.15.67 48       0x80000008  0xB112    2
203.250.16.130 203.250.16.130 212      0x80000006  0x3F44    2

          Summary Net Link States (Area 1)
Link ID      ADV Router    Age      Seq#          Checksum
203.250.13.41 203.250.15.67 602      0x80000002  0x90AA
203.250.15.64 203.250.15.67 620      0x800000E9  0x3E3C
203.250.15.192 203.250.15.67 638      0x800000E5  0xA54E

          Router Link States (Area 0)
Link ID      ADV Router    Age      Seq#          Checksum Link count
203.250.13.41 203.250.13.41 179      0x80000029  0x9ADA    3
203.250.15.67 203.250.15.67 675      0x800001E2  0xDD23    1

          Net Link States (Area 0)
Link ID      ADV Router    Age      Seq#          Checksum
203.250.15.68 203.250.13.41 334      0x80000001  0xB6B5
```

```

Summary Net Link States (Area 0)
Link ID      ADV Router   Age   Seq#      Checksum
203.250.15.0 203.250.15.67 792   0x80000002 0xAEBD

Summary ASB Link States (Area 0)
Link ID      ADV Router   Age   Seq#      Checksum
203.250.16.130 203.250.15.67 579   0x80000001 0xF9AF

AS External Link States
Link ID      ADV Router   Age   Seq#      Checksum Tag
0.0.0.0      203.250.16.130 1787  0x80000001 0x98CE 10
203.250.16.128 203.250.16.130 5     0x80000002 0x93C4 0

```

Show ip ospf statistics

The show ip ospf statistics command shows information for each OSPF SPF calculation that has been run since the process started as shown here:

```

Router#show ip ospf statistics
OSPF process ID 200
-----
Area 0: SPF algorithm executed 10 times
Area 200: SPF algorithm executed 8 times
Summary OSPF SPF statistic
SPF calculation time
Delta
T          Intra   D-Intra  Summ    D-Summ  Ext     D-Ext   Total   Reason
08:17:16  0        0        0       0       0       0       0       R,
08:16:47  0        0        0       0       0       0       0       R, N,
08:16:37  0        0        0       0       0       0       0       R, X
00:04:40  208     40      208    44     220    0       720    R, N,
SN, X
00:03:15  0        112     4       108    8       96     328    R, N,
SN, X
00:02:55  164     40      176    44     188    0       612    R, N,
SN, X
00:01:49  0        4        4       0       4       4       16     R, N,
SN, X
00:01:48  0        0        4       0       4       0       12     R, N,
SN, SA, X
00:01:43  0        0        4       0       4       0       8      R,
00:00:53  164     40      176    44     188    0       612    R, N,
SN, X

```

This table describes the fields shown in the command output that can be used to troubleshoot the OSPF SPF algorithm:

Field	Description
Delta T	The amount of time (ms) when the SPF began its calculation in relation to atomic time.
Intra	The time spent (ms) to process intra-area routes and place them into the IP routing table
D-Intra	The time spent (ms) to remove intra-area routes from the IP routing table
Summ	The time spent (ms) to process inter-area routes and place them into the IP routing table
D-Summ	The time spent (ms) to remove inter-area routes from the IP routing table
Ext	The time spent (ms) to process external and NSSA routes and place them into the IP routing table
D-Ext	The time spent (ms) to remove external and NSSA routes from the IP routing table
Total	Total time (ms) for the OSPF SPF algorithm to complete
Reason	Explanation why the SPF was executed: N—A change in a network LSA (type 2) R—A change in a router LSA (type 1) SA—A change in a Summary ASBR LSA (SA) SN—A change in a Summary Network LSA (SN) X—A change in an External Type-7 LSA (X7)

Figure 29: Fields Shown in the show ip ospf statistics Command

As you can see from the command output and description table, not only can you verify the times the SPF takes over time, but you can also see exactly what caused the router to recalculate routes. This information can prove vital to finding the root cause of an unstable OSPF network.

Debug ip ospf monitor

The debug ip ospf monitor command is useful to see changes in an OSPF network topology as they occur. The following example shows that there was an OSPF SPF recalculation on the router with the router ID of 10.1.1.1 that caused the local database to resynchronize:

```
Router1#debug ip ospf monitor
OSPF: Schedule SPF in area 0.0.0.0
      Change in LS ID 10.1.1.1, LSA type R,
OSPF: schedule SPF: spf_time 1621348044ms wait_interval 10s
```

Debug ip ospf packet

The debug ip ospf packet command lets an administrator view real-time information about each Open Shortest Path First (OSPF) packet received on the local router from neighboring routers as shown in this example:

```
Router#debug ip ospf packet
OSPF: rcv. v:2 t:1 l:48 rid:133.133.133.133
aid:0.0.0.0 chk:C582 aut:0 auk: from FastEthernet1/0
OSPF: rcv. v:2 t:1 l:48 rid:111.111.111.111
aid:0.0.0.0 chk:F3B2 aut:0 auk: from FastEthernet0/0
```

This example shows two ospf type 1 packets, which are hello packets originating from neighbors on two different FastEthernet interfaces. The other important information found in the debug output is as follows:

Field	Description
t:	OSPF Packet Type: 1 – Hello 2 – Data 3 – Link State Request 4 – Link State Update 5 – Link State Ack
rid:	Router ID
aid:	Area ID
aut:	OSPF Authentication: 0 – None 1 – Simple Password 2 – MD5 Hashed Password
seq:	Packet Sequence Number

Figure 30: Fields in the debug ip ospf packet Command

Debug ip ospf adj

A great way to see the OSPF adjacency process is to turn on OSPF adjacency debugging by issuing the debug ip ospf adj command. Once enabled, you can see log messages that show the entire OSPF adjacency process move from the Init to the full states and everything in between. You can also see the election of DR and BDR on broadcast networks such as Ethernet as shown in this output example:

```
Router#debug ip ospf adj
OSPF: Cannot see ourself in hello from 133.133.133.133 on
FastEthernet1/0, state INIT
OSPF: Neighbor change Event on interface FastEthernet1/0
OSPF: DR/BDR election on FastEthernet1/0
OSPF: Elect BDR 2.2.2.2
OSPF: Elect DR 2.2.2.2
OSPF: Elect BDR 0.0.0.0
OSPF: Elect DR 2.2.2.2
DR: 2.2.2.2 (Id) BDR: none
OSPF: Remember old DR 133.133.133.133 (id)
OSPF: Reset old DR on FastEthernet1/0
OSPF: Build router LSA for area 0, router ID 2.2.2.2, seq 0x8000000B
```

Another popular way to use this debug output is to discover OSPF authentication mismatches as shown in this example output:

```
Router#debug ip ospf adj
OSPF adjacency events debugging is on
00:54:04: OSPF: Rcv pkt from 172.12.23.2, Ethernet0 : Mismatch
Authentication type. Input packet specified type 2, we use type 1
```

As you can see, the local router is using type 1 OSPF authentication which is clear text password while the neighbor router is using a type 2 authentication method which is a password with an encrypted MD5 hash.

Debug ip ospf events

Any major OSPF event that occurs locally to the OSPF process can be viewed using the debug ip ospf events command. The possible events include:

- New and deleted adjacencies
- OSPF flooding
- DR/BDR selection
- SPF calculations

The debug output below shows how to use this debug output to discover an OSPF timer mismatch. This mismatch happens to be the dead interval timer:

```
Router#debug ip ospf events
OSPF:hello with invalid timers on interface Ethernet0
hello interval received 10 configured 10
net mask received 255.255.255.0 configured 255.255.255.0
dead interval received 40 configured 30
```

Show ip ospf virtual-links

The CCNP ROUTE exam details how to configure OSPF virtual-links when needed on a network. Essentially, an OSPF virtual link is a tunnel for LSA information when an OSPF area cannot directly connect to area 0. While it is advisable to avoid OSPF virtual-links in production, sometimes it is necessary. If a network administrator suspects that an OSPF virtual-link is not operating properly, the best command to verify functionality is to use the show ip ospf virtual-links command. This will show information such as:

- Current OSPF virtual-link status (up or down)
- OSPF area that is being used to transit to area 0
- OSPF circuit State
- Timer settings
- Adjacency State

Here is an example of this show command:

```
Router#show ip ospf virtual-links
  Virtual Link OSPF_VL0 to router 192.168.10.2 is up
  Run as demand circuit
  DoNotAge LSA allowed.
  Transit area 10, via interface Serial0.1, Cost of using 64
  Transmit Delay is 1 sec, State POINT_TO_POINT,
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
  Hello due in 00:00:09
  Adjacency State FULL (Hello suppressed)
  Index 2/5, retransmission queue length 0, number of retransmission 1
  First 0x0(0)/0x0(0) Next 0x0(0)/0x0(0)
  Last retransmission scan length is 1, maximum is 1
  Last retransmission scan time is 0 msec, maximum is 0 msec
```

Troubleshooting BGP

The Border Gateway Protocol (BGP) is the routing protocol used on the Internet. The protocol maintains a table of IP networks that specify how to reach public networks based on autonomous system (AS) number. The following sections will cover what needs to be known about BGP to troubleshoot it as well as useful show and debug commands that can be used.

IGP vs. EGP

So far, the routing protocols that have been discussed are considered to be interior gateway protocols (IGP). BGP is currently the one and only exterior gateway protocol (EGP). The main difference between the two has to do with the ownership to remote networks the protocol discovers. IGPs handle routing within a single autonomous system (AS) while EGPs handle routing between autonomous systems.

BGP Operation

This section covers how BGP operates and inserts routes into the routing table when configured on a network:

- **Incoming Route Information** – BGP routers receive routes from BGP peers. These peers do not have to be directly connected which is different from OSPF and EIGRP.
- **BGP Tables:**
 - ▶ **BGP Neighbor table** – contains neighbor status information.
 - ▶ **BGP table** – contains network subnets learned from BGP peers. The best path metric determined by BGP using various factors (weight, local preference, etc.) will be placed into the routing table.
 - ▶ **Routing Table** – the table used by the router for path selection to remote networks.
 - **Injecting and Redistributing Routes** – routes can be inserted into the BGP table either through route injection or through route redistribution methods. Keep in mind that the route must be in the IP routing table for it to be injected/redistributed.
 - **Route Installation** – BGP chooses the best route to a network based on multiple configurable factors called BGP metrics. The best route chosen to a remote network is then inserted into the routing table.
 - **Outgoing Route Information** – routes that have been determined to be the best routes to their network subnets are then advertised to the local routers BGP peers.

BGP Metric

When BGP has more than one path to a remote network, it must choose one (candidate route) to be placed into the IP routing table. To accomplish this goal, BGP looks at the following metric parameters to determine the optimal path. These metrics are preferred in order. As soon as one path has a metric better than the others, that route is automatically chosen. If the metrics are identical, BGP continues to compare the remaining metrics until there is a difference in one of them.

Here are the metric parameters, in the order BGP reviews them:

1. Weight
2. Local Preference
3. Multi-exit Discriminator
4. Origin
5. AS Path
6. Next Hop
7. Community

Troubleshooting BGP Neighbors

BGP neighbor relationships are manually configured by a network administrator. The protocol communicates to potential peers on TCP port 179. Once a BGP peer has been setup, the peers will share routing information and eventually reach and Established state.

Reasons Why A BGP Peering would Fail

There are several reasons why BGP peers might not reach an established state. These reasons include.

- Factors between peers such as the manual configuration of AS numbers must be identical on both peer routers. The peer source IP addresses must also match the network statements on both sides.
- Lack of IP connectivity between peer routers. Even though there might be a connection at layer 1 and 2, BGP runs over TCP so administrators must insure proper IP connectivity and to make sure nothing is blocking TCP port 179 which BGP runs on.

Essential BGP Troubleshooting Commands

The following commands should be known by the TSHOOT candidate for troubleshooting problems associated with BGP.

Show ip bgp

The show ip bgp command displays entries in the BGP routing table for one network prefix or the entire BGP routing table. This command presents a network administrator with the following information useful for troubleshooting:

- **Local router ID** - IP address of the Router from a BGP perspective
- **Status code** - Status of the network entry in the BGP table. Possible states include:
 - ▶ **s** - suppressed table entry
 - ▶ ***** - valid table entry
 - ▶ **>** - Best table entry for a specific IP network
 - ▶ **i** - Learned table entry from iBGP

- **Origin code** - The origin of the network entry. Possible origins include:
 - ▶ **i** - Originated from an IGP
 - ▶ **e** - Originated from an EGP (BGP)
 - ▶ **?** - Origin is unknown. Commonly when routes are redistributed from an IGP into BGP
- **Network** - BGP learned network
- **Next Hop** - IP address of the next hop to a network
- **BGP attributes** - such as Metric, LocPrf, Weight and Path

Below is an example of the show command output:

```
Router# show ip bgp
BGP table version is 716977, local router ID is 192.168.32.1
Status codes: s suppressed, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete
Network          Next Hop          Metric LocPrf Weight Path
* i3.0.0.0        192.168.22.1      0      100    0      1800 1239 ?
*>i              192.168.16.1      0      100    0      1800 1239 ?
* i6.0.0.0        192.168.22.1      0      100    0      1800 690 568 ?
*>i              192.168.16.1      0      100    0      1800 690 568 ?
* i7.0.0.0        192.168.22.1      0      100    0      1800 701 35 ?
*>i              192.168.16.1      0      100    0      1800 701 35 ?
*                  172.16.72.24      0      1878   704    701 35 ?
* i8.0.0.0        192.168.22.1      0      100    0      1800 690 560 ?
*>i              192.168.16.1      0      100    0      1800 690 560 ?
*                  172.16.72.24      0      1878   704    701 560 ?
* i13.0.0.0       192.168.22.1      0      100    0      1800 690 200 ?
*>i              192.168.16.1      0      100    0      1800 690 200 ?
*                  172.16.72.24      0      1878   704    701 200 ?
* i15.0.0.0       192.168.22.1      0      100    0      1800 174 ?
*>i              192.168.16.1      0      100    0      1800 174 ?
* i16.0.0.0       192.168.22.1      0      100    0      1800 701 i
*>i              192.168.16.1      0      100    0      1800 701 i
*                  172.16.72.24      0      1878   704    701 i
```


Show ip bgp summary

The show ip bgp summary command is useful to quickly see all of the local routers peerings as well as to verify the current state and how long the peering has been established as shown in the example here:

```

Router# show ip bgp summary
BGP table version is 717029, main routing table version 717029
19073 network entries (37544 paths) using 3542756 bytes of memory
691 BGP path attribute entries using 57200 bytes of memory
Neighbor      V   AS   MsgRcvd MsgSent  TblVer  InQ  OutQ  Up/Down  State/PfxRcd
192.168.16.1  4   1755  32642   2973    717029  0    0     1:27:11
192.168.17.1  4   1755  4790    2973    717029  0    0     1:27:51
192.168.18.1  4   1755  7722    3024    717029  0    0     1:28:13
192.168.19.1  4   1755    0        0        0    0    0     2d02      Active
192.168.20.1  4   1755  3673    3049    717029  0    0     2:50:10  Idle (PfxRcd)
192.168.21.1  4   1755  3741    3048    717029  0    0     12:24:43
192.168.22.1  4   1755  33129   3051    717029  0    0     12:24:48
192.168.23.1  4   1755    0        0        0    0    0     2d02      Active
192.168.24.1  4   1755    0        0        0    0    0     2d02      Active
192.168.25.1  4   1755    0        0        0    0    0     2d02      Active
192.168.26.1  4   1755    0        0        0    0    0     2d02      Active
192.168.27.1  4   1755  4269    3049    717029  0    0     12:39:33
192.168.28.1  4   1755  3037    3050    717029  0    0     2:08:15
172.16.72.24  4   1878  11635   13300   717028  0    0     0:50:39
172.16.72.36  4   1001    0        0        0    0    0     never     Idle (Admin)
Router#

```

Figure 31: Output from the <show ip bgp summary> Command

Show ip bgp neighbors <ip_address>

To view detailed information regarding a specific BGP peer, issue the show ip bgp neighbors command. The output lists information about peers including:

- BGP neighbor IP address and AS number
- BGP state
- Amount of time the BGP connection has been established at TCP 179
- Timers such as:
 - ▶ Hold time
 - ▶ Keepalive time
- Number of error messages (notifications) sent to this peer
- Number of times a TCP connection has been established. If this counter increments often, then there is likely a layer 1-3 issue between the peer

Below is an example of this show command output for a BGP peer at 172.16.10.1:

```
Router# show ip bgp neighbors 172.16.10.1

BGP neighbor is 172.16.10.1, remote AS 10, external link
  Index 1, Offset 0, Mask 0x2
  Inbound soft reconfiguration allowed
  BGP version 4, remote router ID 172.16.10.1
  BGP state = Established, table version = 27, up for 00:06:12
  Last read 00:00:12, hold time is 180, keepalive interval is 60
seconds
  Minimum time between advertisement runs is 30 seconds
  Received 19 messages, 0 notifications, 0 in queue
  Sent 17 messages, 0 notifications, 0 in queue
  Inbound path policy configured
  Route map for incoming advertisements is testing
  Connections established 2; dropped 1
Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Local host: 172.16.232.181, Local port: 11002
Foreign host: 172.16.10.1, Foreign port: 179

Enqueued packets for retransmit: 0, input: 0, saved: 0

Event Timers (current time is 0x530C294):
Timer           Starts      Wakeups      Next
Retrans         12          0            0x0
TimeWait        0           0            0x0
AckHold         12          10           0x0
SendWnd         0           0            0x0
KeepAlive       0           0            0x0
GiveUp          0           0            0x0
PmtuAger        0           0            0x0

iss: 133981889  snduna: 133982166  sndnxt: 133982166  sndwnd:
16108
irs: 3317025518  rcvnxt: 3317025810  rcvwnd: 16093  delrcvwnd:
291

SRTT: 441 ms, RTTO: 2784 ms, RTV: 951 ms, KRTT: 0 ms
minRTT: 0 ms, maxRTT: 300 ms, ACK hold: 300 ms
Flags: higher precedence, nagle

Datagrams (max data segment is 1460 bytes):
Rcvd: 15 (out of order: 0), with data: 12, total data bytes: 291
Sent: 23 (retransmit: 0), with data: 11, total data bytes: 276
```

Debug ip bgp events

A way to see important BGP events such as BGP status changes is to use the debug ip bgp events command. The example below shows peers moving from an Established state to idle and then back to Active, OpenSent, OpenConfirm and finally back to an Established state:

```
Router#debug ip bgp events
BGP events debugging is on

BGP: reset all neighbors due to User reset
BGP: 172.16.1.2 went from Established to Idle
BGP: 172.16.255.6 went from Established to Idle
BGP: 172.16.255.14 went from Established to Idle
BGP: 172.16.1.2 went from Idle to Active
BGP: 172.16.255.6 went from Idle to Active
BGP: 172.16.255.14 went from Idle to Active
BGP: 172.16.255.6 went from Active to OpenSent
BGP: 172.16.255.6 went from OpenSent to OpenConfirm
BGP: 172.16.255.6 went from OpenConfirm to Established
BGP: 172.16.255.6 computing updates, neighbor version 0, table version
1,
  starting at 0.0.0.0
BGP: 172.16.255.6 update run completed, ran for 0ms, neighbor version 0,
  start version 1, throttled to 1, check point net 0.0.0.0
BGP: 172.16.255.14 went from Active to OpenSent
BGP: 172.16.255.14 went from OpenSent to OpenConfirm
BGP: 172.16.255.14 went from OpenConfirm to Established
BGP: 172.16.255.14 computing updates, neighbor version 0, table version
1,
  starting at 0.0.0.0
BGP: 172.16.255.14 update run completed, ran for 0ms, neighbor version
0,
  start version 1, throttled to 1, check point net 0.0.0.0
BGP: 172.16.1.2 went from Active to OpenSent
BGP: 172.16.1.2 went from OpenSent to OpenConfirm
BGP: 172.16.1.2 went from OpenConfirm to Established
BGP: 172.16.1.2 computing updates, neighbor version 0, table version 1,
starting at 0.0.0.0
BGP: 172.16.1.2 update run completed, ran for 0ms, neighbor version 0,
  start version 1, throttled to 1, check point net 0.0.0.0
BGP: 172.16.255.6 computing updates, neighbor version 1, table version
9,
  starting at 0.0.0.0
BGP: 172.16.255.6 update run completed, ran for 0ms, neighbor version 1,
  start version 9, throttled to 9, check point net 0.0.0.0
BGP: scanning routing tables
BGP: scanning routing tables
BGP: scanning routing tables
```

Debug ip bgp updates

The debug ip bgp updates command displays in real-time the received networks and attributes from a configured BGP neighbor.

```
Router#debug ip bgp updates
BGP(0): 172.16.10.1 computing updates, afi 0, neighbor version 0, table
version 5, starting at 0.0.0.0
BGP(0): 172.16.10.1 send UPDATE (format) 192.168.9.0/24, n ext
192.168.1.2, metric 0, path
BGP(0): 172.16.10.1 1 updates enqueued (average=52, maximum=52)
BGP(0): 172.16.10.1 update run completed, afi 0, ran for 0 ms, neighbor
version 0, start version 5, throttled to 5
BGP: 172.16.10.1 initial update completed
BGP(0): 172.16.10.1 rcvd UPDATE w/ attr: nexthop 192.168.1 .1, origin
i, metric 0, path 200 100 ISP-C#
BGP(0): 172.16.10.1 rcvd 192.168.4.0/24
BGP(0): Revise route installing 192.168.4.0/24 -> 192.168. 1.1 to main
IP table
```

Troubleshooting Route Redistribution

Route redistribution is the process of injecting routes from different dynamic routing protocols or statically configured routes. The method of redistribution lets routing protocols share routing information between each other. This section covers reasons to redistribute routes, common redistribution traits and commands used to troubleshoot redistributed routes.

Why an Administrator would Redistribute Routes

At some point in time, network administrators will come across a reason why it is necessary to redistribute routes from one protocol into another. Some of the more common reasons are:

- **Transitioning to a more advanced routing protocol** – there will be a transition period where you are running the old and new routing protocols, these will have to work together until you get all routes moved over to the advanced routing protocol.
- **Merger between one or more business** – this is a common practice when organizations merge and they happen to be running different routing protocols. Typically, network administrators will choose a single IGP and transition the other protocol so the entire organization runs a single one.
- **Network has different network administration groups** – some organizations are very separated even to the point where there are different IT staff and network requirements. Because of this, a single business might be separated by defined boundaries that run different protocols, which will need to be redistributed into each other.

The Boundary Router

A router that is responsible for the redistribution of one routing protocol into another is called a **boundary router**. The router sits on the edge and must run both routing protocols as shown in this figure:

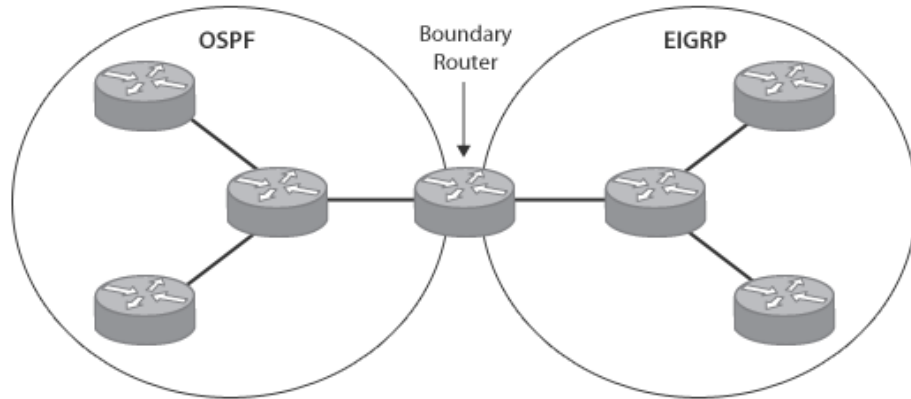


Figure 32: Boundary Routers Explained

Injecting Metrics Into Redistributed Routes

As we are now well aware, routing protocols use metrics to help choose the best paths to a network. The routing protocol uses various different methods to choose the best path to be placed into the routing table. When redistributing routes, the injected routes need to have a **seed metric** associated to it that the specific routing protocol understands. Network administrators can modify the seed metric by modifying either the default metric for all redistributed routes or by modifying the metric for a specific neighbor. If a seed metric is not defined, the default metric is used. Keep in mind however that some routing protocols (RIP and EIGRP) have a default seed metric that is “unreachable” so these protocols must be manually changed from the default seed metric.

Essential Route Redistribution Troubleshooting Commands

The following commands and descriptions are necessary tools for troubleshooting redistribution problems on boundary routers.

Show ip route <routing_protocol>

If the router is redistributing either static, connected or dynamic routes learned through other routing protocols, the show ip route <routing_protocol> command can be used to verify the redistributed routes are being placed into the routing table. This example shows that routes are being redistributed into OSPF. The “O E2” portion of the output means that OSPF “O” is putting this into the routing table but OSPF learned about this route because it was redistributed from a different source (E2):

```
R2#show ip route ospf
5.0.0.0/24 is subnetted, 3 subnets
O E2 5.1.1.0 [110/20] via 172.12.123.1, 00:00:08, Serial0.123
O E2 5.2.1.0 [110/20] via 172.12.123.1, 00:00:08, Serial0.123
O E2 5.3.1.0 [110/20] via 172.12.123.1, 00:00:08, Serial0.123
10.0.0.0/24 is subnetted, 1 subnets
O E2 10.1.1.0 [110/20] via 172.12.123.1, 00:00:08, Serial0.123
```

OSPF has two different types of redistributed external routes being E1 or E2. A type E1 route has a metric that is the sum of the internal OSPF cost and the external redistributed cost. A type E2 route has a metric equal only to the redistributed cost. By default, OSPF will use type E2 for all redistributed routes.

Show ip route profile

A Cisco IOS routing protocol feature used to help keep track of protocol-independent route flaps, network failures and restorations, many network administrators enable the ip route profile command, which begins tracking these statistics. To display routing table change statistics, use the show ip route profile command shown here:

```
Router#show ip route profile
```

Change/ interval	Fwd-path change	Prefix add	Nexthop Change	Pathcount Change	Prefix refresh
0	87	87	89	89	89
1	0	0	0	0	0
2	0	0	0	0	0
3	0	0	0	0	0
4	0	0	0	0	0
5	0	0	0	0	0
10	0	0	0	0	0
15	0	0	0	0	0
20	2	2	0	0	0
25	0	0	0	0	0

Below are descriptions of the most important fields of this show command output:

- **Fwd-path Change** – the number of times the forwarding path changes. This is the accumulation of prefix-add, next-hop change, and pathcount change statistics.
- **Prefix add** – a counter showing that a new prefix was added to the routing table.
- **Nexthop Change** – a counter showing that a prefix is not added or removed, but the next hop changes. This statistic is only seen with recursive routes that are installed in the routing table.
- **Pathcount Change** – a counter showing that the number of paths in the routing table has changed. This counter increments as a result of an increase in the number of paths for an Interior Gateway Protocol (IGP) prefix in the routing table.
- **Prefix refresh** – a counter showing that standard routing table maintenance changes.

Troubleshooting DHCP

DHCP is a service that can be configured on Cisco layer 3 devices such as routers and multilayer switches. In addition, Cisco devices can also forward DHCP requests to other subnets and become DHCP clients themselves. This section covers what DHCP is and how to troubleshoot it using various CLI commands.

Purpose of DHCP

Dynamic Host Control Protocol (DHCP) lets end devices such as computers and IP phones request information required for them to function on an IP network. The main benefit is to reduce the work necessary to maintain an IP network. The alternative to DHCP is to statically assign network information on each end device. When changes are made to the IP network, network administrators would have to physically find and modify the devices that were affected by the network change.

DHCP Operation

Cisco devices such as Layer 3 switches and routers can use DHCP in the following ways:

- **DHCP Client** – Layer 3 interfaces such as router Ethernet interfaces can obtain their IP address and other information from an external DHCP server. To do this, the network administrator configures the following on the interface:

```
Router(config-if)#ip address dhcp
```

- **DHCP Relay** – Networks with multiple subnets often enable DHCP on many of the IP subnets in use. DHCP client devices send out a broadcast on the subnet they are attached to in order to request an IP address from the DHCP server. As we know, each VLAN on a network is a separate broadcast domain. This means that your network would have to have a DHCP server on every VLAN. If the DHCP server is located on a different VLAN, the server would never get the DHCP request because the broadcast cannot leave the VLAN. As a network grows, this method becomes difficult to manage. The ideal situation is to have a single source where DHCP information is maintained. The way around this is to configure a **DHCP helper-address** on all IP subnets where DHCP is being used. The DHCP helper-address command identifies DHCP broadcast requests from client. It sends the request to the DHCP server in a directed broadcast on the behalf of the client. The server responds to the client in a unicast message and the client then sends the DHCP request as a unicast to the DHCP server. The command to configure DHCP relay on an interface is:

```
Router(config-if)#ip helper-address <IP_address>
```

- **DHCP Server** – A layer 3 switch or router can also be used as a DHCP server. DHCP is a separate service that must be enabled using the following command:

```
Router(config)#service dhcp
```

Once the DHCP service is enabled, a network administrator can configure DHCP information such as:

- DHCP IP address pools
- Excluded IP addresses
- Network definitions – IP subnets with prefix length to define DHCP subnets
- Default gateway
- Domain name
- DNS server(s)
- NetBIOS server(s)
- DHCP lease times

DHCP Messages

There is a specific DHCP client/server communication process that network administrators should be familiar with. Each step of the DHCP process uses unique messages that can help the administrator narrow down where a potential problem is. The DHCP messages are in the following table:

Message	Source -> Destination	Description
DHCPDISCOVER	Client -> Server	Broadcast attempt to find DHCP server using UDP 67
DHCPOFFER	Server -> Client	Response from server receiving the DHCPDISCOVER request. Response uses UDP 68
DHCPREQUEST	Client -> Server	Request for DHCP configuration parameters from a client to a specific DHCP server
DHCPDECLINE	Client -> Server	Message sent to the server to inform the server that an IP address is already in use on the IP subnet
DHCPACK	Server -> Client	Network configuration information being sent from the server to the client
DHCPNAK	Server -> Client	A DHCP request deny from the server to the client
DHCPLEASE	Client -> Server	A release of network information based on the DHCP lease time configured. The IP address can then be placed back in the DHCP pool to be handed out to another client
DHCPINFORM	Client -> Server	Request for DHCP configuration parameters from and access-server requesting IP information for a remote client

Figure 33: DHCP Messages

Common DHCP Problems

The following are common DHCP problems that network administrators will need to understand how to troubleshoot.

Misconfiguration

Network administrators are human and will always be a major source of network problems. Many times a misconfiguration of the DHCP server such as the wrong default-network or subnet can cause a DHCP client to receive an address but since it is the incorrect IP settings for a subnet, can cause the device to not function on the network.

Router Not Forwarding DHCP Broadcasts

A common problem for new engineers is to understand that clients request DHCP information from DHCP servers in the form of a broadcast. Unless the IP subnet has a locally configured DHCP server or an IP helper-address, the broadcast requests will never reach the DHCP server.

Full DHCP Pool

As more and more clients join a specific IP subnet, make sure you keep a close eye on the DHCP IP address pool that has been configured. Also, you can modify DHCP lease information to release IP addresses more quickly if a particular subnet (such as a Wi-Fi hotspot or other public network) has many users that appear and disappear often. If a DHCP pool is full, any new DHCP request will be rejected with a DHCPNAK message.

DHCP Troubleshooting Commands

Now that you have an understanding of what Cisco devices can do with DHCP and the potential pitfalls associated with DHCP, this next section covers show, clear and debug commands used to troubleshoot and verify the health of DHCP services.

Show and clear ip dhcp binding

To view current DHCP leased addresses and the MAC address of the client they are leased to, use the show ip dhcp binding command. This command also details lease expiration information as well as the type of DHCP lease. The following output shows the command being used to view a specific IP address within a DHCP pool:

```
Router#show ip dhcp binding 172.16.1.11
IP address      Hardware address  Lease expiration  Type
172.16.1.11    00a0.9802.32de    Feb 01 2010 12:00 AM  Automatic
```

If an administrator wants to clear that specific lease out before the lease expiration, they can issue a clear ip dhcp binding and include the ip address as shown here:

```
Router#clear ip dhcp binding 172.16.1.11
```

Show and clear ip dhcp conflict

Often times when using DHCP, a conflict will be detected either by the client or by the server when an IP address is discovered to already be in use. To see a listing of recently discovered conflicts, use the show ip dhcp conflict command as seen here:

```
Router#show ip dhcp conflict
IP address      Detection Method  Detection time
172.16.1.32     Ping              Feb 16 2010 12:28 PM
172.16.1.64     Gratuitous ARP    Feb 23 2010 08:12 AM
```

The command output lists the IP address found to be in conflict and when it was detected. In addition, the detection method lets an administrator know if the conflict was detected by the server (uses ping as detection method) or the client (uses gratuitous ARP as detection method).

To clear all of the listed conflicts in the DHCP database, you can issue the clear ip dhcp conflict * command. The * means that all conflicts will be purged from the conflict table:

```
Router#clear ip dhcp conflict *
```

Debug ip dhcp server events

To see real time DHCP server service events, enable the debug ip dhcp server events command line option. The debug output will display a lot of DHCP server information, such as:

- DHCP leases
- Client/Server communication
- IP address allocations
- IP to MAC mappings

Here is an example of the output seen using the debug command:

```
Router#debug ip dhcp server events
DHCPD: checking for expired leases.
DHCPD: Sending notification of TERMINATION:
DHCPD: address 10.2.1.7 mask 255.255.255.0
DHCPD: reason flags: RELEASE
DHCPD: htype 1 chaddr 0004.dd69.fd01
DHCPD: lease time remaining (secs) = 315841
DHCPD: returned 10.2.1.7 to address pool vlan10.
DHCPD: Sending notification of DISCOVER:
DHCPD: htype 1 chaddr 0004.dd69.fd01
DHCPD: remote id 020a00000a01040500000000
DHCPD: circuit id 00000000
DHCPD: Seeing if there is an internally specified pool class:
DHCPD: htype 1 chaddr 0004.dd69.fd01
DHCPD: remote id 020a00000a01040500000000
DHCPD: circuit id 00000000
DHCPD: Adding binding to radix tree (10.2.1.8)
DHCPD: Adding binding to hash tree
DHCPD: assigned IP address 10.2.1.8 to client 0063.6973.636f.2d30.3030
.342e.6464.3639.2e66.6430.312d.4661.302f.30.
DHCPD: Sending notification of ASSIGNMENT:
DHCPD: address 10.2.1.8 mask 255.255.255.0
DHCPD: htype 1 chaddr 0004.dd69.fd01
DHCPD: lease time remaining (secs) = 432000
```

Debug ip dhcp server packet

The debug ip dhcp server packet command displays client/server DHCP messages such as DHCPDISCOVER and DHCPOFFER messages taken from the snapshot of debug output shown here:

```
Router#debug ip dhcp server packet
DHCPD:DHCPDISCOVER received from client 0063.6973.636f.2d30.3030.312e.
3432.6339.2e65.6337.352d.4574.31 on interface FastEthernet0.
DHCPD:assigned IP address 192.168.33.81 to client
0063.6973.636f.2d30.3030.312e.3432.6339.2e65.6337.352d.4574.31.
DHCPD:Sending DHCPOFFER to client
0063.6973.636f.2d30.3030.312e.3432.6339.2e65.6337.35
2d.4574.31(192.168.33.81)
```

Troubleshooting NAT

Whether you work on large or small networks, at some point in your network administration career, you will deal with a network that uses network address translation (NAT). This section covers the basics of troubleshooting this IP service.

Purpose of NAT

Network Address Translation (NAT) is a method that translates one or more “inside” IP address into one or more “global” IP addresses. It is common to map multiple RFC 1918 IP private addresses to one or several publicly routed IP address on the Internet. This process conceives public IP addresses because you do not need a one-to-one public IP address to PC ratio. NAT also adds at least a small layer of security because each outgoing or incoming request must go through a translation choke point where security can be enforced.

Four Types of NAT

Configuring NAT on a Cisco router, you can setup the following different NAT types:

- Statically map one inside address to one global address
- Dynamically map inside addresses to a rotating pool of global address
- Dynamically map inside addresses and TCP port to one or a pool of global addresses
- Dynamically map global addresses to a rotating pool of inside addresses

NAT Topology

The overall topology of NAT can be broken down into four separate segments within a topology. Your specific NAT setup may or may not use all of these as shown in the following NAT topology figure:

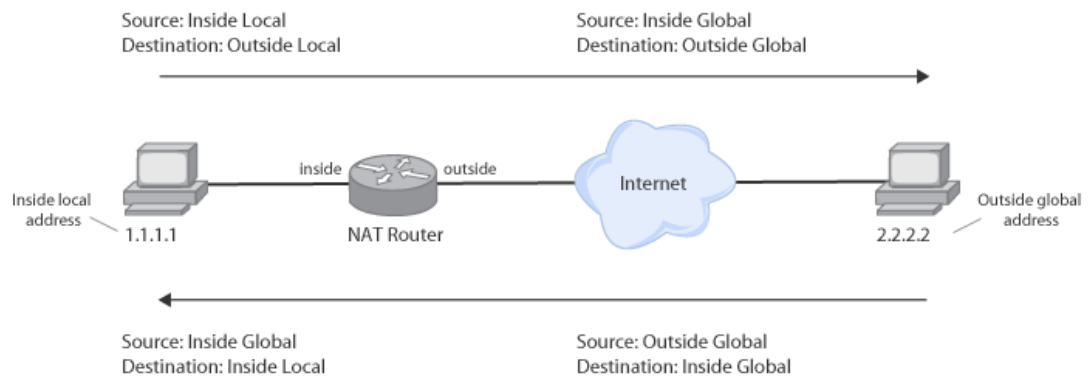


Figure 34: The NAT Topology

Here are the definitions of each of the NAT topology sections:

- **Inside Local** – IP address that is inside the physical network and it is local to the internal devices.
- **Inside Global** – IP address that your router is using to connect to the public network.
- **Outside Local** – for static mappings from the inside to the outside.
- **Outside Global** – public IP address that you're connecting to.

Common NAT Problems

The TSHOOT exam will cover how to identify and troubleshoot common NAT problems. These problems include situations such as the following:

- **Misconfiguration** – many NAT problems can be traced back to a misconfiguration by a network administrator. The following are two of the most common configuration mistakes:
 - ▶ **Incorrect ACL Configuration** – an ACL that is created and applied to a NAT translation can accidentally block some or all of the IP addresses/ports required by the business. A poorly configured ACL might be the culprit if some NAT services are working while others are not.
 - ▶ **Incorrect Interface identification** – When configuring NAT, the network administrator must identify the inside and outside interfaces. It is common for an administrator to flip the inside and outside NAT interfaces which would render NAT completely inoperable.

NAT over VPN

When setting up either IPSec or L2TP VPN connections, you can run into a problem when packets are encrypted and then forced to a NAT translation. For example, IPSec tunnels in transport mode leave the IP header unencrypted. This IP header information must not be modified from the source. If it is modified (which is what NAT will do), once it reaches the destination, the remote side IPSec termination point will see that the IP header has changed and drop the packet. The result is a VPN tunnel that never gets established.

To get around this problem, Cisco invented NAT Traversal (NAT-T) which protects the original IPSec encoded packet by encapsulating it with another layer of UDP and IP headers. These layers can then be stripped off at the remote side revealing the unmodified IPSec IP header.

Latency due to NAT Processing Overhead

If your router is heavily utilized, a network administrator must realize that the NAT choke-point can be a CPU bottleneck. Translating thousands of addresses simultaneously can impact the CPU of even newer routing equipment. This processing bottleneck can cause a delay in the delivery or even drop packets. Keep a close eye on your memory and CPU utilization of your NAT router.

Applications that Do Not Play Well with NAT

One of the requirements of NAT is knowing what TCP/UDP ports a particular application uses. Some applications randomly choose ports for use in an application stream (such as RTP for voice communication). A second example is an application written to contain return IP addresses (or DNS name). That application then uses that IP address to return information back to the source. Because NAT translates addressing, the return traffic will not make it to its intended destination. Some of these application limitations can be overcome while others must avoid NAT or they will be required to be rewritten to support address translation.

NAT hiding True Source or Destination

A pet peeve that many network administrators have with NAT is when they want to troubleshoot a traffic flow from end-to-end. One of the problems with address translation is that tracking down translations in the routers NAT translation table increases the complexity of what typically is a simple troubleshooting task.

NAT troubleshooting Commands

The following show, clear and debug commands are used when identifying, troubleshooting and resolving NAT related problems on a Cisco router.

Show and Clear ip nat translation

The show ip nat translation command is by far the most popular NAT command out there. This shows the current router translation table as shown here:

```
Router#show ip nat translation
Pro Inside global      Inside local      Outside local    Outside global
udp 192.168.17.26:1220  10.100.10.55:1220  1.1.1.132:53    1.1.1.132:53
tcp 192.168.17.26:11012 192.168.1.89:11012 1.1.2.220:23    1.1.2.220:23
tcp 192.168.17.26:1067  10.100.10.55:1067  1.1.2.161:23    1.1.2.161:23
Router#
```

Figure 35: Output from the <show and clear ip nat translation> Command

If a network administrator had a problem with the translation table and it needed to be cleared of all translations so the table can start fresh, you can simply issue the clear ip nat translation * command where * means to clear out ALL translations currently in the table:

```
Router#clear ip nat translation *
```

Now if we were to issue another show ip nat translation, there will be nothing in the table until a new packet requires it as shown here:

```
Router#show ip nat translation
Pro Inside global      Inside local      Outside local    Outside global
```

Show ip nat statistics

The show ip nat statistics command gives an administrator a quick view of NAT statistics such as the number of static, dynamic and extended translations. In addition, it clearly labels the outside and inside interfaces configured on the router:

```
Router# show ip nat statistics
Total translations: 2 (0 static, 2 dynamic; 0 extended)
Outside interfaces: Serial0
Inside interfaces: Ethernet1
Hits: 135 Misses: 5
Expired translations: 2
Dynamic mappings:
-- Inside Source
access-list 1 pool net-208 refcount 2
  pool net-208: netmask 255.255.255.240
    start 172.16.233.208 end 172.16.233.221
    type generic, total addresses 14, allocated 2 (14%), misses 0
```

Debug ip nat

The debug ip nat command lets a network administrator view translations in real-time. The following example shows how the source (s=) ip address are translated (->) and sent to the destination (d=) address when flowing through the router. It then shows how the return packet translation from the source (s=) converts the destination (d=) back into (->) to the original inside local IP address:

```
Router# debug ip nat
NAT: s=192.168.1.95->172.31.233.209, d=172.31.2.132 [6825]
NAT: s=172.31.2.132, d=172.31.233.209->192.168.1.95 [21852]
NAT: s=192.168.1.95->172.31.233.209, d=172.31.1.161 [6826]
NAT*: s=172.31.1.161, d=172.31.233.209->192.168.1.95 [23311]
NAT*: s=192.168.1.95->172.31.233.209, d=172.31.1.161 [6827]
NAT*: s=192.168.1.95->172.31.233.209, d=172.31.1.161 [6828]
NAT*: s=172.31.1.161, d=172.31.233.209->192.168.1.95 [23313]
NAT*: s=172.31.1.161, d=172.31.233.209->192.168.1.95 [23325]
```

Troubleshooting VoIP

As enterprise data networks matured through the 1990's and into the new millennium, the reliability and stability grew as well. Add in the ability to priority packet streams and it soon became clear that an IP network could easily handle the duties of transporting data and voice traffic on a single network. Most modern networks are now leveraging a single network to handle both services. This section covers the requirements and pitfalls when supporting VoIP.

Voice Network Requirements

Voice traffic has more stringent requirements when compared to data applications because of the increase sensitivity to high latency. Remember that voice traffic must arrive quickly, in tact and in a uniform way for it to be properly heard at its destination. The following are voice requirements base on these needs.

- **Jitter** – the variation in the delay of packets at the destination. When an IPT phone call is established the user begins speaking into the phone. The information is placed onto the wire in evenly spaced UDP packets. These UDP packets are then sent across the network and to the destination phone. Somewhere along the network from the sending phone to the receiving phone can be network congestion or incorrect queuing mechanisms that can throw the evenly spaced packets off to where they become bunched up or out of sequence. This causes the voice quality to stutter at the end device phone.
- **Delay** – there will always be network delay when transporting any type of voice or data traffic. For voice traffic, we want eliminate the amount of variable delay. **Variable delay**, is what QoS attempts to control on a network. Variable delay refers bottleneck situations when time-sensitive traffic sits in a queue and waits for other packets to be sent out of the interface before your voice packet can be sent. By incorporating quality of service (QoS), we can help control the amount of variable delay.
- **Packet loss** – if queue bottlenecks get to the point where the queue fills up, packet loss occurs. If this happens, you can implement QoS to begin discarding less-critical packets that can be identified using QoS classification and marking.

Cisco recommends network requirements based on these three network symptoms. If your network can meet or beat the following criteria, your voice/video applications should not experience any problems.

- End-to-end delay: 150 ms or less
- Jitter: 30 ms or less
- Packet loss: 1 percent loss or less

Voice Network Requirements

The TSHOOT exam focuses specifically on (surprise) Cisco voice hardware and its requirements. When configuring Cisco IP phones on a network, there are several required and optional network requirements that network administrators must include as possible sources of problems when troubleshooting voice problems on a network. These requirements are:

- **Voice/Data VLAN on access port** – a Cisco IP phone must be put into the correct VLAN for it to operate. This is done at the access-switch. It is common to configure a special Voice VLAN. In addition, many Cisco IP phones have a secondary port on the back of the phone that is used to connect a PC. The PC data traffic should be on a separate data VLAN. For instance, we have a phone in an office along with a PC. The PC plugs into the phone and the phone into the data jack that connects to our access switch on port fa0/10. VLAN 10 is the voice VLAN and VLAN 20 is the data VLAN. Here is how the VLAN configuration should look for this port:

```
Switch(config-if)#switchport voice vlan 10
Switch(config-if)#switchport access vlan 20
```

- **DHCP** – best practice designs utilize DHCP so phones can receive IP address and TFTP information so they can reach out to the TFTP server and download the phone configuration files which contain information such as what phone numbers are to be configured on that specific phone.

- **TFTP** – As mentioned earlier, a Cisco IP phone relies heavily on TFTP servers for downloading phone configuration information as well as phone firmware updates if there are any.
- **NTP** – Another best practice service that IP phones should use is a network time protocol server (NTP). These servers help to synchronize clocks on networked devices. That way, everyone in the organization has the exact same time on their telephone displays. It is also very useful to have accurate timestamps for logging purposes.
- **CDP** – The Cisco Discovery Protocol (CDP) is a Cisco proprietary service that runs at layer 2. In regards to Cisco IP phones, it is used to identify end devices as Cisco IP phones so the switch will know what traffic to put on the Voice VLAN and what to put on a data VLAN.
- **Power over Ethernet (PoE)** – Cisco phones can be plugged directly into a wall power outlet using a power brick. Alternatively, Cisco PoE capable switches can power up the phones over the Ethernet cabling. A network administrator must be mindful of the capabilities of the PoE switch as well as the power requirements of the Cisco IP phone. Some switches support only a Cisco proprietary inline-power method (ILP) while other switches support a newer IEEE 802.3ab power method and ILP. Some phones will require more power than the switch can handle. There are also potential problems with oversubscribing the power of a PoE switch.

Phone Boot Process

In order to troubleshoot Cisco IP phones on a data network, it is important to understand the phone boot process. Here are the steps in order for booting up and bringing a phone on a Cisco network:

1. Phone powers up (power brick or PoE)
2. Phone loads firmware
3. Switch identifies it as a phone (using CDP) and puts it on a voice VLAN
4. Phone request IP information by sending DHCPDISCOVER broadcast. DHCP server responds with IP information and the IP of the TFTP server
5. The phone uses TFTP to contact the newly learned TFTP server where it downloads the phone configuration information unique to this phone based on MAC address. This configuration information contains details such as:
 - a. IP address of the Call Processing Unit (Such as the CUCM)
 - b. Telephone numbers assigned to that phone
 - c. Softkey setup configuration
6. The Cisco IP phone registers with the Call Processing Unit and is ready to make and receive voice calls

Common Voice Problems

The following are common voice problems that TSHOOT candidates should be able to identify given specific symptoms:

Not Receiving Necessary Services (i.e. CDP, DHCP or TFTP)

Many times, there are Network services that must be operating properly for Cisco IP phones to be able to connect, receive configuration information and ultimately connect to the call processing unit such as CUCM or CUCM Express. If a large number of phones are having difficulties connecting to the call-processing unit, one of the first things to check is to verify that all the supporting services are working.

Cisco IP phones use CDP so the switch knows to connect it to a voice VLAN. CDP can be disabled either on a single port or an entire switch. You can verify the operation of CDP by logging into a switch and typing `show cdp neighbor`. If you have IP phones on the network, you should be able to see them as shown here:

```
Switch#show cdp neighbor
```

```
Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater, P - Phone,
                  D - Remote, C - CVT/A, M - Two-port Mac Relay
```

Device ID	Local Intrfce	Holdtme	Capability	Platform	Port ID
SEP002290593E0F	Gig 0/43	155	H P M	IP Phone	Port 1
SEP0022905A3354	Gig 0/25	140	H P M	IP Phone	Port 1
SEP0021A02374BB	Gig 0/20	135	H P M	IP Phone	Port 1
SEP00506303B71C	Gig 0/21	178	H P	IP Phone	Port 1
SEP00229059FF88	Gig 0/22	140	H P M	IP Phone	Port 1
SEP0022905301F5	Gig 0/18	172	H P M	IP Phone	Port 1
SEP0022455A004A	Gig 0/1	149	H P M	IP Phone	Port 1

If the phones are not receiving an IP address, this can be a problem with the DHCP server itself or possibly the DHCP relay configuration setting using the `ip helper-address` command. Lastly, if your phone is getting an IP address but not receiving extension numbers, verify that your TFTP server is up and running and configured to had out the proper phone configuration files for the specific phone models your network is running.

Misconfiguration of QoS

As was mentioned earlier, QoS is used to help keep a consistent flow of voice traffic so packets don't get caught up in incoming or outgoing interface queues. We'll look at two ways to configure QoS for voice.

Modular QoS CLI (MQC)

There are dozens of books available that cover one single topic being QoS. That is because there are literally dozens of ways to implement QoS depending on the hardware/software and traffic types you are attempting to provide higher levels of network service for. Cisco is attempting to streamline the QoS process by introducing **Modular QoS CLI (MQC)**. MQC is Cisco's attempt to provide a modular framework for deploying QoS on a variety of Cisco platforms. At its core, MQC requires network engineers to configure the following:

1. **Configure a class-map** – the class-map is used to identify interesting traffic.
2. **Configure a policy-map** – the policy-map details steps to perform on a class of traffic once it is identified using the class-map.
3. **Apply the policy-map** – apply the map to the desired interfaces.

In the next section, you will learn how to view show commands that detail the configured class-maps, policy-maps and interfaces that have policy-maps applied to them.

AutoQoS

Despite Cisco's attempt to simplify QoS configurations by using MQC, it can still be very difficult to understand exactly what traffic should be identified, what policies to apply to the identified traffic, and finally, at what interfaces the policies should be applied. If you simply want to configure Cisco's best practice QoS policies for voice traffic, than AutoQoS might be a great option for your network.

AutoQoS is basically a built-in script that when applied on a router, sets up QoS policies automatically according to best-practice methodologies. This takes much of the complexity, time and operating cost to implement QoS on a network. Configuring AutoQoS is really quite easy. One "gotcha" that network administrators must look out for are the network requirements that must be met prior to enabling QoS on your network. These requirements are:

- CEF must be Enabled
- No QoS policy previously applied to an interface
- Correct bandwidth configured on interface
- IP address is required on interfaces 768 Kbps or below

Misconfiguration of Voice VLANs on Access Ports

This was brought up previously in this guide but needs to be mentioned here again. Remember that Voice VLANs that also must be used to carry data traffic act as a "trunk" connection on the access port. The data VLAN is configured using the switchport access vlan <VLAN> command while the voice VLAN is configured using the switchport voice vlan <VLAN> command.

Insufficient Power from PoE Source

Cisco IP phone models require different amounts of power when using PoE. Make sure that the PoE switch you are using is able to send the proper amount of Wattage to the phone for proper booting and operation. In addition, modular switches such as the Catalyst 4500 or 6500 series can be installed with PoE modules. One thing to keep in mind however is that attempting to add too many phones can cause the 4500/6500 series power supplies to become exhausted. To view the amount of power required by a Cisco IP phone as well as the amount of power supplied to it on a switchport, use the following command:

```
4506-switch# show power inline gigabitEthernet 3/2
Interface Admin Oper      Power(Watts) Device      Class
          From PS      To Device
-----
Gi3/2    auto   on        13.5        12.0        Cisco IP Phone 7965 3

Interface AdminPowerMax
          (Watts)
-----
Gi3/2    15.4
```

As you can see from the output, the 7965 phone is receiving 13.5 Watts of power and using 12.0 Watts for operation.

Voice Troubleshooting Commands

Now that we have an understanding of how to design and configure our network for voice, let's look at some of the more common commands used to troubleshoot configurations that have been setup on our routers to support voice in the form of MQC and AutoQoS setups:

Show class-map

The show class-map command lists all of the class-maps defined on the router. Remember that a class-map is used to identify and group traffic so policies can later be applied to them:

```
Router#show class-map
  Class Map class-1
  Match access-group 1
  Class Map class-2
  Match protocol ip
  Class Map class-3
  Match input-interface Ethernet1/0
```

Show policy-map

The show policy-map command displays the contents of all policy-maps currently configured on a router:

```
Router# show policy-map
Policy Map policy-1
  Weighted Fair Queuing
    Class class-1
      Bandwidth 937 (kbps) Max thresh 64 (packets)
    Class class-2
      Bandwidth 937 (kbps) Max thresh 64 (packets)
Policy Map policy-2
  Weighted Fair Queuing
    Class class-3
      Bandwidth 300 (kbps) Max thresh 64 (packets)
```

Show policy-map <interface>

The show policy-map <interface> command displays the policy-maps that are configured on a specific interface. In addition the keyword "output" is used. There are two different optional keywords that can be used here. They are:

- **Input** – indicates that the statistics for the attached input policy will be displayed.
- **Output** – indicates that the statistics for the attached output policy will be displayed.

Here is an example of the displayed output of this command:

```
Router# show policy-map interface output e1/1

Ethernet1/1 output : policy-1
  Weighted Fair Queueing
    Class class-1
      Output Queue: Conversation 264
      Bandwidth 937 (kbps) Max Threshold 64 (packets)
      (total/discards/tail drops) 11548/0/0
    Class class-2
      Output Queue: Conversation 265
      Bandwidth 937 (kbps) Max Threshold 64 (packets)
      (total/discards/tail drops) 11546/0/0
```

Show auto qos voip

The show auto qos voip command shows the default settings for AutoQoS as shown in this example:

```
Switch# show auto qos voip
AutoQoS is enabled with defaults as follows:
Cos-dscp map:
      cos:  0  1  2  3  4  5  6  7
-----
      dscp:  0 10 18 26 34 46 48 56
Cos-queue map:
CoS Value      : 0  1  2  3  4  5  6  7
Priority Queue : 1  1  1  3  3  4  3  3
WRR Queue      : 1  2  3  4
Bandwidth      : 20  1 80  0
```

Troubleshooting Video

Much like voice, video must also be considered more sensitive to latency and jitter compared to data traffic. In fact, video is even more sensitive than voice. The table below shows a comparison of video requirements for different Cisco video products available today:

Video Issues	Cisco Unified Video Advantage	Cisco TelePresence	Cisco Video Surveillance
Jitter	10 ms	10 ms	10 ms
One-way Delay	200 ms	150 ms	500 ms
Packet loss	0.05%	0.05%	0.5%

Figure 36: Video Requirements by Cisco Products

In regards to QoS and video, troubleshooting video is identical to troubleshooting voice. If you are configuring QoS for video, you will probably want to use MQC because AutoQoS is optimized more for the handling of voice rather than video. Also, because there are varying video requirements as shown above, you will want more control with your QoS policies which MQC will give you.

Troubleshooting Multicasting for Video Applications

Multicasting is a network service that allows for the delivery of information from a single source to multiple recipients that request the data stream. The multicast stream is copied on the network when needed and sent to clients that “join” the multicast group. Multicast is commonly found on networks that stream video. This is because it is a very efficient method of video transport. IP multicasting uses IP addresses in the range of 224.0.0.0 to 239.255.255.255. This is also known as Class D addressing space.

IP multicasting is disabled by default on Cisco routers. To enable it, you do the following:

```
Router(config)# ip multicast-routing
```

The next required configuration step is to configure protocol independent multicast (PIM). The PIM protocol keeps track of the current IP multicast service mode of receiver-initiated membership. There are three types of PIM:

- **PIM Dense Mode** – the multicast router interface PIM is configured on believes that all routers want to forward multicast packets for a group.
- **PIM Sparse Mode** – the multicast router interface PIM is configured on believes that routers do not want to forward multicast packets for a group. This mode waits until it receives a request for that multicast group before forwarding.
- **PIM Sparse-Dense Mode** – the multicast router interface PIM is configured on chooses to use either sparse or dense modes based on the multicast group setup.

Here is an example of how to configure PIM sparse-dense mode on the FastEthernet interface of a router:

```
Router(config-if)# ip pim sparse-dense-mode
```

IGMP is a service within multicast that is responsible for controlling and limiting the flow of multicast traffic throughout a network with the use of multicast queries. There are currently two widely used versions of IGMP multicasting that you should be familiar with:

- **IGMPv1** – hosts can join multicast groups. There is no way for clients to leave the group however. Instead, routers use a time-out based mechanism to discover the groups that are of no interest to the members.
- **IGMPv2** – host leave messages were added to the protocol. This allows group membership termination to be quickly reported to the routing protocol.

A layer 2 switch feature called IGMP snooping can be enabled on switchports to help make multicasting more efficient. As the name implies, the switch will listen for IGMPv1 and v2 join and leave requests and build a table containing this multicast group information for the switch. The forwarding of multicast information at layer 2 then becomes much more efficient because only those devices that are in the layer 2 multicast table are forwarded the group packets. Here is an example of a multicast table for a particular layer 2 VLAN:

```
Switch#show mac-address-table multicast vlan 1
Multicast Entries
vlan      mac address      type      ports
-----+-----+-----+-----
      1      0100.5e01.0101   igmp     Switch,Gi6/1
      1      0100.5e01.0102   igmp     Switch,Gi6/1
      1      0100.5e01.0103   igmp     Switch,Gi6/1
      1      0100.5e01.0104   igmp     Switch,Gi6/1
      1      0100.5e01.0105   igmp     Switch,Gi6/1
      1      0100.5e01.0106   igmp     Switch,Gi6/1
Switch#
```

Common Video Problems

This section covers some of the more common root causes when video problems are found on the network.

Bandwidth

Video applications are difficult to handle on a network because they are very “bursty” by nature. It is vitally important that there be a large proportion of bandwidth available when a network has heavy usage of various video applications.

QoS

Along the same lines as bandwidth issues, video applications are highly sensitive to latency. Video packets that get caught up in queues can lead to poor performance and dropped video streams. A sound QoS policy for video on a network goes a long way to fixing many video problems.

Multicast

Video applications often rely on multicasting to transport video streams to their destination. It is important to insure the routers along the path are properly configured for multicasting and the interfaces are configured with PIM. In addition, you should verify the IGMP snooping is enabled on the proper VLANs. To do this, issue the `show ip igmp interface <vlan>` command as shown in the following example:

```
Switch# show ip igmp interface vlan 101
IGMP snooping is globally enabled
IGMP snooping is enabled on this Vlan
IGMP snooping immediate-leave is disabled on this Vlan
IGMP snooping mrouter learn mode is pim-dvmrp on this Vlan
IGMP snooping is running in IGMP-ONLY mode on this VLAN
Switch#
```

Multicast Video Troubleshooting Commands

The following commands are used to troubleshoot multicast which is used in many video streaming applications that have a single source and multiple destinations.

Show ip igmp interface <interface>

The show ip igmp interface <interface> command can be used to verify that IGMP is enabled on router interfaces. If IGMP is enabled on the interface, the output shows the multicast designated router (DR) and which multicast groups end stations attached to this interface belong to:

```
Router#show ip igmp interface
FastEthernet1/0 is up, line protocol is up
  Internet address is 10.10.10.114, subnet mask is 255.255.255.0
  IGMP is enabled on interface
  IGMP query interval is 60 seconds
  Inbound IGMP access group is not set
  Multicast routing is enabled on interface
  Multicast TTL threshold is 0
  Multicast designated router (DR) is 10.10.10.33
  No multicast groups joined
FastEthernet1/1 is up, line protocol is up
  Internet address is 10.10.11.129, subnet mask is 255.255.255.0
  IGMP is enabled on interface
  IGMP query interval is 60 seconds
  Inbound IGMP access group is not set
  Multicast routing is enabled on interface
  Multicast TTL threshold is 0
  Multicast designated router (DR) is 10.10.11.131
  Multicast groups joined: 225.2.2.2 226.2.2.2
```

Show ip igmp groups

To display the multicast groups with receivers that are directly connected to the router and that were learned through Internet Group Management Protocol (IGMP), use the show ip igmp groups command in user EXEC or privileged EXEC mode.

```
Router#show ip igmp groups
IGMP Connected Group Membership
Group Address      Interface      Uptime      Expires      Last Reporter
239.255.255.254   Ethernet3/1   1w0d        00:02:19    172.21.200.159
224.0.1.40        Ethernet3/1   1w0d        00:02:15    172.21.200.1
224.0.1.40        Ethernet3/3   1w0d        never        172.16.214.251
224.0.1.1         Ethernet3/1   1w0d        00:02:11    172.21.200.11
224.9.9.2         Ethernet3/1   1w0d        00:02:10    172.21.200.155
232.1.1.1         Ethernet3/1   5d21h       stopped     172.21.200.206
```

Troubleshooting Wireless Networks

The TSHOOT exam focuses on how to troubleshoot wireless networks from a wired point of view. There is no need to have a deep understanding of wireless spectrums or omni vs. directional antennas. Instead, you must understand what various wireless components/architectures do and an overview of 802.1x authentication methods.

All other troubleshooting services such as DHCP, PoE and QoS have already been discussed in this guide. The TSHOOT exam does not require you to be an expert on this subject but rather help you to have a general understanding of what a network administrator might see in an enterprise network.

Wireless Components

- **Wireless Clients** – these are end devices such as wireless PCs, laptops and handheld devices. Most new devices have 802.11a/b/g/n radios built-in. Modern operating systems have software built-in to the operating system that assists in the setup of wireless access including SSID and authentication methods.
- **Autonomous Access Point (aAP)** – an autonomous access point (aAP) is a stand-alone device that is a single “hot-spot” of wireless connectivity. The wireless intelligence of the network device is built into the access-point itself. It connects directly to an access layer switch for LAN connectivity.
- **LWAPP** – Lightweight Access Point Protocol (LWAPP) a protocol for a centralized WLAN architecture that tunnels wireless client traffic from the lightweight access point radio to a centralized WLAN controller (WLC). The wireless intelligence of the network architecture is located at the centralized WLC that controls multiple lightweight access points.
- **Wireless LAN Controller (WLC)** – as stated earlier, the WLC is the intelligence in an LWAPP architecture. The WLC is an appliance that controls wireless connectivity to multiple lightweight access points. **NOTE:** By default the WLC does not send any broadcast or multicast traffic out to the WLAN client devices. Keep this in mind when troubleshooting problems such as wireless devices that are having problems connecting to multicast video streams. There are some WLC terms that TSHOOT candidates should understand for the exam:
 - ▶ **Ports** – Physical Ethernet interfaces that connect the WLC to the LAN which is typically a core or distribution layer switch such as the Catalyst 6500 series.
 - ▶ **Interfaces** – A logical VLAN mapping that connects wireless VLANs with the wired LAN infrastructure.
 - ▶ **WLANs** – The wireless portion of the WLC setup. WLANs contain wireless services such as 802.1x authentication, Wireless QoS and wireless security in the form of SSIDs and encryption methods.
- **Wireless Control System (WCS)** – The Wireless Controller System (WCS) is found in large wireless deployments. It is a server application that helps to monitor, troubleshoot, and report on wireless network components including access points and WLC appliances.

Wireless Architectures

There are two main wireless architectures that Cisco recommends to organizations. The first is autonomous mode for small hot spot deployments and the other is split MAC mode for large enterprise environments.

Autonomous Mode

As discussed earlier, the aAP mode is simply an access point that has one or more 802.11 WiFi radios to connect to wireless devices. The intelligence is built-in to the aAP and it simply connects back to an access switch for LAN connectivity. This design architecture works well for small hot spot deployments but it does not scale well because each aAP has its own configuration that needs to be managed. Here is a diagram of this simple wireless architecture:



Figure 37: A Simple Wireless Architecture

Controller Mode (aka – LWAPP split MAC mode)

The LWAPP split MAC architecture breaks the wireless network into two separate devices, the lightweight access point, and the WLC. These are linked via the LWAPP protocol across a network to provide the same functionality of radio services. This design also is much easier to deploy and manage than autonomous access points when in a large enterprise environment. Here is a diagram that shows the split MAC architecture:




Figure 38: A Split MAC Wireless Architecture

In this type of architecture the lightweight access points have the following responsibilities:

- Initial setup communication between a client and AP
- Frame beacons
- Buffering and transmission of frames for wireless devices that are running in power-save mode
- Monitoring each of WiFi channels for things like external noise and interference. aAP's also discover other legitimate and rogue access points
- Encryption and decryption of 802.11 frames

And the WCS handles these responsibilities:

- Authentication
- Wireless association and mobility to other LWAPs
- 802.11 frame translation and bridging
- Places wireless traffic onto a LAN switch

802.1x Authentication

One major security difference between administrating a wired vs. a wireless network has to do with authentication. If an end user wants to connect to a wired network, they must have physical access to a switchport off of an access switch. Wireless access is different in the fact that the physical security is not as effective. Now you can have non-authorized people sitting in the coffee shop next door or out in the parking lot having access to your WiFi network. Because of this, it is vital that wireless users be properly authenticated to protect the entire LAN from being accessed by outsiders.

802.1x Authentication can be used on a Cisco network with the corporation of a **Cisco Access Control Server (ACS)**. 802.1x is an authentication protocol that can be enabled on a wireless network. It can be setup to prompt wireless users for a username and password, when credentials are entered, that information is sent to a centralized database, which is possibly controlled by an ACS. The ACS works to accept or deny access to the network.

In a split MAC environment, the WLC can be configured to use an Extensible Access Protocol (EAP) to communicate with the ACS server. EAP can be configured in one of the following five methods for a particular SSID:

- **EAP-TLS** – an IETF standard Wireless clients use PKI to secure communication to a RADIUS authentication server such as Cisco ACS.
- **EAP-PEAP** – similar in nature to EAP-TLS. EAP-PEAP however allows for better backwards compatibility with legacy EAP methods. It does this by encapsulating EAP messages within TLS headers.
- **EAP-TTLS** – similar to EAP-PEAP in the fact that it tunnels TLS but EAP-TTLS encapsulates TLS messages inside the payload rather than the TLS itself.
- **LEAP** – a Cisco proprietary method of using rotating WEP keys to better secure WEP. LEAP is widely know to be crackable and is not recommended any longer.
- **EAP-FAST** – a Cisco proprietary method to address the security problems found in LEAP.

Troubleshooting Advanced Services

In this section, we will cover troubleshooting some of the more advanced IP services on Cisco routers. All of these services can be used to baseline IP traffic for monitoring and troubleshooting purposes. The three services that will be covered are:

- NetFlow
- IP SLA
- NBAR

An administrator studying for the TSHOOT exam must understand the purpose of each of these services as well as which show commands to use in troubleshooting situations.

NetFlow Overview

NetFlow is a service found in router IOS images that tracks IP flow statistics across an interface or SVI. NetFlow can be used to baseline and track information such as:

- Top talkers by IP address
- Top talkers by protocol/application
- ToS Markings used in QoS
- Data Volumes
- Interface data Rates
- Interface utilization

NetFlow is configured on the interface by using the `ip flow ingress` (or `egress` with NetFlow v9). You will also want to set a source interface for the router such as a loopback interface. This is done using the `ip flow-export source <interface>` command. Collected flows can then be exported to a NetFlow collector using the `ip flow-export destination <IP_address> <port>` command. **NOTE** that NetFlow does not have a standard port so make sure you configure the routers to use the same port that your collector is listening on. Also make sure that this port is not being blocked by an ACL or firewall from the source to the destination.

NetFlow Troubleshooting Commands

Once you have NetFlow configured, you can view the statistics using the IOS command line by using the `show ip cache flow` command as shown here:

```
Router#show ip cache flow
IP packet size distribution (230151 total packets):
  1-32  64  96  128  160  192  224  256  288  320  352  384  416  448
480
    .929 .070 .000 .000 .000 .000 .000 .000 .000 .000 .000 .000 .000 .000
.000 .000

    512  544  576 1024 1536 2048 2560 3072 3584 4096 4608
    .000 .000 .000 .000 .000 .000 .000 .000 .000 .000 .000
```

```

IP Flow Switching Cache, 4456448 bytes
 65509 active, 27 inactive, 820628747 added
 955454490 aged polls, 0 flow alloc failures
Exporting flows to 1.1.15.1 (2057)
 820563238 flows exported in 34485239 udp datagrams, 0 failed
last clearing of statistics 00:00:03

```

Protocol	Total	Flows	Packets	Bytes	Packets	Active(Sec)	Idle(Sec)
-----	Flows	/Sec	/Flow	/Pkt	/Sec	/Flow	/Flow
TCP-BGP	71	0.0	1	49	0.0	2.5	15.8
UDP-other	17	0.0	1	328	0.0	0.0	15.7
ICMP	18966	6.7	10	28	72.9	0.1	22.9
Total:	19054	6.7	10	28	72.9	0.1	22.9

Here you can see useful information such as the IP packet size distribution (i.e. 92.9% of packets seen are 1-32 bytes). In addition, the command breaks out flows by TCP/UDP or IP port.

Overview of IP Service Level Agreement (IP SLA)

The Cisco IP Service Level Agreement (IP SLA) is a feature that can be enabled to monitor important servers/applications and network links to help track uptime for service level agreement contracts. The service can monitor up to the session layer of the OSI model (port).

IP SLA works by having the router it is configured on simulate traffic that the destination should respond to in a particular way such as Pinging a device or checking to see if TCP port 80 is open. These probes can be used to track performance and uptime statistics over a period of time. IP SLA is configured by using the ip sla monitor responder family of commands.

IP SLA Troubleshooting Commands

The following commands can be useful when troubleshooting IP SLA configurations on a router.

Show ip sla monitor responder

The show ip sla monitor responder command displays recent monitor probes and breaks the out by messages received and messages failed as shown here:

```

IP SLA Monitor Responder is: Enabled
Number of control message received: 15 Number of errors: 1
Recent sources:
192.168.10.254 [19:11:49.035 UTC Sat May 15 2010]
192.168.10.254 [19:10:49.023 UTC Sat May 15 2010]
192.168.10.254 [19:09:48.707 UTC Sat May 15 2010]
192.168.10.254 [19:08:48.687 UTC Sat May 15 2010]
192.168.10.254 [19:07:48.671 UTC Sat May 15 2010]
Recent error sources:
192.168.10.254 [19:10:49.023 UTC Sat May 15 2010] RTT_AUTH_FAIL

```

Show ip sla statistics

The show ip sla statistics command displays the current state of IP SLA monitoring. This includes the following pieces of information:

- Current operation state
- Number of probes attempted
- How long (in seconds) until the monitoring ends
- Latest probe completion time
- Probe round trip time (in ms)

Here is an example of the output of this show command:

```
Router#show ip sla statistics
      Current Operational State
Entry Number: 3
Modification Time: *22:15:43.000 UTC Sat May 15 2010
Diagnostics Text:
Last Time this Entry was Reset: Never
Number of Octets in use by this Entry: 1332
Number of Operations Attempted: 2
Current Seconds Left in Life: 3511
Operational State of Entry: active
Latest Completion Time (milliseconds): 544
Latest Operation Start Time: *22:16:43.000 UTC Sat May 15 2010
Latest Oper Sense: ok
Latest Sense Description: 200 OK
Total RTT: 544
DNS RTT: 12
TCP Connection RTT: 28
HTTP Transaction RTT: 504
HTTP Message Size: 9707
```

Overview of Network-based Application Recognition (NBAR)

Network-based Application Recognition (NBAR) is a service that can be enabled on Cisco IOS routers to serve one of two purposes:

- Classify traffic for use in QoS
- Protocol discovery for baselining IP traffic

This section will focus on the use of NBAR for discovery and baseline usage. By enabling NBAR on an interface using the ip nbar protocol-discovery command, the service will begin identifying the protocols and ports that are the top talkers of traffic passing through.

Keep in mind that NBAR is not incredibly intelligent. It will not be able to recognize applications that use non-standard ports. Also, if any applications pass traffic using well known ports, it cannot differentiate between this traffic. For example, if an application streamed music over TCP 80, NBAR would think that this is simply HTTP traffic and cannot split the two up into separate groups.

NBAR Troubleshooting Commands

The following command is used to view statistics on routers configured for NBAR protocol discover on one or more interfaces.

Show ip nbar protocol-discovery

If your router is configured for NBAR protocol discovery than the show ip nbar protocol-discovery command will give statistics for the past 5 minutes on each interface NBAR is configured on. Here is an example of this output:

```
Router#show ip nbar protocol-discovery
FastEthernet1/0
```

Protocol	Input		Output	
	Packet Count	Byte Count	Packet Count	Byte Count
	5 minute bit rate (bps)		5 minute bit rate (bps)	
igmp	316773	26340105	0	0
streamwork	3000	4437	0	7367
rsvp	2301891	3000	339213	0
ntp	279538	319106191	14644	673624
.	0	0	0	0
.	8979	906550	7714	694260
.	0	0	0	0
Total	17203819	19161397327	151684936	50967034611
		4179000		6620000

As you can see, the output breaks out the data by protocol and port as well as gives in/out byte statistics for the interface.

Troubleshooting Security

Security plays a major role in modern networks and will continue to grow. It used to be that security was pushed out to the edge where virus software and local authentication handled the bulk of the work. It has become commonplace now to integrate security features into the network that can provide a uniform blanket of security across all end devices. This section covers network security structures and some IOS troubleshooting commands that can be used on Cisco IOS devices.

Securing Router and Switch Planes

Cisco divides a router into three major “operational planes” that are responsible for different tasks. These three operational planes are:

- Management plane
- Control plane
- Data plane

We will discuss what each plane does and how security is used to protect against malicious behavior on a network.

Management Plane Security

The management plane is the processes of the router and IOS that provide access to the IOS device for network administration purposes. There are three main types of management plane access into IOS devices. They are:

- **Command line interface (CLI)**
 - ▶ Console port
 - ▶ Telnet
 - ▶ SSH
- **Web graphical user interface (GUI)**
 - ▶ HTTP
 - ▶ HTTPS
- **Simple Network Management Protocol (SNMP)**
 - ▶ Uses read/write strings for authentication

Access to the management plane CLI and web GUI are controlled by the use of authentication usernames and passwords. In addition, specific users can be granted a certain level of authorization access. Finally, authenticated and authorized users access and adds/changes/deletions can be logged and stored through the use of accounting tools. The authentication, authorization and accounting are known as **AAA**.

AAA Security

AAA on Cisco devices is commonly controlled by a Cisco Access Control Server (ACS). The ACS server is used to maintain a database of users, authorization access and is a logging facility for accounting records. While some smaller deployments will simply have a username and password configured locally on IOS devices, larger businesses use ACS servers and configure AAA on routers/switches to communicate to one or more ACS servers using the TACACS+ or RADIUS protocols.

Management Plane Security Troubleshooting Commands

The AAA authentication piece of the management plane is often a source of TSHOOT questions. Network administrators must be able to understand how to read debug information to see where a potential problem lies along the authentication line. Remember that when using an ACS server, open communication must exist between the IOS device and the ACS server. If you are having communication problems, make sure there is not a firewall or ACL preventing traffic from getting through.

Debug AAA authentication

The following command displays typical output from the debug aaa authentication command. Here you see that a user is authenticating against a TACACS+ server that is communicating with a Cisco ACS server. The user successfully authenticates (PASS) and is granted authorization up to the appropriate level that is assigned to that User ID:

```
Router#debug aaa authentication
AAA/AUTHEN: create_user user='' ruser='' port='tty19' rem_
addr='172.31.60.15' authen_type=1 service=1 priv=1
AAA/AUTHEN/START (0): port='tty19' list='' action=LOGIN
service=LOGIN
AAA/AUTHEN/START (0): using "default" list
AAA/AUTHEN/START (50996740): Method=TACACS+
TAC+ (50996740): received authen response status = GETUSER
AAA/AUTHEN (50996740): status = GETUSER
AAA/AUTHEN/CONT (50996740): continue_login
AAA/AUTHEN (50996740): status = GETUSER
AAA/AUTHEN (50996740): Method=TACACS+
TAC+: send AUTHEN/CONT packet
TAC+ (50996740): received authen response status = GETPASS
AAA/AUTHEN (50996740): status = GETPASS
AAA/AUTHEN/CONT (50996740): continue_login
AAA/AUTHEN (50996740): status = GETPASS
AAA/AUTHEN (50996740): Method=TACACS+
TAC+: send AUTHEN/CONT packet
TAC+ (50996740): received authen response status = PASS
AAA/AUTHEN (50996740): status = PASS
```

Control Plane Security

The control operational plane is the software that determines the path selection for the transport of traffic from point A to point B at layers 2 and 3. At layer 2, the Spanning Tree Protocol is the control plane software. At layer 3, routing protocols and other tables such as the ARP table are part of the control plane. In addition, services such as DHCP, which assist end devices with receiving correct addressing information, are part of the control plane. This section covers some of the security features of the control plane.

Securing Spanning Tree

The Spanning Tree Protocol (STP) is a protocol that is designed to prevent loops at the switch layer (layer 2). The protocol is responsible for talking to other layer 2 switches using bridge protocol data units (BPDUs). Using various factors, the protocol selects a single switch within a VLAN to be the root bridge over all the others. All of the root bridge ports are called root ports and put into a forwarding state which means that traffic can be sent and received on them. Other non-root switches will run STP and put the ports that have the shortest path to the root bridge in a forwarding state and the other ports in a blocking state. Ports in a blocking state do not send or receive data frames but will become active if a forwarding link were to fail. By having a single path to the root bridge, STP guarantees that the network has a loop-free topology.

But as we all know, things can happen both accidentally and on purpose. That's why there are some STP security measures that should be implemented to prevent from an accidental loop at layer 2 which can create broadcast storms and ultimately impact network performance. The following are two of the most important STP security measures to implement.

STP BPDU Guard

As stated previously, the Spanning Tree Protocol uses BPDUs to communicate with neighboring switches to determine things such as the root bridge and forwarding/blocking ports. By default, any switchport that comes online must go through a spanning tree process that cycles through the following stages while communicating with neighbor switches:

- Listening
- Learning
- Blocking or Forwarding

The time it takes for a switch port to reach a forwarding state is 30 seconds. Keep in mind that this process runs even if an end device is connected such as a PC. That means that when a PC is plugged into a switchport, 30 seconds must pass before the switch determines the port should be placed in a forwarding state and data traffic is allowed to flow. Because of this, you can configure ports you know will only have end devices (non-STP devices) on them in spanning-tree portfast mode. This mode will still cycle through the STP phases but the difference is that the port is immediately put into forwarding mode when it comes up.

This is great, but what if a switch is accidentally placed onto a port configured for portfast or if an Ethernet cable is connected between two switchports creating a physical loop where both ports are forwarding? This is a serious security problem that can bring your network to its knees. That is why it is highly recommended to enable BPDU guard on all ports that are configured in portfast mode. Here's an example of how to configure an access port for portfast and BPDU guard:

```
Switch#configure terminal
Switch(config)#interface gi2/10
Switch(config-if)#switchport mode access
Switch(config-if)#spanning-tree portfast
Switch(config-if)#spanning-tree bpduguard enable
```

BPDU guard listens on the switchport for BPDUs and when it sees one, it puts the port in a disabled mode (err-disable) as shown here using the following show command:

```
Switch#show interfaces gigabitethernet 2/10 status
Port    Name    Status          Vlan    Duplex  Speed    Type
Gi2/10          err-disabled   100      full    1000    1000BaseSX
```

A network administrator then has to manually login to the switch and re-enable it for it to come back online by issuing a shutdown and no shutdown on the interface. Because no BPDUs should ever been seen on access ports, the switch recognizes something is wrong and proactively prevents any problems by simply turning the port off until the problem can be investigated further.

STP Root Guard

One of the other issues with standard STP is the fact that there are no controls to strictly handle the STP topology of a network. Left untouched, a switch that is either unstable or in an undesirable topology location may be chosen to be the root. **Root guard** prevents a switch from becoming the root bridge. Root guard is performed on a per-port basis and insures that the root guard configured port maintains a designated state as shown in this example:

```
Switch#configure terminal
Switch(config)#interface gi1/1
Switch(config-if)#spanning-tree guard root
```

Routing Service Security

Protecting routing protocols comes down to authentication. If routing protocols are left unauthenticated, other users with access to the wired network could potentially connect a router and begin manipulating routes using the same routing protocol you are running. If your network is large and there are many users, it is best to protect your equipment from rogue routers trying to form neighbor relationships. In addition router next-hop gateway protocols can authenticate against each other to prevent unwanted outsiders from causing problems. Let's briefly examine how to configure EIGRP, OSPF and some of the gateway redundancy protocol authentication mechanisms.

Another important note for RP security is to use the passive-interface command to prevent unintended interfaces from forming neighbor relationships.

Routing Protocol Authentication

Routing protocols such as EIGRP and OSPF have mechanisms used so neighbors must properly authenticate before exchanging routing table information. This prevents routes from being exchanged with unknown or improperly configured routers.

In regards to EIGRP, here is an example of how to properly configure authentication using a key chain called "My_Chain", key #1 and key-string of "my_password". **NOTE:** that EIGRP only supports MD5 hashed passwords. Here is the example configuration:

```
Router#configure terminal
Router(config)#key chain My_Chain
Router(config-keychain)#key 1
Router(config-keychain-key)#key-string my_password
Router(config-keychain-key)#exit
Router(config-keychain)#exit
Router(config)#interface fa0/1
Router(config-if)#ip authentication mode eigrp 10 md5
Router(config-if)#ip authentication key-chain eigrp 10 MYCHAIN
```

From an OSPF authentication perspective, there are considered to be three different types:

- **Null Authentication** – Type 0, which means no authentication information is passed to neighbors.
- **Plain Text Authentication** – Type 1, which passes passwords in clear-text to neighbors.
- **MD5 Authentication** – Type 2, which passes passwords using an MD5 form of encryption.

Gateway Redundancy Authentication

Just like routing protocols, gateway redundancy protocols should also use authentication. Similar to router neighbor relationships, gateway redundancy protocols rely on more than one router to function. Because of the dependence on other hardware, it is important that a level of authentication be built-in to the protocol so the device can be fully trusted.

Troubleshooting authentication problems again boils down mainly to configuration problems. HSRP for example can be configured for either plain-text or MD5 encrypted passwords. The methods must match otherwise authentication will fail.

Securing DHCP

DHCP servers are a great method for providing address information to clients dynamically on a network. Having a centralized database where an administrator can control things like subnets, pools, lease times and other settings helps ease the administration burden. However, hackers can use the fact that clients request information from a server using a broadcast message to their advantage. One thing that hackers can do is to place their own DHCP server on the network. When clients on the same subnet as the spoofed DHCP server request an IP address, the fake DHCP server can respond before the real one does. The DHCP server can then give clients the IP address of the DHCP server as the default gateway so now all routed data from these client devices gets sent to the spoofed DHCP server where data can be collected and used for malicious purposes.

In response to this, Cisco switches can be configured for **DHCP snooping**. This feature can be configured to either trust or not to trust specific switch ports. If a port is not trusted, then it will never receive DHCP responses even if it is on the same IP subnet as a client device.

Securing ARP

ARP spoofing is when a device sends an ARP broadcast message to the entire VLAN announcing its IP address. Hackers can attempt to send out an ARP broadcasting that it is the default gateway IP address. This can cause other devices in the VLAN to send all of their gateway bound traffic to the attacking device instead of the real default gateway. This is referred to as a **man-in-the-middle** attack. To prevent this, network administrator can use port security and 802.1x features to prevent the attacker from ever getting on the network. The engineer can also use and ARP entry rules to protect against the spoof attack.

Static ARP

Instead of dynamically learning the MAC address to IP address entry, you can use **static ARP** commands to assign the MAC to Static address. This static ARP entry takes precedence over any dynamically learned entry. Note that this security method greatly increases the switch management overhead required. An example of a static ARP is as follows:

```
Switch#configure terminal
Switch(config)#arp 10.1.1.100 010e.1234.5678.1234 arpa
```

Dynamic ARP Inspection

Dynamic ARP Inspection (DAI) uses VLAN ACLs to filter ARP traffic on a particular VLAN. ARP inspection can be configured to allow ARP traffic to a specific MAC address only and deny all other traffic. A hacker would not be able to become the man in the middle using this security feature. To configure DAI, you need to first make an ARP ACL mapping the IP address to a MAC address. You can then enable ARP inspection on a single VLAN or multiple VLANs as shown in the syntax section. Here is an example of how a DAI filter would be configured using an ACL named my_filter and applying it to VLAN 10:

```
Switch#configure terminal
Switch(config)#ip arp inspection filter my_filter vlan 10
Switch(config)#ip arp inspection vlan 10
```

Control Plane Security Troubleshooting Commands

The following commands are a sample of the commands that can be used to troubleshoot control plane security features. These commands can be used to verify that the security features are functioning properly on a network.

Debug ip eigrp packets

If you are having problems forming EIGRP neighbor relationships and have authentication configured, run the debug eigrp packets command to see if there is a problem when neighbors try to match authentication methods and passwords as shown here:

```
Router#debug eigrp packets
EIGRP: ignored packet from 192.169.1.2 (invalid authentication)
EIGRP: ignored packet from 192.169.1.6 (invalid authentication)
```

As you can see from this example, the authentication is invalid which likely means that the key-string on both of the routers does not match.

Debug ip ospf adj

The debug ip ospf adj command lets administrators view authentication logs in real time to troubleshoot any misconfigurations including authentication type or password mismatches between neighbors. The following debug output shows that the neighbors have an authentication type mismatch between them which is preventing authentication from occurring and ultimately, preventing the neighbor relationship from forming:

```
Router# debug ip ospf adj
OSPF: Rcv pkt from 192.16.64.2, Serial0 : Mismatch Authentication type.
```

Show ip ospf interface <interface>

The show ip ospf interface <interface> command lists interface information on OSPF configured interfaces. This includes showing if the link has authentication enabled and which type. This example shows that MD5 authentication is enabled on the link:

```
Router# show ip ospf interface serial0
Serial0 is up, line protocol is up
Internet Address 192.16.64.1/24, Area 0
Process ID 10, Router ID 172.16.10.36 , Network Type POINT_TO_POINT,
Cost: 64
Transmit Delay is 1 sec, State POINT_TO_POINT,
```

```

Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
  Hello due in 00:00:05
Index 2/2, flood queue length 0
Next 0x0(0)/0x0(0)
Last flood scan length is 1, maximum is 1
Last flood scan time is 0 msec, maximum is 4 msec
Neighbor Count is 1, Adjacent neighbor count is 1
  Adjacent with neighbor 70.70.70.70
Suppress hello for 0 neighbor(s)
Message digest authentication enabled
  Youngest key id is 1

```

Debug standby errors

The debug standby errors command will show any HSRP errors when communicating to peers. In the example output below, we see that one peer is configured for plain-text authentication while the other is configured for MD5:

```

Router#debug standby errors
A:Jun 16 12:14:50.337:HSRP:Et0/1 Grp 0 Auth failed for Hello pkt from
10.21.0.5, MD5  cfgd but no tlv
B:Jun 16 12:16:34.287:HSRP:Et0/1 Grp 0 Auth failed for Hello pkt from
10.21.0.4, Text auth  failed

```

Show ip arp inspection statistics

If you have configured DAI on an interface or vlan and want to check the forwarding and dropped statistics, an administrator can issue a show ip arp inspection statistics command. The following output shows statistics for DAI configured on VLAN 10:

```

SwitchA#show ip arp inspection statistics vlan 10
SwitchA#

```

Vlan	Forwarded	Dropped	DHCP Drops	ACL Drops
10	2	2	2	0

Vlan	DHCP Permits	ACL Permits	Source MAC Failures
10	2	0	0

Vlan	Dest MAC Failures	IP Validation Failures
10	0	0

Data Plane Security

The data plane is where the actual transport of packets resides. This is application data such as HTTP, Telnet, and voice/video protocols. Protecting data usually means that you want to allow only traffic that should be allowed to traverse your network. Most large networks use security-targeted appliances such as an ASA firewall or Cisco IPS solution. Cisco routers and switches also have tools built-in to the IOS to protect the data plane. This section will explain some of those tools and what commands can be used to troubleshoot them.

Access Control Lists

When network administrators first think about how a router can protect the data plane, access control lists usually is the first thing that pops into their minds. ACLs can be used for many situations but one of the more popular reasons is to configure lists of permit and deny traffic and then apply that access-list to an interface or SVI using the access-group <acl> in (or out) command.

IOS Firewall Feature Set

As part of your router IOS purchase, you can choose to purchase a license of the **IOS Firewall Feature Set**, which was formally called the Context-based access control (CBAC). This special IOS allows the IOS to act more like a true firewall. It breaks up interfaces on the route into trusted and untrusted networks. Trusted networks can communicate to untrusted networks and this data flow is stored in a state table. Returning traffic is then matched to that state table and is permitted or denied based on the IOS firewall rules that are in place. All other traffic originating from an untrusted network attempting to go into a trusted network will be denied.

VPN Tunnels

If an administrator needs to transport sensitive data over a foreign network, it is advisable to encrypt and tunnel the data so it cannot be read or tampered with. By using a VPN method such as IPSec, you can insure that hackers cannot sit somewhere in between the source and destination and intercept data or inject malicious code. This is referred to as the **man-in-the-middle** attack. VPN tunnels have endpoints where tunnels terminate. As a network administrator, you want both ends of the VPN termination endpoint to be on trusted networks.

802.1x Authentication

The 802.1x standard uses Extensible Authentication Protocol (EAP) to authenticate end users prior to giving them access to the network. Authentication happens through a RADIUS server such as the Cisco ACS server. Once authenticated, the user has access to the network. If authentication fails, the end device cannot connect to the switch.

Data Plane Security Troubleshooting Commands

Show ip access-lists <name>

The show ip access-lists <name> command displays the rules defined within a specific named access list. If you want to see all configured access-lists, simply do not list a specific ACL name:

```
Router#show ip access-lists dev1
Standard IP access list dev1
100 permit 10.4.4.0, wildcard bits 0.0.0.255
105 permit 10.5.5.0, wildcard bits 0.0.0.255
115 permit 10.6.0.0, wildcard bits 0.0.0.255
120 permit 10.7.5.0, wildcard bits 0.0.0.255
125 permit 172.16.0.0, wildcard bits 0.0.255.255
```

Also note the use of sequence numbers within named access lists. This helps administrators to add ACL entries within an existing ACL without having to remove the entire thing.

Show ip inspect all

The show ip inspect all command lists all of the IOS Firewall inspect configurations in addition to all currently tracked sessions located at the bottom. You can see the source and destination IP and port numbers as well as a description if the stream is using a well-known port. Here is an example of the command output:

```
Router#show ip inspect all
Session audit trail is disabled
one-minute (sampling period) thresholds are [400:500] connections
max-incomplete sessions thresholds are [400:500]
max-incomplete tcp connections per host is 50. 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
Inspection Rule Configuration
  Inspection name all
    tcp timeout 3600
    udp timeout 30
    ftp timeout 3600
Interface Configuration
Interface Ethernet0
  Inbound inspection rule is all
    tcp timeout 3600
    udp timeout 30
    ftp timeout 3600
  Outgoing inspection rule is not set
  Inbound access list is not set
  Outgoing access list is not set
Established Sessions
Session 25A6E1C (30.0.0.1:46065)=>(40.0.0.1:21) ftp SIS_OPEN
Session 25A34A0 (40.0.0.1:20)=>(30.0.0.1:46072) ftp-data SIS_OPEN
```

Show ip inspect session

A slightly different IOS Firewall inspect command is the show ip inspect session command that show the currently inspected established sessions:

```
Router#show ip inspect session
Established Sessions
Session 25A3318 (10.0.0.1:20)=>(10.1.0.1:46068) ftp-data SIS_OPEN
Session 25A6E1C (10.1.0.1:46065)=>(10.0.0.1:21) ftp SIS_OPEN
```

Show dot1x interface <interface>

To display 802.1x settings and statistics for an interface it is configured on, use the show dot1x interface <interface> command. Here is an example of the output from this command:

```
Switch#show dot1x interface fastethernet6/1
Dot1x Info for FastEthernet6/1
-----
PAE = AUTHENTICATOR
PortControl = AUTO
ControlDirection = Both
HostMode = SINGLE_HOST
ReAuthentication = Disabled
QuietPeriod = 60
ServerTimeout = 30
SuppTimeout = 30
ReAuthPeriod = 3600 (Locally configured)
ReAuthMax = 2 MaxReq = 2
TxPeriod = 30
RateLimitPeriod = 0
Critical-Auth = Enabled
Critical Recovery Action = Reinitialize
Critical-Auth VLAN = 3
Dot1x Authenticator Client List Empty
Authorized By = Critical-Auth
Vlan Policy = 3
```

Note that you can verify information such as if and how much time until the user must re-authenticate on the port, the maximum number of re-authentications allowed and what action the port will take if 802.1x failure occurs. Also note the HostMode section settings. If it is set to SINGLE_HOST that means only one device can authenticate at a time on that port. If this setting said MULTIPLE_HOSTS than more than one would be allowed simultaneously.

Troubleshooting IPv6

The push toward readdressing devices to IPv6 is due to the demand for IPv4 outgrowing the supply for it. There simply aren't enough IPv4 addresses to go around. NAT has helped to slow the need for a full migration to IPv6 but eventually, our networks will have to change to support the new protocol. This section will cover the basics of IPv6 as well as troubleshooting IPv6 routing protocols.

IPv6 Overview

This section will cover the basics of IPv6 to give you a general understanding of how the protocol works compared to IPv4.

IPv6 Header

The IPv6 header is much more simplified compared to an IPv4 header. It only contains far fewer fields. The field sizes and descriptions are:

- 4 bit - Version
- 8 bit – Priority/Traffic class
- 20 bit – Flow/QoS management
- 16 bit – Payload length
- 8 bit – Next header
- 8 bit - TTL
- 128 bit - Source address
- 128 bit - Destination address

IPv6 Traffic Types

Within the IPv6 protocol, there are three types of address:

- **Unicast** – A one source to one destination mode of transport.
- **Multicast** – A one source to many destinations mode of transport.
- **Anycast** – A single IPv6 address is assigned to multiple devices. A client that requests information from this IPv6 address will communicate with the device closest to itself.

NOTE: IPv6 does not use broadcasts! Instead, broadcast traffic as we know it in IPv4 is more likely to use a multicast method in IPv6 transport.

IPv6 Routing Methods

Just like IPv4, IPv6 has both static and dynamic methods for routing to remote networks.

These methods include:

- Static IPv6 routes
- RIPng
- OSPFv3
- IS-IS for IPv6
- Multiprotocol BGP
- EIGRP configured for IPv6

The TSHOOT exam focuses mainly on being able to troubleshoot RIPng and OSPFv3 using show and debug commands.

IPv6 to IPv4 Communication

As stated earlier, the IPv6 headers are much different than IPv4 headers. That means that the protocols are inoperable. When transitioning from IPv4 to IPv6, you will need the two protocols to interact with one another in some fashion until a complete migration has occurred. There are two primary methods used when IPv6 and IPv4 devices can co-exist on the same network.

- **Dual Stack** – IP dual stack is a method for a device such as a router to run two IP protocols simultaneously. Communication between end devices is then determined by what protocol the devices can use. End devices can be running IPv4, IPv6 or both protocols. If all of the devices in the communication stream can talk IPv6, then this is the preferred method of transport. If one of the devices can only communicate using IPv4, the routing of packets falls back to this protocol.
- **IPv6 to IPv4 tunnel** – this is not so much a way for IPv4 to interoperate with IPv6 but rather a way to tunnel IPv6 addressed packets across an IPv4 network. If you are migrating your network and have two IPv6 portions sandwiched between an IPv4 network, you need a way to tunnel the IPv6 traffic across the IPv4 network. This is common when you have migrated several distribution blocks from v4 to v6 but the core routers remain at v4. Setting up a tunnel is very easy and involves creating tunnel interfaces on the edges of the IPv4 networks that create a complete tunnel through the IPv4 network.

IPv6 Address Format

Most network engineers today are well versed in the IPv4 addressing format. When it comes time to learn the IPv6 addressing format, many find the scheme to be cumbersome and overwhelming. This section will help to show some of the tricks for reading and simplifying IPv6 addresses.

An IPv6 address has an 8 octet hexadecimal format that are separated by the colon (:) symbol. Each 4 character section is a 16 bit block. $16 \times 8 = 128$. Below is an example of a full IPv6 address:

```
FE80:0000:0000:0000:0204:00CF:002C:8524
```

Pretty big right? At this point, many engineers start to panic and long for the days of their much more simplified IPv4 addressing scheme. Fortunately, there are tricks to trim down the length by collapsing the address by using the following two rules:

- Leading zeros within each 16 bit block can be omitted.
- Contiguous blocks that have all zeros can be represented with a double colon (::). This trick can only be performed one time to an address.

So looking back at our original IPv6 address, we can shrink it down to the following using the two rules we just learned:

```
FE80::204:CF:2C:8524
```

Removing all of those 0s in the address simplifies it and makes it much easier to read.

IPv6 OSPFv3 Troubleshooting Commands

This section covers the show and debug commands that should be used to troubleshoot OSPFv3 routing. You will notice that the output of these commands is similar to regular OSPFv2 that runs on IPv4. One possibly confusing aspect of OSPFv3 commands is that you'll notice the OSPF router ID's (RID) is in an IPv4 format. Remember however that the RID is not actually an IP address but rather a separate address that simply grabbed an IPv4 IP address out of convenience. In the case of OSPFv3, it will use the highest numbered loopback address or a better option is to manually configure a RID by using the **router-id** <RID> command when configuring OSPFv3.

Here are the show and debug commands that TSHOOT candidates must be familiar with:

Show ipv6 ospf

The show ipv6 ospf command details local OSPFv3 information such as the router's RID, any route redistribution methods, and timers, number and types of areas. Here is an example of this output:

```
Router#show ipv6 ospf
Routing Process "ospfv3 1" with ID 172.16.3.3
  It is an autonomous system boundary router
  Redistributing External Routes from,
    static
  SPF schedule delay 5 secs, Hold time between two SPFs 10 secs
  Minimum LSA interval 5 secs. Minimum LSA arrival 1 secs
  LSA group pacing timer 240 secs
  Interface flood pacing timer 33 msec
  Retransmission pacing timer 66 msec
  Number of external LSA 1. Checksum Sum 0x218D
  Number of areas in this router is 1. 1 normal 0 stub 0 nssa
    Area 1
      Number of interfaces in this area is 2
      SPF algorithm executed 9 times
      Number of LSA 15. Checksum Sum 0x67581
      Number of DCbitless LSA 0
      Number of indication LSA 0
      Number of DoNotAge LSA 0
      Flood list length 0
```

Show ipv6 ospf interface

The show ipv6 ospf interface command lists specific information on local interfaces configured to route OSPF traffic. The interfaces list the following important information:

- Interface status
- Network Type
- Interface cost
- Interface type (DR/BDR, POINT-TO-POINT, etc)
- Timers
- Designated router (if a broadcast interface)

Here is the output of this command:

```

Router#show ipv6 ospf interface
Ethernet1/0 is up, line protocol is up
  Link Local Address FE80::A8BB:CCFF:FE00:6601, Interface ID 6
  Area 0, Process ID 1, Instance ID 0, Router ID 10.0.0.1
  Network Type BROADCAST, Cost: 10
  Transmit Delay is 1 sec, State DR, Priority 1
  Designated Router (ID) 10.0.0.1, local address
FE80::A8BB:CCFF:FE00:6601
  No backup designated router on this network
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:09
  Index 1/1/1, flood queue length 0
  Next 0x0(0)/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)
Serial11/0 is up, line protocol is up
  Link Local Address FE80::A8BB:CCFF:FE00:6600, Interface ID 46
  Area 1, Process ID 1, Instance ID 0, Router ID 10.0.0.1
  Network Type POINT_TO_POINT, Cost: 64
  MD5 authentication (Area) SPI 500, secure socket UP (errors: 0)
  Transmit Delay is 1 sec, State POINT_TO_POINT,
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:09
  Index 1/1/2, flood queue length 0
  Next 0x0(0)/0x0(0)/0x0(0)
  Last flood scan length is 1, maximum is 5
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 10.0.0.1
  Suppress hello for 0 neighbor(s)

```

Show ipv6 ospf neighbor

You'll notice immediately that the show ipv6 ospf neighbor command is identical to the ipv4 cousin. Remember that OSPFv3 works nearly identically to v2 except for the packet header types. Remember, OSPFv3 still uses the 4 octet Neighbor ID to identify OSPFv3 neighbors and you cannot differentiate this from its older v2 show command brother. Here is the output of this command:

```

Router#show ipv6 ospf neighbor
Neighbor ID   Pri  State           Dead Time   Interface ID  Interface
10.1.1.1     1    FULL/DROTHER   00:00:30   3             FE0/1
10.1.3.1     1    FULL/BDR       00:00:31   3             FE0/1

```

Debug ipv6 ospf adj AND debug ip ipv6 ospf hello

When troubleshooting routing problems using debug commands, sometimes it is helpful to enable multiple debug commands at the same time. Troubleshooting ospf using debug ipv6 ospf adj and debug ipv6 ospf hello is a great example of how the output of both of these commands helps to understand what is happening with OSPFv3 in real time. The debug ipv6 ospf adj details the state of the OSPF neighbor relationship process flowing from the ATTEMPT to the FULL state. We also know that the movement of state changes between two neighboring routers is accomplished through the use of OSPFv3 hello packets. Therefore, it is nice to turn on the debug ipv6 ospf hello events to see the exchange of hellos between the neighbors to insure they are getting through. Here is an example of turning on both of these debug commands:

```

Router#debug ipv6 ospf adj
Router#debug ipv6 ospf hello
OSPFv3: Rcv hello from 10.1.1.10 area 1 from Serial0/0
FE50::202:F3FF:FE5A:E40
interface ID 4
OSPFv3: 2 Way Communication to 10.1.1.10 on Serial0/0, state 2WAY
OSPFv3: Neighbor change Event on interface Serial0/0
OSPFv3: DR/BDR election on Serial0/0
OSPFv3: Elect BDR 10.1.1.10
OSPFv3: Elect DR 10.1.1.10
          DR: 10.1.1.10 (Id) BDR: 10.1.1.10 (Id)
OSPFv3: Send DBD to 10.1.1.10 on Serial0/0 seq 0xF78 opt 0x0013 flag
0x7 len 28
OSPFv3: Rcv hello from 10.1.3.1 area 1 from Serial0/0
FE50::201:42FF:FA39:E500
interface ID 4
OSPFv3: 2 Way Communication to 10.1.3.1 on Serial0/0, state 2WAY
OSPFv3: Neighbor change Event on interface Serial0/0
OSPFv3: DR/BDR election on Serial0/0
OSPFv3: Elect BDR 10.1.1.10
OSPFv3: Elect DR 10.1.3.1
          DR: 10.1.3.1 (Id) BDR: 10.1.1.10 (Id)
OSPFv3: Send DBD to 10.1.3.1 on Serial0/0 seq 0x1C93 opt 0x0013 flag
0x7 len 28
OSPFv3: Remember old DR 10.1.1.10 (id)
OSPFv3: End of hello processing
OSPFv3: Send DBD to 10.1.1.10 on Serial0/0 seq 0xF78 opt 0x0013 flag
0x7 len 28
OSPFv3: Retransmitting DBD to 10.1.1.10 on Serial0/0 [1]
OSPFv3: Send DBD to 10.1.3.1 on Serial0/0 seq 0x1C93 opt 0x0013 flag
0x7 len 28
OSPFv3: Retransmitting DBD to 10.1.3.1 on Serial0/0 [1]
OSPFv3: Send DBD to 10.1.1.10 on Serial0/0 seq 0xF78 opt 0x0013 flag

```

```
0x7 len 28
OSPFv3: Retransmitting DBD to 10.1.1.10 on Serial0/0 [2]
OSPFv3: Rcv hello from 10.1.1.10 area 1 from Serial0/0
FE80::202:FDFE:FE5A:E40
interface ID 4
OSPFv3: End of hello processing
```

IPv6 RIPng Troubleshooting Commands

Much like OSPFv3, RIPng shares most of its routing behavior characteristics from RIPv2 which is the latest RIP protocol for routing IPv4. It still uses the Bellman-Ford algorithm and has a maximum hop-count of 15 layer 3 devices. That means that this protocol is a decent choice for small to medium sized networks only.

Below are show and debug commands displayed to give you an idea of what information can be used for troubleshooting purposes.

Show ipv6 rip

The show ipv6 rip command shows local information about the RIPng protocol and its settings. Some useful information includes:

- Port – RIPng runs over UDP 521 by default
- Timers – Update and Holddown
- Split horizon/poison reverse settings
- Local interfaces configured to run RIPng
- Any redistributed protocols

Here is the output of this command:

```
Router#show ipv6 rip
RIP process "process1", port 521, multicast-group FF01::A9, pid 62
  Administrative distance is 120. Maximum paths is 1
  Updates every 5 seconds, expire after 15
  Holddown lasts 10 seconds, garbage collect after 30
  Split horizon is on; poison reverse is off
  Default routes are generated
  Periodic updates 223, trigger updates 1
Interfaces:
  Ethernet0/0
Redistribution:
  Redistributing protocol bgp 65001 route-map bgp-to-rip
```

Show ipv6 route rip

The show ipv6 route rip command again should be very familiar except for the different structure of the address. Here's an example:

```
Router#show ipv6 route rip
IPv6 Routing Table - 17 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B - BGP
        U - Per-user Static route
        I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
        O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
R   2001:0db8:21::/32 [120/2]
    via FE80::A8BB:CCFF:FE00:A00, Ethernet0/0
R   2001:0db8:12::/32 [120/2]
    via FE80::A8BB:CCFF:FE00:A00, Ethernet0/0
R   2001:0db8:33::/32 [120/2]
    via FE80::A8BB:CCFF:FE00:A00, Ethernet0/0
```

Debug ipv6 rip

Finally, the debug ipv6 rip command displays sent and received information to and from RIPng neighbors. Remember that RIP sends a complete routing table to neighbors every 30 seconds by default. The output also contains information such as source and destination IP and port numbers as shown in this example:

```
Router#debug ipv6 rip
RIPng: Sending multicast update on Ethernet0/0 for process1
      src=FE80::A8BB:CCFF:FE00:B00
      dst=FF02::9 (Ethernet0/0)
      sport=521, dport=521, length=112
      command=2, version=1, mbz=0, #rte=5
      tag=0, metric=1, prefix=2001:0db8::/64
      tag=4, metric=1, prefix=2001:0db8:1::/16
      tag=4, metric=1, prefix=2001:0db8:2::/16
      tag=4, metric=1, prefix=2001:0db8:3::/16
      tag=0, metric=1, prefix=::/0
RIPng: Next RIB walk in 10032
RIPng: response received from FE80::A8BB:CCFF:FE00:A00 on Ethernet0/0
for process1
      src=FE80::A8BB:CCFF:FE00:A00 (Ethernet0/0)
      dst=FF02::9
      sport=521, dport=521, length=92
      command=2, version=1, mbz=0, #rte=4
      tag=0, metric=1, prefix=2001:0db8::/64
      tag=0, metric=1, prefix=2001:0db8:1::/32
      tag=0, metric=1, prefix=2001:0db8:2::/32
      tag=0, metric=1, prefix=2001:0db8:3::/32
```

Practice Questions

Chapter 1

1. Which well-known network maintenance model is a model defined by the ISO? Choose the best answer.
 - A. FCAPS
 - B. ITIL
 - C. TMN
 - D. Cisco Lifecycle Services

2. Which of the following is NOT considered to be a routine maintenance task? Choose the best answer.
 - A. Configuration changes
 - B. Network security breach
 - C. Monitoring network performance
 - D. Replacement of failed hardware

3. Which type of network diagram shows the interconnection of network segments, protocols used and how clients/servers interface with the network? Choose the best answer.
 - A. Physical topology diagram
 - B. Three-tiered network diagram
 - C. Logical topology diagram
 - D. Distributed topology diagram

4. A network engineer is investigating a trouble ticket and has found problems with a router configuration where changes have recently been made by another engineer. Instead of removing the changes, the engineer wants to roll-back and use a configuration that was archived onto an FTP server. To do this the engineer uses the following command: `configure replace ftp://192.168.1.199/R1-config-2`. What does this command do? Choose the best answer.
 - A. Replaces the current router startup-configuration with the archived configuration.
 - B. Replaces the current router running-configuration with the archived configuration.
 - C. Merges the current router startup-configuration with the archived configuration.
 - D. Merges the current router running-configuration with the archived configuration.

5. The “Shoot from the hip” troubleshooting approach typically skips over which two troubleshooting steps that a more structured approach would step through? Choose two.
- A. Problem reported
 - B. Examine information
 - C. Eliminate Potential Causes
 - D. Collect information
 - E. Hypothesize underlying cause

Chapter 2

1. A network administrator believes he has fixed a performance problem on a switch that was causing the input interface counter to increment. The administrator wants to monitor the port overnight to see if the problem has gone away. What command should be run to more easily verify that the problem has been fixed? Choose the best answer.
- A. Issue the terminal monitor command.
 - B. Issue the logging buffered command.
 - C. Issue the clear counters command.
 - D. Issue the clear ip route command.
2. What are the two SNMP community string types? Choose two.
- A. Read-only
 - B. Write-only
 - C. Append-only
 - D. Read-write
 - E. Read-mark
3. Which two show commands on a switch display the VLANs each switch port belongs to? Choose two.
- A. show ip arp
 - B. show interfaces trunk
 - C. show cdp neighbors
 - D. show vlan brief

4. A junior network engineer has configured two PCs in the lab that connect to a layer 2 switch. One PC is configured for VLAN 10 and the other is on VLAN 20. Each PC is in the same /24 network of 172.16.10.X. The engineer does not understand why he cannot ping from one PC to the other. What two things should you tell him? Choose two.
- A. A layer 3 interface needs to be configured to route between VLANs.
 - B. Devices that reside in separate VLANs must be on separate IP subnets.
 - C. A layer 2 interface needs to be configured to route between VLANs.
 - D. Devices that reside in separate VLANs must be on the same IP subnet.
5. At which plane is troubleshooting different when working on a layer 3 switch compared to a layer 3 router? Choose the best answer.
- A. Route plane
 - B. Data plane
 - C. Control plane
 - D. Switch plane
6. A network engineer is configuring an IP address on a FastEthernet port of a multilayer switch. Here is the result of that command. `MLS(config-if)#ip address 192.168.1.1 255.255.255.0%` IP addresses may not be configured on L2 links. Which of the statements below is the cause of the problem? Choose the best answer.
- A. The switch hardware does not support assigning IP addresses directly on physical interfaces. Instead, the engineer must configure SVIs and add layer 3 addressing to them.
 - B. The switch software does not support assigning IP addresses directly on physical interfaces. Instead, the engineer must configure SVIs and add layer 3 addressing to them.
 - C. The switchport is set for layer 2 mode. To make it a layer 3 physical port, run this command.
`MLS(config-if)# no switchport`
 - D. The switch is set for layer 2 mode. To make it a layer 3 switch, run this command
`MLS(config)# no switchport`

7. A standby router running HSRP sends and receives hello messages every 2 seconds. By default, how many hello messages can the standby router miss before moving from standby to active mode? Choose the best answer.
- A. 3 hellos
 - B. 5 hellos
 - C. 2 hellos
 - D. 10 hellos
8. A network engineer issues a show standby brief command and received the following output:
- ```
Router# show standby brief
Interface Grp Prio P State Active Standby Virtual IPFa0/0 0 100 P
Active local 192.168.0.11 192.168.0.1
```
- What does the 'P' indicate? Choose the best answer.
- A. If this router is the primary or standby switch. A 'P' means that it is currently primary.
  - B. If the router HSRP authentication is privately encrypted. A 'P' means that it is encrypted using an MD5 hash.
  - C. If the router HSRP configuration is set to preempt. A 'P' means that this router will preempt if it ever goes down and back up for some reason.
  - D. If the HSRP is configured using private IP address space. A 'P' means that it is using private IP space.
9. A frame enters a port on a layer 2 Catalyst switch, enters the \_\_\_\_\_ and exits out another port. Choose the best answer.
- A. Route-switch processor (RSP)
  - B. Backplane
  - C. CPU
  - D. ASIC
10. You run the following show command on a router configured for frame-relay. Router# show frame-relay map Serial1/0.2 (up): point-to-point dlci, dlci 111(0xB6,0x2C60), broadcast status defined, active Serial1/0.1 (up): point-to-point dlci, dlci 112(0xB5,0x2C50), broadcast status defined, active What does this command show? Choose the best answer.
- A. The DLCI to IP address mapping.
  - B. The frame-relay sub-interface to MAC address mapping.
  - C. The frame-relay sub-interface to DLCI address mapping.
  - D. The MAC address to DLCI mapping.

11. You are asked to verify the seed metric on the boundary router. What is a seed metric? Choose the best answer.
- A. The metric assigned by a network engineer that is to be injected into another routing protocol when redistributing routes.
  - B. The metric assigned by a network engineer that is to be injected into other routers in different OSPF areas.
  - C. The metric assigned by a network engineer that is to be injected into another routing protocol when redistributing distance vector routes.
  - D. The metric assigned by a network engineer that is to be injected into another routing protocol when redistributing link state routes.
12. How do two different routing protocols exchange routes when performing redistribution? Choose the best answer.
- A. By pulling routes out of the IP routing table that are placed there by the native routing protocol.
  - B. The two protocols place all of their routes into a separate routing information base (RIB). The protocol that is doing the redistribution into itself then uses its native algorithm to choose routes to be placed into the IP routing table.
  - C. The two protocols use a third protocol called CEF that created solely for the purpose of redistribution. CEF then uses its native algorithm to choose routes to be placed into the IP routing table.
  - D. It depends on which routing protocols are being redistributed. If they are both distance vector protocols, then the protocol that is doing the redistribution into itself pulls the foreign routes directly from the IP routing table. If they are any other combination, then the two protocols place all of their routes into a separate routing information base (RIB). The protocol that is doing the redistribution into itself then uses its native algorithm to choose routes to be placed into the IP routing table.
13. You are having problems with routes flapping on your router. What IOS configuration can be used to help monitor these by keeping track of the times that each flap occurs? Choose the best answer.
- A. ip route monitor
  - B. ip route triggered
  - C. ip route flap
  - D. ip route profile

14. A network engineer is troubleshooting a router running BGP. The router is in the BGP Established state with its peer. If everything is working properly, when running debug commands, which messages should the engineer see being sent and received? Choose two.
- A. Keepalive messages
  - B. Open messages
  - C. Notification messages
  - D. Update messages
15. You are troubleshooting a routing problem on your edge router running both EIGRP and iBGP. Both protocols know how to reach the same remote network. Which route will be put into the IP routing table? Choose the best answer.
- A. Because BGP is an Exterior Gateway protocol and EIGRP is an Interior Gateway protocol, both routes will be placed into the IP routing table.
  - B. The EIGRP learned route will be placed into the routing table.
  - C. The AD of iBGP is 20 while the AD of EIGRP is 90. The higher AD is preferred so the EIGRP route will be placed into the IP routing table.
  - D. The AD of iBGP is 20 while the AD of EIGRP is 90. The lower AD is preferred so the iBGP route will be placed into the IP routing table.
16. A network engineer is concerned about high CPU utilization on a router. The engineer thinks the cause might be an excessive number of TCP connections. What show command can be used to verify this theory? Choose the best answer.
- A. show processes cpu
  - B. show tcp statistics
  - C. show ip tcp buffers
  - D. show processes cpu history
17. How do Cisco IP phones learn what VLAN is the voice VLAN? Choose the best answer.
- A. As soon as the phone receives an IP address from the DHCP server, it sends a multicast to the Cisco Unified Communications Manager (CUCM). The response from the CUCM contains information about the Voice VLAN.
  - B. As soon as the phone receives an IP address from the DHCP server, it sends a multicast to the gateway. The response from the gateway contains information about the Voice VLAN.
  - C. The Voice VLAN is learned via CDP.
  - D. The VLAN is learned via DHCP.

18. An application on the network is using the multicast address of 239.10.10.10. What is the corresponding MAC address for the multicast address? Choose the best answer.
- A. 0100.5E0A.0A0A
  - B. 0100.5E10.1010
  - C. 010E.EF0A.0A0A
  - D. 0100.EF10.1010
19. What IPv6 OSPF show command displays IPv6 link local address, area ID, process ID, router ID, and cost? Choose the best answer.
- A. show ipv6 ospf neighbor
  - B. show ipv6 ospf interface
  - C. show ipv6 ospf
  - D. show ipv6 ospf processes
20. What are the two modes that wireless access points can operate in? Choose two.
- A. Independent mode.
  - B. Mobile mode
  - C. Split-MAC mode
  - D. Autonomous mode
  - E. Low-power mode

## Answers & Explanations

### Chapter 1

#### 1. Answer: A

**Explanation A.** Correct. FCAPS is a network maintenance model defined by the International Organization of Standardization (ISO).

Explanation B. Incorrect. IITL is not a network maintenance model defined by the International Organization of Standardization (ISO).

Explanation C. Incorrect. TMN is not a network maintenance model defined by the International Organization of Standardization (ISO).

Explanation D. Incorrect. The Cisco Lifecycle Services is not a network maintenance model defined by the International Organization of Standardization (ISO).

#### 2. Answer: B

Explanation A. Incorrect. Configuration changes will always be required on a network where people are constantly requiring additional functionality or improved security.

**Explanation B.** Correct. This is not a routine maintenance task.

Explanation C. Incorrect. The monitoring of network performance should be done on a regular basis.

Explanation D. Incorrect. Hardware no matter how reliable will fail and therefore is part of routine maintenance procedures.

#### 3. Answer: C

Explanation A. Incorrect. A physical topology shows cabling and hardware information only.

Explanation B. Incorrect. A three-tiered network diagram is another term used for a physical diagram.

**Explanation C.** Correct. The Logical topology diagram shows how the network functions from a protocol and configuration point of view.

Explanation D. Incorrect. There is no such thing as a distributed topology diagram.

#### 4. Answer: B

Explanation A. Incorrect. The command given does not replace the current startup-config with an archived configuration.

**Explanation B.** Correct. The command completely replaces the current running-config with the archived configuration specified.

Explanation C. Incorrect. The command given does not merge configurations.

Explanation D. Incorrect. The command given does not merge configurations.

## 5. Answers: B, C

Explanation A. Incorrect. The shoot from the hip approach does not skip this step.

**Explanation B.** Correct. The shoot from the hip troubleshooting approach often skips this step.

**Explanation C.** Correct. The shoot from the hip troubleshooting approach often skips this step.

Explanation D. Incorrect. The shoot from the hip approach does not skip this step.

Explanation E. Incorrect. The shoot from the hip approach does not skip this step.

## Chapter 2

### 1. Answer: C

Explanation A. Incorrect. The terminal monitor command enables network administrators to view console logging information when logged in via telnet or SSH. The command will not help in this situation.

Explanation B. Incorrect. The logging buffer command enables logged messages to be stored on the IOS device's memory. The command will not help in this situation.

**Explanation C.** Correct. The clear counters command is useful because it resets the input errors counter back to 0. If the administrator checks the counters the next morning and they're still at 0, then the duplex problem was certainly the cause.

Explanation D. Incorrect. The clear ip route command will clear out any dynamically learned routes on an IOS router. The command will not help in this situation.

### 2. Answers: A, D

**Explanation A.** Correct. The two SNMP community string types are read-only and read-write.

Explanation B. Incorrect. Write-only is not an SNMP community string type.

Explanation C. Incorrect. Append-only is not an SNMP community string type.

**Explanation D.** Correct. The two SNMP community string types are read-only and read-write.

Explanation E. Incorrect. Read-mark is not an SNMP community string type.

### 3. Answers: B, D

Explanation A. Incorrect. The show ip arp command will not display information about the VLANs each switchport belongs to.

**Explanation B.** Correct. The two commands are show interfaces trunk and show vlan brief.

Explanation C. Incorrect. The show cdp neighbors command will not display information about the VLANs each switchport belongs to.

**Explanation D.** Correct. The two commands are show interfaces trunk and show vlan brief.



#### 4. Answers: A, B

**Explanation A.** Correct. If the engineer wants the two PC's to be on separate VLANs, a layer 3 interface must be used as gateways for routing traffic from one VLAN to another.

**Explanation B.** Correct. If multiple VLANs are used, this is a broadcast segment. Because of this, a separate IP subnet must be used on each VLAN.

Explanation C. Incorrect. A layer 3 interface is needed.

Explanation D. Incorrect. This is the mistake the Junior engineer did. Instead, devices that are in different VLANs must be in different subnets.

#### 5. Answer: B

Explanation A. Incorrect. There is no such thing as the route plane.

**Explanation B.** Correct. Troubleshooting on a layer 3 switch vs. a layer 3 router is different when investigating the data plane.

Explanation C. Incorrect. Troubleshooting on a layer 3 switch vs. a layer 3 router is identical when troubleshooting the control plane.

Explanation D. Incorrect. There is no such thing as the switch plane.

#### 6. Answer: C

Explanation A. Incorrect. Cisco multilayer switch hardware allows you to configure IP addressing directly on the interface.

Explanation B. Incorrect. Cisco multilayer switch software allows you to configure IP addressing directly on the interface.

**Explanation C.** Correct. By default, multilayer switch physical ports are configured to be layer 2 switchports. To change this setting, you can configure the port with the no switchport command.

Explanation D. Incorrect. By default, multilayer switch physical ports are configured to be layer 2 switchports. To change this setting, you can configure the port with the no switchport command.

#### 7. Answer: B

Explanation A. Incorrect. The router can miss more than 3 hellos before HSRP will move an interface from standby to active mode.

**Explanation B.** Correct. The default is 5 hellos missed (10 seconds).

Explanation C. Incorrect. The router can miss more than 2 hellos before HSRP will move an interface from standby to active mode.

Explanation D. Incorrect. The router will miss fewer than 10 hellos before HSRP will move an interface from standby to active mode.

**8. Answer: C**

Explanation A. Incorrect. The 'P' does not stand for primary.

Explanation B. Incorrect. The 'P' does not mean that authentication is encrypted between two routers.

**Explanation C.** Correct. The 'P' stands for preempt.

Explanation D. Incorrect. The 'P' does not mean the HSRP interface is using private address space.

**9. Answer: B**

Explanation A. Incorrect. The frame will not be sent to an RSP.

**Explanation B.** Correct. All frames enter in one port, move across the backplane and are forwarded out another port onto its intended destination.

Explanation C. Incorrect. The frame will not be sent to the CPU.

Explanation D. Incorrect. The frame will not be sent to an ASIC.

**10. Answer: C**

Explanation A. Incorrect. This command does not show a table listing DLCI to IP address mappings.

Explanation B. Incorrect. This command does not show a table listing frame-relay subinterface to MAC address mappings.

**Explanation C.** Correct. This command shows the serial sub-interface to DLCI mapping.

Explanation D. Incorrect. This command does not show a table listing MAC address to DLCI mappings.

**11. Answer: A**

**Explanation A.** Correct. The seed metric is the default metric a routing protocol uses within itself when any other routes are redistributed into it.

Explanation B. Incorrect. The answer does not accurately describe a seed metric.

Explanation C. Incorrect. The answer does not accurately describe a seed metric.

Explanation D. Incorrect. The answer does not accurately describe a seed metric.

**12. Answer: A**

**Explanation A.** Correct. Routing protocols do not exchange routes directly between themselves. Instead the protocol that is doing the redistribution into itself pulls the foreign routes directly from the IP routing table.

Explanation B. Incorrect. The answer does not accurately describe how two different routing protocols exchange routes when performing redistribution. They do not combine all routes into a separate RIB.

Explanation C. Incorrect. The answer does not accurately describe how two different routing protocols exchange routes when performing redistribution. CEF is not used in redistribution.

Explanation D. Incorrect. The answer does not accurately describe how two different routing protocols exchange routes when performing redistribution.

**13. Answer: D**

Explanation A. Incorrect. This is not a valid IOS command.

Explanation B. Incorrect. This is not a valid IOS command.

Explanation C. Incorrect. This is not a valid IOS command.

**Explanation D.** Correct. This command helps to keep statistics on a router for occurrences of route flapping. These statistics include forward-path changes, next-hop changes, path count changes and others.

**14. Answers: A, D**

**Explanation A.** Correct. Keepalive messages are sent back and forth to peers to insure that the connection is still functional.

Explanation B. Incorrect. Once two peers are in the BGP Established state, open messages will no longer be exchanged.

Explanation C. Incorrect. If Notification messages are seen, there is some sort of problem and the BGP session will transition back to an Idle state.

**Explanation D.** Correct. Update messages are used to exchange routing information between peers to insure everything is up to date.

**15. Answer: B**

Explanation A. Incorrect. The fact that BGP is an EGP and EIGRP is an IGP does not factor into which route will be placed into the routing table.

**Explanation B.** Correct. EIGRP will have an administrative distance of 90 while iBGP will have an administrative distance of 200. The lower AD is more preferred so the EIGRP learned route will be placed into the IP routing table.

Explanation C. Incorrect. The Administrative Distance of iBGP is not correct. Also, the lower AD is the preferred protocol.

Explanation D. Incorrect. The Administrative Distance of iBGP is not correct.

**16. Answer: B**

Explanation A. Incorrect. This output of this command will not give you the information you are looking for.

**Explanation B.** Correct. This command show information about TCP connections including number of connections initiated, accepted, established, and closed.

Explanation C. Incorrect. This is not a valid IOS command.

Explanation D. Incorrect. This output of this command will not give you the information you are looking for.

**17. Answer: C**

Explanation A. Incorrect. This is not the correct way that IP phones learn what VLAN is the voice VLAN.

Explanation B. Incorrect. This is not the correct way that IP phones learn what VLAN is the voice VLAN.

**Explanation C.** Correct. CDP is used by Cisco phones to communicate to the local access switch. The switch then tells the phone which Voice VLAN it should be on.

Explanation D. Incorrect. This is not the correct way that IP phones learn what VLAN is the voice VLAN.

**18. Answer: A**

**Explanation A.** Correct. The first 6 digits of all multicast addresses are 0100.5E. The last 6 digits are the hex form of the last 3 octets of the address (10.10.10).

Explanation B. Incorrect. The first 6 digits of all multicast addresses are 0100.5E. The last 6 digits are the hex form of the last 3 octets of the address which does not turn out to be 10.1010.

Explanation C. Incorrect. The first 6 digits of all multicast addresses are 0100.5E.

Explanation D. Incorrect. The first 6 digits of all multicast addresses are 0100.5E.

**19. Answer: B**

Explanation A. Incorrect. This command will not display the IPv6 link local address, area ID, process ID, router ID, and cost.

**Explanation B.** Correct. The correct command is show ipv6 ospf interface.

Explanation C. Incorrect. This command will not display the IPv6 link local address, area ID, process ID, router ID, and cost.

Explanation D. Incorrect. This command will not display the IPv6 link local address, area ID, process ID, router ID, and cost.

**20. Answers: C, D**

Explanation A. Incorrect. The names for the two modes of operation for wireless access points are autonomous and split-MAC mode.

Explanation B. Incorrect. This is not a wireless AP mode.

**Explanation C.** Correct. This is one of the two wireless modes Cisco wireless access points can operate in.

**Explanation D.** Correct. This is one of the two wireless modes Cisco wireless access points can operate in.

Explanation E. Incorrect. This is not a wireless AP mode.