



CCNP TSHOOT 642-832

Quick Reference

Brent Stewart

Chapter 1	
Maintenance.....	3
Chapter 2	
Troubleshooting Methodology.....	16
Chapter 3	
Troubleshooting Tools.....	22
Chapter 4	
Troubleshooting Switches	43
Chapter 5	
Troubleshooting Routing	55
Chapter 6	
Troubleshooting Security Features	66

About the Author

Brent Stewart, CCNP, CCDP, CCSI, MCSE, is the manager of Connectivity Services at CommScope. He is responsible for designing and managing a large-scale worldwide voice, video, and data network. Previously he was a course director for Global Knowledge, participated in the development of BSCI with Cisco, and has written and taught extensively on CCNA and CCNP. Brent lives in Hickory, NC, with his beautiful wife, Karen, and their mischievous children Benjamin, Kaitlyn, Madelyn, and William.

About the Technical Editor

‘Rhette (Margaret) Marsh, CCIE No. 17476 Routing and Switching, CCNP, CCDP, CCNA, CCDA, CISSP, Marsh has been working in the networking and security industry for more than ten years and has extensive experience with internet-network design, IPv6, forensics, and greyhat work. She currently is a design consultant for Cisco in San Jose, CA, and works primarily with the Department of Defense and contractors. Prior to this, she worked extensively both in the financial industry as a routing and switching and design/security consultant and also in an attack attribution and forensics context. ‘Rhette is working toward her Security and Design CCIEs. In her copious free time, she enjoys number theory, arcane literature, cycling, hiking in the redwoods, sea kayaking, and her mellow cat, lexx.

Chapter 1

Maintenance

Maintenance might seem separate from the process of troubleshooting but imagine it as the other side of the same coin. Any device that is well maintained will be more reliable, suffers fewer problems, and will be easier and quicker to repair. Network owners, such as businesses and governments, want computer systems that are consistently available. Good troubleshooting technique minimizes the length of time of an outage, but good maintenance technique reduces outages.

You must select the appropriate tools and techniques for the network you maintain, based on law, company policy, and your experience. You need to understand, whichever elements you incorporate into your strategy, that a structured approach to maintenance is a key part of reducing unplanned outages.

Methodology

Network maintenance involves many different kinds of tasks, such as

- Installing new equipment
- Adjusting settings to support new service
- Securing the network
- Restoring service
- Backing up configs
- Planning new or upgraded service

NOTE:

TSHOOT doesn't assume a specific approach to maintenance.

Organizations might produce documentation and monitor their networks in unique ways. TSHOOT focuses on understanding the general practices that are used to successfully maintain a network.

- Building redundancy and disaster recovery
- Documentation
- Responding to user complaints

Many activities are reactive, and it is easy for interrupt-driven issues to monopolize your time. Defining a preventative maintenance schedule can help you avoid “firefighting.” Taking a more structured approach—as opposed for waiting for the phone to ring—can also help you recognize problems earlier and respond to them more efficiently. A broader perspective toward the network also provides an opportunity to align costs with the organization’s goals and budget effectively.

Several generic maintenance frameworks are available. Some organizations embrace a specific methodology, but many organizations pick, choose, and customize pieces that fit their environment. The important point is to have a documented approach to maintenance. If your organization doesn’t have a documented strategy, you might want to research some of these models.

- IT Infrastructure Library (ITIL)
- FCAPS
- Telecommunications Management Network (TMN)
- Cisco Lifecycle Services/PPDIOO
- Microsoft Operations Framework

After you choose a specific model, map the model onto processes you can use to maintain the network and then select the tools that you use.

Common Tasks

Although organizations that own networks have different expectations, the management of every network still includes some basic components. Planning and accomplishing these tasks repetitively and competently is a key to successful network management.

Some common tasks include

- Adds, moves, and changes
- Compiling documentation
- Preparing for disaster
- Capacity planning/utilization monitoring
- Troubleshooting
- Proactive scheduled maintenance
- Rollback plans for each change
- Lab testing in a controlled environment before each change is put into production to minimize risk

Preventative maintenance is the process of anticipating potential sources of failure and dealing with the problem before it occurs. It is probably not possible to anticipate every source of failure, but careful thought might help you identify candidates. One technique to identify issues is to look at prior records of trouble, such as trouble tickets, ISP records, network monitoring systems, or purchase records. Use this information to categorize and rank the experience of your network.

Organizations are typically willing to accept small periods of scheduled downtime to offset the probability of long periods of unscheduled downtime. Using the data collected from your experience, consider the steps that can be taken during this window of time. Operating systems can be patched or upgraded to more stable and secure versions.

Redundancy can be tested to ensure smooth failover. Additionally, normal business changes (such as new circuits) can be accomplished during this period to minimize disruption.

Most large organizations use a system of change controls to enforce a thought-out approach to configuration changes. Change control involves producing a document that describes the change to be made, who will make it, when the change will be made, and who will be affected. A well-written change control document will also have some notes about how the new configuration can be “backed out” if something goes wrong. This change control is then approved by management.

Change control systems help the business balance the need to update network components and configurations against the risk of changes. Change control systems also protect the network administrator—if each change is well thought out and thoroughly communicated, the business has the opportunity to accept the risks inherent in change.

Documentation reduces troubleshooting time and smoothes project communication as networks are changed and upgraded. Although time consuming, it is impossible to over emphasize the importance of accurate and up-to-date documentation. Well-maintained documentation includes details such as

- Configuration templates or standards
- Configuration history
- Equipment inventory (including serial number and support contract information)
- Circuit inventory (including circuit ID and service provider contact)
- IP address assignment
- Network drawings
- Communication plan
- Out-of-band communication details
- Expected traffic patterns

Templates can be a fill-in-the-blanks version of a complete configuration or can be *snippets* that show how your organization handles specific issues, such as IPsec tunnels. Either way, templates provide an opportunity for consistency and enable technicians to more quickly move from interpreting to troubleshooting. Consider, for instance, access-lists and how easily they might be confused. Access-list 100 might be typically related to permitting SNMP to certain destinations but on some devices is used to filtering traffic on the public interface. Understanding the ramifications of confusion in this example, it is easy to see the benefit of standardizing things such as labels. (And in this case, it is probably best to use named access-lists, not numbered.)

The documentation for the communication plan should include contact information for internal IT and management contacts, and vendor and service provider information. The plan should also specify who should be contacted, in what circumstances, and how often. For instance, should a technician update the business contract or the Network Operations Center? Is there a proscribed after-action review?

Often the individual documentation elements are combined, such as IP addresses and circuit IDs on the Network diagram, or simplified, such as a TFTP server directory to keep configuration history.

Documentation should also include a disaster recovery plan. Disasters come in many sizes, so it pays to consider several cases. If the problem is related to a single piece of equipment, consider Cisco SmartNet maintenance as a way to guarantee backup hardware is onsite quickly. Even in the case where a spare is procured, you need a backup of the configuration and IOS. If getting a spare involves a service contract, you probably also need the serial number. Someone onsite needs a console cable and a laptop with a serial port. Larger disasters, such as a fire, might require replacing equipment from memory. It's a good idea to also have a record of the installed cards and licenses. Finally, consider the staff at the site. Is there someone there who can be talked through copying a config or do you need a technician to go to the site?

A final common piece to managing the network is to have some form of network monitoring. Network monitors take many forms, from simple no-frills systems to complex central management. These systems are available from a variety of vendors and through open source. Regardless of which system you use, you need to pull data showing utilization, availability, performance, and errors. The system should alert the staff through emails or SMS messages so that you are aware

of problems before the phone rings.

After the monitoring system is in place, you need to periodically characterize performance as a snapshot. A *snapshot* describes the expected performance of a system and enables you to compare later performance and recognize change. For instance, changes in jitter or in dropped packets might indicate that a WAN link is oversubscribed. In addition, a functional baseline for performance metrics serves as a critical diagnostic tool for security breaches and zero-day attacks and worms. Without thorough knowledge of typical behavior on a given network, aberrant traffic analyses become a subjective art.

Tools

Most network administrators have a variety of tools in their toolbox. Some of the basic tools include a configuration history, device logs, and documentation. As the number of devices maintained grows, tools that collect data about the performance of the network and tools that collect user issues become increasingly important.

Configurations

A configuration history is built by saving the device configuration to a central point periodically or after each change. IOS supports a variety of different remote targets. FTP and TFTP are commonly used because implementations are bundled with many operating systems, and free open-source versions are readily available.

Blackburn-rtr01#copy run ?

```
archive:      Copy to archive: file system
flash:        Copy to flash: file system
ftp:          Copy to ftp: file system
http:         Copy to http: file system
https:        Copy to https: file system
```



```

idconf          Load an IDConf configuration file
null:          Copy to null: file system
nvram:         Copy to nvram: file system
pram:          Copy to pram: file system
rcp:           Copy to rcp: file system
running-config Update (merge with) current system configuration
scp:           Copy to scp: file system
slot0:         Copy to slot0: file system
startup-config Copy to startup configuration
syslog:        Copy to syslog: file system
system:        Copy to system: file system
tftp:          Copy to tftp: file system
tmpsys:        Copy to tmpsys: file system
xmodem:        Copy to xmodem: file system
ymodem:        Copy to ymodem: file system

```

One way to build a configuration history is to save your configuration after each change. Saving the file with the date attached makes it easy to sort later, and adding a .txt makes it easy for Windows-based machines to open the file. In the following example, the TFTP server has a directory for each site and the configuration is saved with the date:

```

Blackburn-rtr01#copy run tftp
Address or name of remote host []? 192.168.255.10
Destination filename [blackburn-rtr01-config]? blackburn/blackburn-rtr01-09-08-25.txt
!!
820 bytes copied in 2.628 secs (312 bytes/sec)

```

Logging events and alerts to Syslog is another important tool. Syslog is a facility that receives alerts from network equipment and stores them in a common log. Again, many version of syslog are available. Events are logged based on a severity scale, from zero to seven. Choosing a logging level tells the router to transmit events at that level and lower. To set up

syslog support on an IOS device, the logging keyword is used, as shown here:

```
Blackburn-rtr01(config)#logging trap ?
<0-7>          Logging severity level
alerts         Immediate action needed          (severity=1)
critical       Critical conditions             (severity=2)
debugging      Debugging messages                    (severity=7)
emergencies    System is unusable                     (severity=0)
errors         Error conditions                       (severity=3)
informational  Informational messages                 (severity=6)
notifications  Normal but significant conditions      (severity=5)
warnings      Warning conditions                     (severity=4)
<cr>

Blackburn-rtr01(config)#logging on
Blackburn-rtr01(config)#logging 192.168.255.10
Blackburn-rtr01(config)#logging trap informational
```

As the rate of log entries grows (because there are more devices or because the sensitivity is changed), finding the appropriate information in the logs becomes more cumbersome. One way to make it easier to tie events together in the log is to have accurate time on each device so that log entries have a consistent time. Time stamps become vital in forensics and post mortems, where sequence and patterns of events evolve into chains of evidence.

Time is synchronized on network devices using the network time protocol (NTP). Setting up NTP is straightforward; specify the NTP server with the command **ntp server** *<ip address>*. Time servers are organized by *stratums*, where stratum 1 clocks are super precise atomic clocks, stratum 2 devices get their time from stratum 1, stratum 3 devices ask stratum 2, and so on. Public stratum-1 devices are listed on the Internet; it is considered a courtesy that each organization has a minimal number of connections to a stratum-1 device and that other clocks in the organization pull from these stratum-2 devices.

Another time-related logging issue to consider is time zone. Will your organization log using local time zones, the time zone of headquarters, or set all devices to GMT? The following example demonstrates the time zone set to GMT, logging set, and the router set to use a remote NTP server:

```
service timestamps debug datetime msec localtime
service timestamps log datetime msec localtime
ntp server 192.168.1.1
clock timezone GMT 0 0
service timestamps debug datetime msec localtime
service timestamps log datetime msec localtime
```

Cisco IOS supports an Archive and Restore feature that makes maintaining a configuration history and logs easier. The archive function maintains a current copy of the configuration and a set of previous configurations. The archive can be maintained within the router or at an accessible URL. The restore function enables the router to smoothly revert to any of the saved configurations.

Setting up the archive function involves going into the archive configuration mode. The path command specifies a backup location, and time-period is used to periodically backup the configuration. If write-memory is specified, an archive copy will be made whenever the configuration is saved. Archive copies have a version number, such as “-1” on the end. This version number is reset with each router reset, so it would be hard to use this as a long-term archive. The path can include \$h for the hostname and \$t for time, so it is possible to time stamp each saved file. Using the time stamp is impractical with a Windows TFTP server, however, because the time stamp includes colons. In the next example the filename is *host-name.txt* and results in Blackburn-rtr01 saving files named Blackburn-rtr01.txt-1 and Blackburn-rtr01.txt-2. The example is set to back up at the maximum periodic interval, so most backups happen because the administrator saves the configuration:

```
archive
path tftp://192.168.255.10/$h.txt
write-memory
time-period 525600
```

The router uses a standard name structure for all saved files, counting up to 14 and then cycling back to 1. This is hard to use as a complete configuration history. One possible solution is to save the archive to flash and to have administrators save to TFTP periodically (which automatically updates the flash archive). The periodic backup could be set to run once a week, just in case someone forgot to “copy run start”:

```
archive
  path flash://$h
  write-memory
  time-period 10080
```

Archive can help troubleshoot in two ways. First, archive can compare differences between different versions of the config: archive config differences. Second, Archive can also be used to supplement syslog with all commands executed on the router. In archive configuration mode, enter log config mode. **logging enable** turns on command capture; **hidekeys** prevents logging passwords. Normally the log of commands is kept in memory on the router, but **Notify syslog** exports the commands to syslog. This configuration is shown here:

```
archive
  path flash://$h
  write-memory
  time-period 10080
  log config
    logging enable
    hidekeys
    notify syslog
```

To review the archive files, use the command **show archive**:

```
Blackburn-rtr01#show archive
```

The next archive file will be named tftp://192.168.255.10/Blackburn-rtr01-7

```

Archive #  Name
0
1      tftp://192.168.255.10/Blackburn-rtr01-1
2      tftp://192.168.255.10/Blackburn-rtr01-2
3      tftp://192.168.255.10/Blackburn-rtr01-3
4      tftp://192.168.255.10/Blackburn-rtr01-4
5      tftp://192.168.255.10/Blackburn-rtr01-5
6      tftp://192.168.255.10/Blackburn-rtr01-6 <- Most Recent
7
8
9
10
11
12
13
14

```

Finally, the archiving function adds the ability to restore to a previous configuration. Replacing an old configuration with **copy tftp run** results in the tftp file being merged into the running configuration whereas **copy tftp start** results in a complete replacement but requires a restart.

An archive can be restored with the **configure replace** command. The router compares the running configuration against the archive and builds and applies a list of commands necessary to match the archive. This method avoids reapplying existing commands or rebooting to make the migration:

```

Router#configure replace tftp://192.168.255.10/blackburn-rtr01-5
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 blackburn-rtr01-5 from 192.168.255.10 (via FastEthernet0/0): !
```

One trick when working with a remote router is to use “**reload in 5**” to schedule a reload. If a command inadvertently breaks the connection, the router reboots to the last saved configuration. If everything works, reload cancel prevents the reboot. The same functionality is available with **configure replace filename time** but avoids the reboot. Avoid the roll-back by confirming the change is working with **configure confirm**.

Other Tools

Documentation is a huge part of troubleshooting, and there are many tools that you can use to compile documentation. One of the key things to understand about documentation is that it must be easy and quick to update, or it will quickly grow stale. Microsoft Visio is a common way to show connectivity. A database or spreadsheet is frequently used to track inventory. You can use a ticketing system to list issues and gather trending data. Wikis are a more recent innovation that enables the network staff to produce and edit documentation.

There isn't a definitive way to produce documentation; the important part is to have documentation that is useful in the troubleshooting process. Ideally, the documentation should also feed directly into the disaster recovery process as well, so it should include part numbers, serial numbers, service contracts, and a variety of information that isn't strictly part of the “network” description.

Cisco has a variety of web-based tools that are helpful. The Dynamic Configuration tool is useful in planning hardware configurations; this tool can verify compatibility and build a parts list to help you plan a project. The Feature Navigator verifies that a specific feature is in a particular version of IOS. The Power Calculator calculates the required power supply for PoE installations. Many other tools are available through CCO, so it's worth spending some time understanding the width of the offering.

A final category of tools to consider are the network performance monitoring tools. Typically, monitoring and performance tracking in a small organization is accomplished with a phone—people call when they have problems. As an organization grows, however, it becomes more and more important to recognize problems before they occur. This same information can be used to budget hardware and circuit upgrades.

Monitoring tools typically use SNMP, Netflow, pings, and Syslog data to compile statistics about the current and historical behavior of the network. Typically, networks are monitored for capacity usage, availability, delay, and CPU and memory utilization. Solarwinds, nGenius, OPNET Net Doctor, SP Guru, and WhatsUpGold each make products that fulfill these functions, and MRTG is a similar open source project.

Remember to plan a monitoring system around the service level agreements (SLA) in the environment. Service providers typically offer some performance guarantees, such as minimizing unplanned downtime or minimizing jitter. The business might insist that IT also support SLAs internally. The Network monitoring system should provide information to back up both types of SLAs. Cisco has built in a SLA monitoring tool that can make availability statistics known and monitored for critical links and servers. This is called SLA Monitor and is customizable for MPLS, link utilization, RTT, and others. It is quite useful for critical traffic real-time monitoring and notification. Frequently these statistics are run as a continuous background process between CE nodes between sites, if remote connectivity between critical traffic endpoints is a business driver.

Chapter 2

Troubleshooting Methodology

The responsibilities of a network administrator boil down to four essential measurements: Maximize performance and availability; minimize cost, and time-to-repair.

This chapter focuses on minimizing time-to-repair. The time it takes to restore functionality is predicated on two things: preparation and technique. The previous chapter spoke about the elements of preparation, such as documentation and scheduled preventative maintenance. This chapter focuses on the techniques that you can apply to minimize downtime.

Each of the troubleshooting practices described in this chapter assume that good documentation exists and that appropriate tools are available. Troubleshooting is much more frustrating and time consuming when the necessary preparation isn't accomplished.

Principles

The scientific method is commonly described as a six-step process:

1. Define the problem.
2. Gather information.
3. Hypothesize.
4. Test hypothesis.

NOTE:

Cisco's Troubleshooting test doesn't assume a specific approach. Many approaches and different approaches might be successful in specific situations. The test does advocate a structured approach to troubleshooting, based on the scientific method.

Troubleshooting Methodology

5. Analyze test.
6. Interpret results and, if necessary, generate a new hypothesis.

The first step—problem description—is usually accomplished when a user reports a problem. The initial problem description tends to be vague or overly general (“The Internet is down!”). A troubleshooter’s initial response should therefore be to gather more information and build a more specific description. You can determine symptoms by talking to the user, by personal observation, or by referring to management systems such as Netflow, Syslog, and SNMP monitors.

When you have an adequate description of the problem, you can form a hypothesis. A hypothesis is a hypothetical potential problem whose symptoms would be similar. The hypothesis should commonly suggest a way to prove or disprove itself. For instance, if you suspect that the WAN connection is down, looking at the interface status or pinging a remote device would test that theory.

Test results will either support or refute a theory. A single test result can’t prove a theory but just support it. For example, ping might be used to test a WAN connection. A ping timeout cannot, by itself, be considered definitive. The target might be shut down or have a firewall that drops ICMP. Test results should be confirmed through a number of different lines of evidence. If the tests contradict the hypothesis, start over with a new theory.

After a hypothesis is accepted as a reasonable explanation, you can take action to fix the problem. Of course, any action is another type of test. If the action doesn’t fix the problem, simply develop a new hypothesis and repeat the process.

Structured Troubleshooting

The term *structured troubleshooting* describes any systematic way of collecting information, forming a hypothesis, and testing. In a structured approach, each unsuccessful test rules out entire classes of possible solutions and gracefully suggests the next hypothesis. An unstructured—random—approach usually takes much longer and is less likely to be successful.

Troubleshooting Methodology

A number of techniques have been used successfully, their common feature being a rigorous and thoughtful approach that collects data and analyzes data:

- **Top down:** Start at the OSI application layer and drill down.
- **Bottom up:** Start with the OSI physical layer and work up.
- **Divide and conquer:** Start at the network layer and follow the evidence, developing specific tests of each hypothesis.
- **Follow path:** Consider the “packets perspective” and examine the devices and processes it encounters moving through the network. Understand the order of operations within each device to do this.
- **Spot difference:** Compare the configuration to an older version or to that of a similar device. Diff and WinDiff are tools that make this comparison easy.

FIGURE 2-1
The OSI Model

7	Application
6	Presentation
5	Session
4	Transport
3	Network
2	Data Link
1	Physical

- **Move the problem:** Swap components to see if the problem moves with a device.

There isn't a single “best method,” although a given technician might find one more intuitive or more suitable for a given problem. It's a good idea to be familiar with each technique and to change approaches if necessary.

Two troubleshooting tactics need special mention. Most technicians build up a reservoir of experience, which gives them an intuition about the solution to a given problem. This can be incredibly impressive when it works; the trick is to not let this become a series of random stabs when it doesn't work.

Troubleshooting Methodology

Networkers also look for things that happened at about the same time, on the theory that the similar timing implies causation. This thinking is a logical error: *post hoc ergo proctor hoc*. Sometimes this does provide a clue, but large networks have many things happening contemporaneously every second. This troubleshooting method can easily provide a false lead.

The Troubleshooting Method

Troubleshooting a network falls into a series of steps that mirror the scientific method.

The first step in troubleshooting is to define the problem. Some users, for instance, might report that “The Internet is down,” when what they mean is “My e-mail is taking a long time to download.” Some users over-generalize or exaggerate for effect, but most users lack the technical sophistication to tell which symptoms are relevant. Always start the troubleshooting process by gathering a detailed description of the problem. Ask questions to gather details, such as the names and locations of affected devices. One good way to gather details is to ask about how the problem can be duplicated. (“So, if I browse the web I’ll see this problem?”)

After defining the problem, gather information about the problem. What is the scope? What other devices or locations are affected? When did it start? How can you test the problem?

As information is gathered, one or more theories might begin to form. Develop tests that confirm or refute the theories, and work to find the root cause. Tests can be as simple as pings or as complex as implementing a configuration change; the tests should be aimed at separating valid theories.

When the testing process is complete, take a moment to consider the results. Do the results suggest a configuration or hardware change? Is the problem resolved? If not, reconsider the problem description and the original hypothesis. Either the problem was not completely and accurately described, or the hypothesis was incorrect and needs to be revisited.

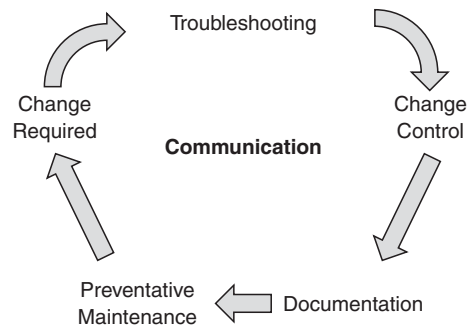
Troubleshooting Methodology

When the problem is resolved, take some time to consider the changes. The state of the network and the problem resolution need to be communicated, and documentation might need to be updated. Past these obvious steps, consider whether the problem found can be in other parts of the network. If the problem were in the configuration, think through the configuration template used in your network and determine if the fix needs to be repeated preemptively on other devices.

Each organization has its own specific methods for working through the break/fix cycle. The important points here are to work logically and methodically, and to view each problem as an opportunity to perfect the larger network.

Integrating Troubleshooting into Maintenance

Every interaction with the network is an opportunity to learn. Smart organizations capture information learned to solve similar problems and to help understand the network in the future. Change control and documentation are the two principal ways that feedback from network changes is incorporated into the maintenance cycle, as shown in Figure 2-2.



Preventative maintenance is ongoing, but changing conditions or reported problems create the need to make a change. Troubleshooting identifies the corrective action to upgrade or repair the network. Throughout these processes, a regular communication with end users is critical to understand the problem and to gather feedback on the solution. Communication with end users, within the team, and with management is pervasive throughout the cycle.

FIGURE 2-2
Maintenance Cycle

Troubleshooting Methodology

Change control is a process found in many organizations with large networks. The change-control process is a formal communication process for requesting and receiving permission. Change control provides an opportunity for management and peers to be aware and consent to the proposed change. The change process encourages the network technician to take a deliberate and thoughtful approach. Finally, the change process creates a record of the change that can be incorporated in documentation.

After a change is made and an issue is resolved, updating documentation must be seen as a part of the clean-up process. Most organizations have records including IPs, inventory, configurations, and topology; changes need to be added to these records. If the change is sufficiently broad, it might also need to be incorporated into standards and templates so that other devices can be preemptively upgraded. As records and standards change, team members need to be educated on the changes.

A *baseline* is a reading of the critical parameters of the network (such as latency and utilization) over a period of time. The baseline serves as a record of normal behavior to help identify how performance has changed. Updating baseline information is part of the documentation process.

Think about troubleshooting as a holistic process. Approach each issue with a rational evidence-based philosophy, make thoughtful changes, and communicate with all the invested groups often.

NOTE:

A number of tools can compile baseline data and monitor the network continuously. Cisco Works, HP Openview, What's Up? and SolarWinds are examples of commercial applications. Cacti and MRTG are two well-known Open Source versions.

Chapter 3

Troubleshooting Tools

Cisco IOS has a number of ways to extract data about the state of the machine. Understanding the capabilities of the operating system and how to use them effectively can reduce time-to-repair and the stress of a network outage.

IOS Filtering Tools

Most of the commands for pulling information from a router are familiar to anyone with Cisco IOS experience. Many people are not familiar with the filtering techniques that enable a troubleshooter to quickly focus.

Some of these filters are command-specific. Consider **show ip route**, which is a familiar command. When used, this command shows a complete routing table (as shown here):

```
Foard-rtr01#show ip route
```

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
```

```
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
```

```
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
```

```
       E1 - OSPF external type 1, E2 - OSPF external type 2
```

```
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
```

```
       ia - IS-IS inter area, * - candidate default, U - per-user static route
```

```
       o - ODR, P - periodic downloaded static route
```

```
Gateway of last resort is 10.100.1.1 to network 0.0.0.0
```

Troubleshooting Tools

```

172.0.0.0/8 is variably subnetted, 7 subnets, 2 masks
B    172.136.157.12/30 [20/0] via 10.1.254.246, 2d01h
S    172.99.120.2/31 [1/0] via 10.100.254.240
B    172.139.78.232/30 [20/0] via 10.1.254.246, 2d01h
B    172.136.88.20/30 [20/0] via 172.176.128.25, 5w2d
B    172.136.41.104/30 [20/0] via 172.176.128.25, 5w2d
B    172.137.230.128/30 [20/0] via 172.176.128.25, 6d18h
B    172.139.83.100/30 [20/0] via 172.176.128.25, 1w5d
172.16.0.0/32 is subnetted, 1 subnets
S    172.16.201.141 [1/0] via 10.100.254.240
192.168.0.0/30 is subnetted, 6 subnets
B    192.168.26.52 [20/0] via 10.1.254.246, 2d01h
B    192.168.241.236 [20/0] via 172.176.128.25, 5w2d
...

```

The output for this command can continue over many pages of information. One way to summarize this information is to ask for a summary using **show ip route summary**.

```

Foard-rtr01#show ip route summary
IP routing table name is Default-IP-Routing-Table(0)
IP routing table maximum-paths is 32
Route Source    Networks    Subnets    Overhead    Memory (bytes)
connected       0           19          1216        2888
static          4           22          1664        3952
bgp 65100       19          385         25856       62428
  External: 382 Internal: 22 Local: 0
internal        45          0           0           52740
Total           68          426         28736       122008
Removing Queue Size 0

```

Troubleshooting Tools

A second routing table filtering option is to ask for a selection of routes. Specifying an address, mask, and the keyword **longer-prefixes** asks for anything that matches the prefix or any routes contained within the prefix. The following example shows all the more-specific routes contained within the 10.1.254.0/24 block:

```
Foard-rtr01#show ip route 10.1.254.0 255.255.255.0 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 10.100.254.240 to network 0.0.0.0
```

```

    10.0.0.0/8 is variably subnetted, 241 subnets, 12 masks
C       10.1.254.244/30 is directly connected, Multilink31
C       10.1.254.246/32 is directly connected, Multilink31
B       10.1.254.252/30 [200/0] via 10.100.1.2, 1d09h
C       10.1.254.232/30 is directly connected, Multilink42
C       10.1.254.234/32 is directly connected, Multilink42
```

The options for filtering available for a given **show** command vary, so it's a good idea to spend some time with the question mark and understand the options available for areas of focus in your organization.

Generic filters can also be applied to all **show** commands. **Show process cpu**, which might be used to look for runaway processes, can be used as an example. First, an example portion of output is shown:

```
Foard-rtr01#show process cpu
CPU utilization for five seconds: 14%/13%; one minute: 14%; five minutes: 14%
```


Troubleshooting Tools

```

PID Runtime(ms)  Invoked    uSecs   5Sec   1Min   5Min TTY Process
  1         292     6405      45 0.00% 0.00% 0.00% 0 Chunk Manager
  2         296    639947      0 0.00% 0.00% 0.00% 0 Load Meter
  3          0         1      0 0.00% 0.00% 0.00% 0 chkpt message ha
  4          0         1      0 0.00% 0.00% 0.00% 0 EDDRI_MAIN
  5    1600592    326740   4898 0.00% 0.04% 0.00% 0 Check heaps
  6         2016    28869      69 0.00% 0.00% 0.00% 0 Pool Manager
  7          0         2      0 0.00% 0.00% 0.00% 0 Timers
  8          0         2      0 0.00% 0.00% 0.00% 0 ATM AutoVC Perio
  9          0         2      0 0.00% 0.00% 0.00% 0 ATM VC Auto Crea
 10          0    53330      0 0.00% 0.00% 0.00% 0 IPC Dynamic Cach
 11          0         1      0 0.00% 0.00% 0.00% 0 IPC Zone Manager
 12          20   3199682      0 0.00% 0.00% 0.00% 0 IPC Periodic Tim
 13          12   3199682      0 0.00% 0.00% 0.00% 0 IPC Deferred Por
 14          0         4      0 0.00% 0.00% 0.00% 0 IPC Seat Manager
 15          0         1      0 0.00% 0.00% 0.00% 0 IPC BackPressure
 16    387428  57716731      6 0.07% 0.01% 0.00% 0 EnvMon

```

...

The pipe (|) character is used to filter output by passing it through logic such as include, exclude, begin, and section. Output is matched against a regular expression.

Following is a table of common regular expression characters.

Character	Usage	Example
^	Begins with	^Fast matches lines that begin with FastEthernet.
\$	Ends with	FastEthernet0/0\$ matches lines that end with FastEthernet0/0.

Troubleshooting Tools

.	Any character	Ethernet./ matches Ethernet 0/0, FastEthernet0/1, and Ethernet ?.
	Or	FastEthernet 0/0 1 matches either FastEthernet0/0 and FastEthernet0/1.
_	Matches beginning, end, or braces	_Ethernet_ matches any line that includes the word “Ethernet.”

A show command piped to include will display any line of output that matches the regular expression. In the following example, the pipe is used to look for any line that includes the text “IP Input”.

```
Foard-rtr01#show process cpu | include IP Input
 87      2755292  47045037          58  0.07%  0.07%  0.07%  0 IP Input
```

The running configuration is another place to see piping work. In the following example, piping to begin starts the output at the telnet ports. This is a lot easier than using the space key to work through a large configuration:

```
Foard-rtr01#show running-configuration | begin vty
line vty 0 4
  exec-timeout 20 0
  password 7 0401001C02010D4106
  logging synchronous
  transport input ssh
  transport output telnet ssh
line vty 5 15
  exec-timeout 20 0
  password 7 0401001C02010D4106
  logging synchronous
  transport input ssh
  transport output telnet ssh
!
ntp source Loopback0
```

Troubleshooting Tools

```
ntp master 5
ntp update-calendar
ntp server 172.31.55.2
ntp peer 10.1.1.123 key 1 source Loopback0
end
```

In the preceding example, piping to `begin` also includes all the text after the part of interest. Piping to `section` shows the indented commands under a line that matches the regular expression. In the following example, the sections found under the keyword `vtty` are shown:

```
Foard-rtr01#show running-config | section vty
line vty 0 4
  exec-timeout 20 0
  password 7 045C021302284D4906
  logging synchronous
  transport input ssh
  transport output telnet ssh
line vty 5 15
  exec-timeout 20 0
  password 7 14101B1E010D2B2C2B
  logging synchronous
  transport input ssh
  transport output telnet ssh
```

The pipe symbol is also used as an OR within a regular expression, as shown in the next examples. Normally, `show ip interface brief` summarizes all the interfaces found on a router. Some routers have a large number of interfaces, making even this simplified display cumbersome. In the following text, some of the interfaces are grouped into multilinks and others are turned off. Finding the detail you need is complicated by the long and confusing output:

```
Foard-rtr01#show ip interface brief
Interface                IP-Address      OK? Method Status          Protocol
FastEthernet0/0          10.87.1.1       YES NVRAM  up              up
```

NOTE:

Piping output can be a great way to focus on relevant details, but `show running-configuration | section` is a lot to type, particularly repeatedly. The `alias` command can make this easier. In configuration mode, create a shortened version of a command as shown next.

```
rtr01(config)#alias
exec srs show
running-configuration
| section
```

Now “srs” is the shortened version of the long and cumbersome command. Type `srs vty` to see the same output as the example.

Troubleshooting Tools

FastEthernet0/0.2	10.76.2.2	YES NVRAM	up	up
FastEthernet0/0.3	10.76.3.2	YES NVRAM	up	up
FastEthernet0/0.4	10.76.4.2	YES NVRAM	up	up
FastEthernet0/0.5	10.76.5.2	YES NVRAM	up	up
FastEthernet0/0.6	10.76.6.2	YES NVRAM	up	up
FastEthernet0/0.7	10.76.7.2	YES NVRAM	up	up
FastEthernet0/0.8	10.76.8.2	YES NVRAM	up	up
FastEthernet0/0.12	10.76.12.2	YES NVRAM	up	up
FastEthernet0/0.120	10.76.12.130	YES NVRAM	up	up
FastEthernet0/0.1000	10.76.0.2	YES NVRAM	up	up
FastEthernet0/1	unassigned	YES NVRAM	administratively down	down
GigabitEthernet0/1	unassigned	YES NVRAM	administratively down	down
FastEthernet0/2	unassigned	YES NVRAM	administratively down	down
GigabitEthernet0/2	unassigned	YES NVRAM	administratively down	down
GigabitEthernet0/3	unassigned	YES NVRAM	administratively down	down
Serial1/0	unassigned	YES NVRAM	administratively down	down
Serial1/0.402	unassigned	YES unset	administratively down	down
Serial1/0.404	10.1.254.237	YES NVRAM	administratively down	down
Serial1/1	unassigned	YES NVRAM	administratively down	down
Serial1/2	unassigned	YES NVRAM	administratively down	down
Serial1/3	unassigned	YES NVRAM	administratively down	down
Serial1/4	unassigned	YES NVRAM	administratively down	down
Serial1/5	unassigned	YES NVRAM	administratively down	down
Serial1/6	unassigned	YES NVRAM	administratively down	down
Serial1/7	unassigned	YES NVRAM	administratively down	down
Serial2/0:0	unassigned	YES NVRAM	up	up
Serial2/1:0	unassigned	YES NVRAM	up	up
Serial2/2:0	unassigned	YES NVRAM	up	up
Serial3/0	unassigned	YES NVRAM	up	up

Troubleshooting Tools

```

Serial3/0.100      172.16.128.26 YES NVRAM up up
Serial3/1          unassigned YES NVRAM down down
Serial4/0:0        unassigned YES NVRAM down down
Serial4/1:0        unassigned YES NVRAM down down
Serial4/2:0        unassigned YES NVRAM down down
Serial4/3:0        unassigned YES NVRAM down down
Serial4/4:0        unassigned YES NVRAM up up
Serial4/5:0        unassigned YES NVRAM up up
Serial4/6:0        unassigned YES NVRAM up up
Serial4/7:0        unassigned YES NVRAM up up
Serial6/0:0        unassigned YES NVRAM down down
Serial6/1:0        unassigned YES NVRAM down down
Serial6/2:0        unassigned YES NVRAM down down
Serial6/3:0        unassigned YES NVRAM down down
Serial6/4:0        unassigned YES NVRAM up up
Serial6/5:0        unassigned YES NVRAM up up
Serial6/6:0        unassigned YES NVRAM up up
Serial6/7:0        unassigned YES NVRAM up up
SSLVPN-VIF0       unassigned NO unset up up
Multilink20       10.1.254.249 YES NVRAM down down
Multilink31       10.1.254.245 YES NVRAM up up
Multilink42       10.1.254.233 YES NVRAM up up
Loopback0         10.1.1.1 YES NVRAM up up
Loopback1         10.254.253.94 YES NVRAM up up

```

To condense the output to the active parts, the following example pipes the output to exclude any lines with the words “unassigned” or “administratively.” Notice how much this simplifies the display:

```

Foard-rtr01# show ip interface brief | exclude unassigned|administratively
Interface                IP-Address      OK? Method Status Protocol

```

Troubleshooting Tools

NOTE:

The **alias** command can make this easier. In configuration mode, create a shortened version of a command as shown here.

```
Router(config)#alias
exec ii show ip
interface brief |
exclude
unassigned|adminis-
tratively
```

Now **ii** is the shortened version of the long and cumbersome command.

FastEthernet0/0	10.87.1.1	YES	NVRAM	up	up
FastEthernet0/0.2	10.76.2.2	YES	NVRAM	up	up
FastEthernet0/0.3	10.76.3.2	YES	NVRAM	up	up
FastEthernet0/0.4	10.76.4.2	YES	NVRAM	up	up
FastEthernet0/0.5	10.76.5.2	YES	NVRAM	up	up
FastEthernet0/0.6	10.76.6.2	YES	NVRAM	up	up
FastEthernet0/0.7	10.76.7.2	YES	NVRAM	up	up
FastEthernet0/0.8	10.76.8.2	YES	NVRAM	up	up
FastEthernet0/0.12	10.76.12.2	YES	NVRAM	up	up
FastEthernet0/0.120	10.76.12.130	YES	NVRAM	up	up
FastEthernet0/0.1000	10.76.0.2	YES	NVRAM	up	up
Serial3/0.100	172.176.128.26	YES	NVRAM	up	up
Multilink20	10.1.254.249	YES	NVRAM	down	down
Multilink31	10.1.254.245	YES	NVRAM	up	up
Multilink42	10.1.254.233	YES	NVRAM	up	up
Loopback0	10.1.1.1	YES	NVRAM	up	up
Loopback1	10.254.253.94	YES	NVRAM	up	up

A second example shows the OR capability by piping the output of **show process cpu** to include lines that start with CPU or include the words **IP Input**:

```
Foard-rtr01#show process cpu | inc ^CPU|IP Input
CPU utilization for five seconds: 14%/13%; one minute: 14%; five minutes: 14%
 87      2755772  47054573      58  0.07%  0.07%  0.07%  0 IP Input
```

Output Redirection

In addition to filtering output, IOS also enables **show** command output to be redirected. Redirecting output enables an administrator to collect information for archiving or to share with other troubleshooters and save it as a text file.

Troubleshooting Tools

Output can be piped to a file using either `redirect` or `tee`. `Redirect` just creates the file, whereas `tee` also displays the content in session. Any filesystem supported by that router is supported, so output can be pointed at `flash`, `tftp`, `ftp`, `http`, and other destinations.

The syntax to use this function is

```
Show command | redirect file
Show command | tee file
```

The next examples show the running configuration being piped to TFTP. In the first example, the output is redirected. The second example tees the output so that it builds the TFTP file and displays on screen.

```
Foard-rtr01#show running-configuration | redirect tftp://tftp/Foard-rtr01-shrun.txt
Translating "tftp"...domain server (10.186.2.30) [OK]
```

```
Foard-rtr01#show running-configuration | tee tftp://tftp/Foard-rtr01-shrun.txt
!
Building configuration...
```

```
Current configuration : 22291 bytes
```

```
...
```

IOS Troubleshooting Tools

Ping and traceroute are the most obvious tools available in IOS to test the network.

Ping tests connectivity and is so commonly used that even end users are passingly familiar with it. A ping response shows that a working path between two end points exists. End systems sometimes have firewalls that prevent response, but generally ping is a reasonable first test of network connectivity:

```
Foard-rtr01#ping 10.186.1.1
```

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

Troubleshooting Tools

```

Sending 5, 100-byte ICMP Echos to 10.186.1.1, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 8/9/12 ms

```

Exclamation marks show a response, but there is a lot of information besides the most obvious part. First, pay attention to the pattern of the response. Alternating success and failure (!.!.!) is a classic sign of a load balancing problem, where one path succeeds and the other fails. Second, pay attention to the response time. Many applications depend on quick response. Voice, for instance, assumes a round-trip time of less than 150 ms. The response time can also clue the troubleshooter to utilization issues. If the response time is much larger than usual that might indicate a heavy traffic load and queuing. If you notice that the minimum and maximum times vary widely, this could also be a sign of queuing because of a heavy load.

Ping can do a lot more than that simple test, however. Privileged mode supports an extended ping that enables every aspect of ping to be controlled. This opens up many more tests that can be accomplished with the humble command.

The following example below an extended ping. Notice that the command **ping**—with no destination specified—is entered in privileged mode. The example sends five pings of 100 bytes, then five of 200 bytes, continuing to 1500 byte pings. The DF bit (do not fragment) is set. A similar ping might be used if you suspect that an intermediate link didn't support the same size MTU as the source and destination. A more detailed explanation of the command is found after the example:

```

Foard-rtr01#ping
Protocol [ip]:
Target IP address: 10.186.1.1
Repeat count [5]:
Datagram size [100]:
Timeout in seconds [2]:
Extended commands [n]: y
Source address or interface: loopback0
Type of service [0]:

```


Troubleshooting Tools

```

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 10.186.1.1, timeout is 2 seconds:
Packet sent with a source address of 10.1.1.1
Packet sent with the DF bit set
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!
Success rate is 100 percent (75/75), round-trip min/avg/max = 8/10/12 ms

```

Remember that defaults are shown in square brackets. Selecting all the defaults is similar to a normal ping.

Sometimes testing involves repeatedly pinging (for instance, when you believe that an interface is flapping up and down). An extended ping with a repeat count of 99999 can be used to interactively test the network over a period of time.

Pings can be set to different packet sizes through the Datagram Size variable. The router can automate testing a range of sizes. To do so, use the extended commands and choose to sweep a range of sizes.

If a router is asked to forward a packet that is larger than the MTU of the transmitting link, the router normally breaks the packet into smaller pieces. Setting the DF bit instructs receiving routers to discard the traffic rather than fragment it.

Using different size packets and setting the DF bit allows testing MTU. When the MTU limit is reached, all subsequent pings will be dropped.

Troubleshooting Tools

Another nice testing technique is to change the source interface. Pings are normally sourced from the transmitting interface. Using an internal interface as the source shows that the receiving device and the intermediate routers understand how to route back to that prefix.

A final idea is to try different Type of Service settings. Many networks now carry voice, video, and prioritized data. Voice is commonly set to ToS 5, so pinging using ToS 5 enables a peek into how the QoS settings are functioning.

Like ping, there is an extended version of traceroute. It has a few of the same capabilities, with one other significant testing ability. Traceroute in IOS uses UDP, and extended traceroute enables setting the UDP port. This can be used to test application performance for applications that use UDP, such as voice. This is important when trying to diagnose the affects of firewalls and access-lists.

An example extended traceroute is shown next. The only choice specified in the example is to use UDP port 16000:

```
Newton-rtr01#traceroute
Protocol [ip]:
Target IP address: 10.200.1.1
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]: 16000
Loose, Strict, Record, Timestamp, Verbose[none]:
Type escape sequence to abort.
Tracing the route to 10.200.1.1
```

Troubleshooting Tools

In the same way the UDP port connectivity can be probed with traceroute, telnet can be used to test TCP ports. Telnet does not offer many options, but by changing the target port, different network services can be tested. The following examples show that email and the web server respond on the appropriate ports:

```
Foard-rtr01#telnet www.example.com 25
Translating "www.example.com"...domain server (10.1.2.2) [OK]
Trying www.example.com (172.16.0.25, 25)... Open
220 www.example.com ESMTP Postfix
```

```
Foard-rtr01#telnet www.example.com 80
<ctrl-c>
HTTP/1.1 400 Bad Request
Content-Type: text/html
Date: Fri, 4 Sep 2009 17:14:29 GMT
Connection: close
Content-Length: 35
```

```
<h1>Bad Request (Invalid Verb)</h1>
```

Hardware Diagnostics

The commands examined so far have dealt with network issues, but sometimes the problem is within the IOS device. Several commands describe the functional state of an IOS device.

If network hardware is suspected, a good place to start troubleshooting is to understand the external environment. The **show environment all** command displays information about the temperature within the device and the state of the power supplies. Especially when troubleshooting remotely it is easy to forget power and air conditioning, but problems in either area can lead to device malfunction:

Troubleshooting Tools

```
Foard-rtr01#sh environment all
```

```
Power Supplies:
```

```
Power Supply 1 is AC Power Supply. Unit is on.
```

```
Power Supply 2 is AC Power Supply. Unit is on.
```

```
Temperature readings:
```

```
NPE Inlet      measured at 25C/77F
```

```
NPE Outlet     measured at 27C/80F
```

```
I/O Cont Inlet measured at 25C/77F
```

```
I/O Cont Outlet measured at 28C/82F
```

```
CPU Die        measured at 43C/109F
```

```
Voltage readings:
```

```
+3.30 V        measured at +3.30 V
```

```
+1.50 V        measured at +1.49 V
```

```
+2.50 V        measured at +2.50 V
```

```
+1.80 V        measured at +1.79 V
```

```
+1.20 V        measured at +1.20 V
```

```
VDD_CPU        measured at +1.28 V
```

```
VDD_MEM        measured at +2.50 V
```

```
VTT            measured at +1.25 V
```

```
+3.45 V        measured at +3.43 V
```

```
-11.95         measured at -12.17 V
```

```
+5.15 V        measured at +4.96 V
```

```
+12.15 V       measured at +12.18 V
```

```
Envm stats saved 0 time(s) since reload
```

Troubleshooting Tools

A complete and accurate inventory is another part of troubleshooting. Of course, this information is much more useful if obtained before a problem occurs and connectivity drops! By comparing the inventory to previous inventories, it is possible to recognize differences (caused, presumably, by hardware failure). If the organization has a Cisco SmartNet maintenance contract, the serial number and part-number information is necessary to obtain spares:

```
foard-rtr01#show inventory
```

```
NAME: "Chassis", DESCR: "Cisco 7206VXR, 6-slot chassis"
```

```
PID: CISC07206VXR      , VID:      , SN: 24323096
```

```
NAME: "NPE-G2 0", DESCR: "Cisco 7200 Series Network Processing Engine NPE-G2"
```

```
PID: NPE-G2           , VID: V03 , SN: JAS1456B4EC
```

```
NAME: "disk2", DESCR: "256MB Compact Flash Disk for NPE-G2"
```

```
PID: MEM-NPE-G2-FLD256 , VID:      , SN:
```

```
NAME: "module 0", DESCR: "I/O Dual FastEthernet Controller"
```

```
PID: C7200-I/O-2FE/E   , VID:      , SN: 21753008
```

```
NAME: "disk0", DESCR: "Cisco 7200 I/O PCMCIA Flash Disk, 48M"
```

```
PID: MEM-I/O-FLD48M    , VID:      , SN:
```

```
NAME: "disk1", DESCR: "Cisco 7200 I/O PCMCIA Flash Disk, 48M"
```

```
PID: MEM-I/O-FLD48M    , VID:      , SN:
```

```
NAME: "module 1", DESCR: "Serial"
```

```
PID: PA-8T-V35=        , VID:      , SN: 49010448
```

```
NAME: "module 2", DESCR: "4 port, software configurable Multichannel T1/E1 with TDM Port Adapter"
```

```
PID: PA-MCX-4TE1       , VID:      , SN: JAS1680Y0EM
```

Troubleshooting Tools

```
NAME: "module 3", DESCR: "Enhanced 2 port T3/E3 clear channel PA"
PID: PA-2T3/E3-EC      , VID: V01 , SN: JAS249200K5
```

```
NAME: "module 4", DESCR: "8 port, software configurable Multichannel T1/E1 without TDM Port Adapter"
PID: PA-MC-8TE1+      , VID:      , SN: JAS1689A2MM
```

```
NAME: "module 6", DESCR: "8 port, software configurable Multichannel T1/E1 without TDM Port Adapter"
PID: PA-MC-8TE1+      , VID:      , SN: JAS1689A2BV
```

```
NAME: "Power Supply 1", DESCR: "Cisco 7200 AC Power Supply"
PID: PWR-7200-AC      , VID:      , SN:
```

```
NAME: "Power Supply 2", DESCR: "Cisco 7200 AC Power Supply"
PID: PWR-7200-AC      , VID:      , SN:
```

A lack of memory can also cause a network issue. The **show memory** command displays the state of memory on a device; focus on the Free column to determine if enough is available. Another sign of memory issues is %SYS-2-MALLOCFAIL messages:

```
foard-rtr01#show memory
```

	Head	Total	Used	Free	Lowest	Largest
Processor	6319860	818832732	74864300	743968432	742841100	727580236
I/O	38000000	67108864	11964260	55144604	55137712	54643068
Transient	37000000	16777216	58244	16718972	16226680	16718696

```
...
```

Hardware issues can also manifest themselves on the interfaces. **Show controller** can show some information about the interface—serial interfaces in particular report things such as cable information here. **Show interface** (shown next) displays a good deal of information about the state of the interface. In particular, pay attention to four measurements:

Troubleshooting Tools

- **Input queue drops:** Signify that the router had more traffic than it could process. Some amount of drops is excusable, but drops could be related to CPU oversaturation. Double-check the processor with the **show processes cpu** command.
- **Output queue drops:** Usually mean that the line is congested.
- **Input errors:** These errors show duplex errors, interface problems, and CRC errors.
- **Output errors:** Usually related to duplex issues.

```

Foard-rtr01#show interface

```

```

FastEthernet0/0 is up, line protocol is up
  Hardware is i82543 (Livengood), address is 000a.f3f7.9808 (bia 000a.f3f7.9808)
  Description: enter port #
  Internet address is 10.100.1.1/16
  MTU 1500 bytes, BW 100000 Kbit/sec, DLY 100 usec,
    reliability 255/255, txload 32/255, rxload 14/255
  Encapsulation 802.1Q Virtual LAN, Vlan ID 1., loopback not set
  Keepalive set (10 sec)
  Full-duplex, 100Mb/s, 100BaseTX/FX
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input 00:00:00, 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 5517000 bits/sec, 2571 packets/sec
  5 minute output rate 12927000 bits/sec, 2550 packets/sec
    1326060749 packets input, 711066620 bytes
      Received 45468700 broadcasts, 0 runts, 0 giants, 0 throttles

```

Troubleshooting Tools

```
148 input errors, 0 CRC, 0 frame, 0 overrun, 148 ignored
0 watchdog
0 input packets with dribble condition detected
1191821108 packets output, 2981100223 bytes, 0 underruns
2 output errors, 0 collisions, 4 interface resets
5634739 unknown protocol drops
0 babbles, 0 late collision, 0 deferred
2 lost carrier, 0 no carrier
0 output buffer failures, 0 output buffers swapped out
```

Working with External Tools

The IOS troubleshooting capabilities are supplemented by external network management tools. Cisco IOS devices support these tools and in many cases supply detailed information to the management system. This section describes the methods used to coordinate with these tools.

Packet Sniffing

Packet capture from a laptop or specialized device enables low-level vision into the exact traffic flowing over a link. Capturing traffic can show errors and underlying protocol traffic. The issue with packet capture is that switches do not forward all traffic out all ports, so it is difficult to find a port from which to see all traffic.

SPAN (Switched Port Analyzer) is a tool within IOS switches to direct copies of packets to a capture port. SPAN is configured by identifying a source port or VLAN from which traffic should be copied. SPAN is then pointed to an output port, to which a packet capture tool is attached. SPAN can capture traffic on a switch and output to a trunked VLAN. A second switch can then capture the VLAN and output it to a port. This configuration is called remote SPAN (RSPAN).

Troubleshooting Tools

The generic configuration of SPAN is

```
Monitor session [session number] [source|destination] [interface|vlan]
```

The following example shows the configuration used when suspicious device is on port F0/1 and a packet capture tool is plugged into port F0/24:

```
Monitor session 1 source interface f0/1
Monitor session 1 destination interface f0/24
```

Router IP Traffic Export (RITE) is similar to SPAN but used by routers to capture traffic to a monitoring port. The following example demonstrates capturing ten percent of the interesting traffic on f0/1 and exporting it to a device with a given MAC:

```
(config)# ip traffic-export profile rite
(config-rite)# interface FastEthernet 0/1
(config-rite)# bidirectional
(config-rite)# mac-address 00a.8aab.90a0
(config-rite)# incoming access-list my_acl
(config-rite)# outgoing sample one-in-every 10
(config)# interface FastEthernet0/0
(config-if)# ip traffic-export apply rite
```

RITE can also be used to export the traffic to a file on the router. From there it can be copied off for inspection on a PC:

```
traffic-export interface fastethernet0/0 copy tftp:
```

Netflow

Netflow collects summaries of traffic information and transmits the summary to a Netflow collector. Netflow is enabled on each monitored interface. Netflow supports a version 5 and version 9; this should be set to match the requirements of your network management system. Finally, Netflow exports information to a target IP address. The commands to accomplish these actions are shown here:

Troubleshooting Tools

```
(config-if)#ip flow ingress
(config)#ip flow-export version [5|9]
(config)#ip flow-export destination [ip-address]
```

In addition to using a monitoring system to track Netflow, an administrator can also peek into the current flows using **show ip cache flow**.

SNMP and EEM

SNMP is another monitoring protocol. Whereas Netflow tracks traffic, SNMP can monitor any type of event or statistic from the device. SNMP is supported by most network monitoring systems. The router also has a tool to react to events through embedded event manager (EEM).

SNMP is set up by identifying a server and listing the events to be monitored. If **snmp-server enable traps** is used without specifying specific events, all traps are monitored:

```
(config)#snmp-server host [ip-address]
(config)#snmp-server enable traps
```

EEM enables custom reactions to events and acts as a supplement to SNMP. Events can be triggered by any SNMP event and actions can include (among others) SNMP, Syslog, IOS commands, and email messages.

A simple example EEM applet is shown next. This applet logs a Syslog message and outputs a message to the console in reaction to an administrator entering configuration mode:

```
Event manager applet CONFIG-STARTED
Event cli pattern "configure terminal" sync on skip no occurs 1
Action 1.0 syslog priority critical msg "Configuration mode was entered"
Action 2.0 syslog priority informational msg "Change control policies apply. Authorized access only."
```

NOTE:

EEM applets are starting to appear on the Internet, both at Cisco.com and at other sites.

Chapter 4

Troubleshooting Switches

Ethernet is ubiquitous in campus networks and Data Centers. Movement to consolidate networks has collapsed storage and virtualization, and telephony has put more traffic on Ethernet. Maintaining this critical infrastructure involves understanding the component pieces: Spanning Tree, VLANs, InterVLAN routing, and gateway redundancy.

Poor forwarding performance on switches is usually associated with cabling and port problems, duplex mismatch, or TCAM issues.

Problems at the physical layer can be seen from **show interface**, **show interface counters** and **show interface counters errors**. Look for the following errors:

Align-Err, runts: Alignment errors are usually associated with cabling, NICs, or duplex mismatch.

- **FCS-Err:** Frame Check Sequence errors are usually associated with a cabling issue.
- **Xmit-Err:** The transmission buffers are full. Commonly associated with switching a faster link to a slower link.
- **Undersize, Giants:** Suspect the transmitting NIC.
- **Single-Col, Multi-Col, Late-Col, Excess-Col:** Collisions are a sign of duplex mismatch.

An example of these commands is shown here.

```
Newton-Sw01#show interface fastethernet1/1
FastEthernet1/2 is up, line protocol is up (connected)
Hardware is C6k 100Mb 802.3, address is 001c.58c8.ac92 (bia 001c.58c8.ac92)
MTU 1500 bytes, BW 100000 Kbit, DLY 100 usec,
    reliability 255/255, txload 1/255, rxload 1/255
```

Troubleshooting Switches

```

Encapsulation ARPA, loopback not set
Keepalive set (10 sec)
Full-duplex, 100Mb/s
input flow-control is off, output flow-control is unsupported
ARP type: ARPA, ARP Timeout 04:00:00
Last input never, output 00:00:43, output hang never
Last clearing of "show interface" counters 6w5d
Input queue: 0/2000/0/0 (size/max/drops/flushes); Total output drops: 0
Queueing strategy: fifo
Output queue: 0/40 (size/max)
1 minute input rate 0 bits/sec, 0 packets/sec
1 minute output rate 7000 bits/sec, 9 packets/sec
 4182737 packets input, 719363170 bytes, 0 no buffer
 Received 5970 broadcasts (174 multicasts)
  0 runts, 0 giants, 0 throttles
  0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
  0 watchdog, 0 multicast, 0 pause input
  0 input packets with dribble condition detected
45957071 packets output, 19815895675 bytes, 0 underruns
  0 output errors, 0 collisions, 0 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

```

Newton-Sw01#**sh interface counters**

Port	InOctets	InUcastPkts	InMcastPkts	InBcastPkts
Fa1/1	6658590	73024	27	95
Fa1/2	719363238	4176768	174	5796

Troubleshooting Switches

```

...
Newton-Sw01#sh interface counters errors

Port          Align-Err    FCS-Err    Xmit-Err    Rcv-Err    UnderSize  OutDiscards
Fa1/1          0            309        0            309        0           0
Fa1/2          0            0          0            0          0           0
...

```

Duplex mismatch is a common cause of forwarding problems. Half-duplex is unusual in modern networks, so duplex mismatch usually occur when one port is set to auto and the other to full. Setting everything to auto is Cisco's recommendation.

Spanning Tree

Redundancy is a common technique to increase availability in computer networks. Ethernet redundancy would look like multiple core switches and multiple paths between workgroup switches and the core. Of course, multiple paths mean loops, and Ethernet lacks a mechanism for dealing with loops.

Spanning Tree is a protocol that detects potential loops and breaks them:

1. Each switch advertises Bridge Protocol Data Units (BPDU) that periodically announces name (bridge ID), current root, and cost to the root. Each switch starts believing it is the root.
2. If a switch receives a BPDU with a different root, it compares roots. If the received BPDU has a lower root, the switch changes root and recalculates cost to the root. The port that received the superior BPDU is the **root port**—the port that leads to the root. Other ports are **designated ports**—ports leading away from the root.

NOTE:

When speed and duplex are auto, Cisco switches also support auto-MDIX. (The switch will adjust the port to be straight through or crossover as needed.)

```

Interface f0/0
  Mdx auto

```

Troubleshooting Switches

Each link has a cost based on its speed, as shown in the following table.

Link Speed	Cost
Ethernet	100
Fast Ethernet	19
Gigabit Ethernet	4
Ten Gigabit Ethernet	2

3. If a switch receives two BPDUs with the same root but different costs, it uses the lower cost port. The port with the higher cost is blocked (it filters all traffic except BPDUs) to prevent a loop. Blocked ports are also called **non-designated**.

At the end of the process there will be one root bridge. Each nonroot switch will have one root port.

Spanning tree status can be seen using the **show spanning-tree [vlan vland-id]** command, as shown here:

```
Newton-Sw01#show spanning-tree vlan 1
```

```
VLAN0001
```

```
Spanning tree enabled protocol rstp
```

```
Root ID    Priority    8192
           Address    001d.4664.7d01
           Cost      4
           Port      641 (GigabitEthernet6/1)
Hello Time  2 sec Max Age 20 sec Forward Delay 15 sec
```

Troubleshooting Switches

```

Bridge ID Priority    32768
Address      001d.46c8.ac01
Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
Aging Time  300

```

```

Interface      Role Sts Cost      Prio.Nbr Type
-----
Fa1/2          Desg FWD 19        128.2   Edge P2p
Fa1/3          Desg FWD 19        128.3   Edge P2p
Fa1/4          Desg FWD 19        128.4   Edge P2p
Fa1/5          Desg FWD 19        128.5   Edge P2p
Fa1/7          Desg FWD 19        128.7   Edge P2p
Fa1/9          Desg FWD 19        128.9   Edge P2p
Fa1/10         Desg FWD 19        128.10  Edge P2p
Fa1/11         Desg FWD 19        128.11  Edge P2p
Fa1/12         Desg FWD 19        128.12  Edge P2p
...

```

The details of received BPDUs can be seen using **show spanning-tree interface [interface] detail**. This command shows root status, cost, and timers:

```
Newton-Sw01#show spanning-tree vlan 1 detail
```

```

VLAN0001 is executing the rstp compatible Spanning Tree protocol
Bridge Identifier has priority 32768, address 001d.4664.cc01
Configured hello time 2, max age 20, forward delay 15, tranmsit hold-count 6
Current root has priority 8192, address 001d.4632.6c01
Root port is 641 (GigabitEthernet6/1), cost of root path is 4
Topology change flag not set, detected flag not set
Number of topology changes 119 last change occurred 25w6d ago
from GigabitEthernet6/1

```

Troubleshooting Switches

```

Times: hold 1, topology change 35, notification 2
      hello 2, max age 20, forward delay 15
Timers: hello 0, topology change 0, notification 0, aging 300

```

```

Port 2 (FastEthernet1/2) of VLAN0001 is designated forwarding
  Port path cost 19, Port priority 128, Port Identifier 128.2.
  Designated root has priority 8192, address 001d.4664.ec01
  Designated bridge has priority 32768, address 001d.4664.cc01
  Designated port id is 128.2, designated path cost 4
  Timers: message age 0, forward delay 0, hold 0
  Number of transitions to forwarding state: 1
  The port is in the portfast mode
  Link type is point-to-point by default
  Bpdu guard is enabled
  Root guard is enabled on the port
  BPDU: sent 9120, received 0

```

...

Before spanning-tree, loops meant that traffic would cycle continuously. Over a short time traffic would accrete in the loop until it consumed all capacity. This is called a broadcast storm. *Broadcast storms* are still a real danger, but spanning tree has mitigated this almost entirely. The danger today is that—through protocol failure or administrative misprogramming—when a broadcast storm forms, few administrators have seen it before and know how to deal with it.

A broadcast storm can be diagnosed when the switches become saturated with traffic. All the traffic lights will be solid, the switch will be slow to respond, and users will complain about network speed.

The only fix for a broadcast storm is to break the loop. If the switches are accessible, it might be possible to fix spanning-tree. Otherwise, the administrator must manually remove redundant links.

Troubleshooting Switches

As previously stated, the purpose of spanning-tree is to select one root path and filter all others. When there are multiple links between two switches it seems intuitive that, rather than turn one off, the switches should use all the links together. This is possible using Etherchannel.

Etherchannel logically combines several physical links between switches and spanning tree treats the bundle as a single port. Up to eight physical lines may be combined in this way.

Etherchannel failures cluster into three groups:

- All ports must be identical (speed, duplex, access or trunk, VLAN). If Etherchannel will not form, look for inconsistencies between ports.
- Both switches must either be configured or a link aggregation protocol (LACP or PAgP) must be used. If only one side is configured for Etherchannel, look for Etherchannel ports that are error-disabled.
- The channel might form, but traffic might still be traveling predominately over a single link. This is because traffic is statistically multiplexed using a three-bit hash. This means that the traffic is split over eight paths, and an etherchannel of three links will split the load in a 2:1:1 ratio. Fix this by using 2, 4, or 8 links. Second, the hash uses a user-selectable Ethernet or IP field. If all traffic comes from a single source and the switch is hashing on source MAC, it will not multiplex. Fix this by selecting a different hashing method.

VLANs

Virtual LANs are logical broadcast domains, administratively assembled from component ports on the switches in the network. Switches are interconnected by Ethernet lines that use 802.1q, a shim header inserted in the Ethernet frame. 802.1q adds a two-byte shim, 12 bits of which are used to identify the VLAN and three bits of which are used to specify Layer-2 class of service. (This is called the 802.1p subfield.)

Troubleshooting Switches

When troubleshooting VLaN switching issues, concentrate on three types of failure:

- **Wiring issues:** Cabling issues, power outage, or bad switch ports
- **Switch issues:** Software bugs, hardware bugs, loops, and ARP issues
- **Logic issues:** Misconfigured VLANs, VTP, trunks, and native VLAN mismatch

Troubleshooting switches often involves using these tables to understand the path traffic takes through the switch. Two commands can help identify the path taken:

- **Show platform forward:** Displays forwarding info from TCAM
- **Traceroute mac:** Shows intermediate MACs from source to destination

Switches keep several mapping tables. Each of these tables is shown in the following table, as well as the IOS command to examine the table.

Table	IOS Command
MAC Address Table: Maps MAC addresses to ports	Show mac-address
VLAN assignments: Maps VLANs to ports Show interface switchport	Show vlan
VLAN Database: Maps names to VLANs	Show vlan
Trunk assignments Show interface trunk Show etherchannel	Show interface switchport

Switched Virtual Interfaces and InterVLAN routing

Routing between VLANs can be accomplished on a Layer 3 switch or on a router. Troubleshooting the control plane (the Layer 3 structures) is identical between the two. This means that OSPF runs identically on the two platforms.

The data plane (the structures and hardware that handle frame forwarding) is different between routers and Layer 3 switches. In both cases, **show ip cef** shows the cef forwarding table, and **show adjacency** shows the Layer 2 headers used in forwarding.

Catalyst 3560, 3750, and 4500 switches can also use **show platform** to see detailed forwarding information.

Catalyst 6500 switches display forwarding details using **show mls cef** commands.

Another difference between routers and Layer 3 switches, in the context of troubleshooting intervlan routing, is the concept of an SVI (Switched Virtual Interface).

Routers forward traffic between ports using Layer 3 information.

Layer 3 switches can have multiple ports in the same vlan and pass traffic between them using MAC information. Layer 3 switches also support SVIs (these look like interface vlan 1) that act as virtual layer-3 ports for a VLAN. Finally, a switch can treat a port as a separate routed port.

From a troubleshooting perspective, routed ports do not run switching protocols like Spanning Tree or Etherchannel. SVIs, on the other hand, are extremely stable. An SVI changes only state to down when all the VLAN ports are down.

First-Hop Redundancy

Hosts are configured with a default gateway—a router address that will pass traffic off the local subnet. The problem is that router failures strand the hosts. The solution is first-hop redundancy protocols, which enable two routers to cooperatively support a single IP, which can then be given to hosts as a default gateway.

Troubleshooting Switches

There are three first-hop redundancy protocols:

- HSRP is an older Cisco proprietary protocol. One router is the active and one is the standby. The routers pass keepalives that enable the standby to recognize failure of the primary router.
- VRRP is an open standard but is otherwise similar to HSRP. Because HSRP works, many organizations have continued to use HSRP.
- GLBP is an open standard, but it enables simultaneous load balancing over as many as four gateways.

Because HSRP is the most common, this section focuses on HSRP. The general configuration and troubleshooting strategy applies well to VRRP and GLBP, however.

HSRP is configured under the interface using **standby** commands. Routers in the same HSRP group share a Mac and IP, so standby is used to identify the group and virtual IP.

By default, each HSRP speaker has a priority of 100. The speaker with the highest priority is the active router. If a new router starts however, HSRP does not change the active router until the failure of the active router. To change this so that the higher priority is instantly recognized, use the **preempt** command. An HSRP snippet is shown here to illustrate the configuration:

```
Interface f0/0
Ip address 10.1.1.2 255.255.255.0
Standby 2 ip 10.1.1.1
Standby 2 priority 120
Standby 2 preempt
```

Troubleshooting Switches

Verify the HSRP state of a router using **show standby**, which summarizes this information to a table (an example is shown next). To see detailed information on HSRP, such as timers and virtual MAC, use **show standby interface**:

```
Maiden-rtr01#show standby
GigabitEthernet0/1 - Group 135
  State is Active
    23 state changes, last state change 25w6d
  Virtual IP address is 135.159.64.1
  Active virtual MAC address is 0000.0c07.ac87
    Local virtual MAC address is 0000.0c07.ac87 (v1 default)
  Hello time 5 sec, hold time 20 sec
    Next hello sent in 0.284 secs
  Preemption enabled
  Active router is local
  Standby router is unknown
  Priority 150 (configured 150)
  Group name is "hsrp-Gi0/1-135" (default)
Richardson-rtr01#show standby interface gi0/1
Global          Config: 0000
Gi0/1 If hw     BCM1125 Internal MAC (27), State 0x210040
Gi0/1 If hw     Config: 0000
Gi0/1 If hw     Flags: 0000
Gi0/1 If sw     Config: 0000
Gi0/1 If sw     Flags: 0000
Gi0/1 Grp 135   Config: 0072, IP_PRI, PRIORITY, PREEMPT, TIMERS
Gi0/1 Grp 135   Flags: 0000
```

Troubleshooting Switches

```
HSRP virtual IP Hash Table (global)
103 172.25.96.1    Gi0/1    Grp 135
```

```
HSRP MAC Address Table
43 Gi0/1 0000.0c07.ac87
    Gi0/1 Grp 135
```

show standby brief is mirrored with **show vrrp brief** and **show glbp brief**. Similarly, **show standby interface** and **debug standby** have equivalents for the other first-hop redundancy protocols.

Chapter 5

Troubleshooting Routing

This section reviews troubleshooting for common routing protocols. A more theoretical explanation of the working of the protocols is available in the *BSCI Quick Reference Guide*.

Network Layer Connectivity

Routers use three tables to make routing decisions: the routing table, ARP table, and CEF mappings

The routing table is visible using **show ip route**. Each entry in the routing table has an output interface or next hop. Packets are routed per the routing table, matching the longest prefix match first and then by other metrics determined by that IGP's algorithm.

When a determination of the next hop has been made, the router needs to turn this information into a destination Layer 2 address. For this purpose, mapping tables are maintained that match Layer 2 and Layer 3 addresses. The ARP table (**show ip arp**) and the frame-relay map (**show frame-relay map**) are examples of this.

Cisco Express Forwarding (CEF) is the common switching method found on most Cisco gear. CEF combines information from the routing table and the various mapping tables to optimize routing and to optimize the construction of new Layer 2 headers. CEF entries may be viewed using **show ip cef** and associated commands.

Routing Protocols

Routing protocols are mechanisms that enable routers to share information about the structure of the network. Regardless of the protocol, troubleshooting routing protocol issues have some basic logic that is true for any routing protocol.

Troubleshooting routing issues always starts with looking at the routing table. Use **ping** to test connectivity, **show ip route** to inspect the routing table to see if the route is present, and **tracert** to inspect how traffic is forwarding. **show ip protocols** displays information about the current routing protocols, such as autonomous system and timer values.

Troubleshooting routing issues can be summarized by answering three basic questions:

1. Is the correct route advertised?
2. Is the correct route communicated?
3. Is there a more desirable path (lower AD or longer prefix length)?

EIGRP

After determining that there is a routing problem in EIGRP using the routing table or ping, follow the three basic steps to troubleshooting.

EIGRP stores information in three tables that can be interrogated.

Table	Command
Interface table: Lists EIGRP-enabled interfaces	Show ip eigrp interface
Neighbor table: Lists discovered neighbors	Show ip eigrp neighbors
Topology table: Complete list of received EIGRP routes	Show ip eigrp topology

Is the Correct Route Advertised?

Verify that the router attached to the destination subnet is advertising the route. There are several ways to see the advertised subnets; two good ways are either direct interrogation of the running configuration using **show running-config | section eigrp** or by reviewing the protocol settings using **show ip protocol** (shown here):

```
Hickory-rtr01#show ip protocol
```

```
Routing Protocol is "eigrp 10"
```

```
  Outgoing update filter list for all interfaces is not set
```

```
  Incoming update filter list for all interfaces is not set
```

```
  Default networks flagged in outgoing updates
```

```
  Default networks accepted from incoming updates
```

```
  EIGRP metric weight K1=1, K2=0, K3=1, K4=0, K5=0
```

```
  EIGRP maximum hopcount 100
```

```
  EIGRP maximum metric variance 1
```

```
  Redistributing: eigrp 100, bgp 65096
```

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

```
  Automatic network summarization is not in effect
```

```
  Maximum path: 4
```

```
  Routing for Networks:
```

```
    10.0.0.0
```

```
  Passive Interface(s):
```

```
    GigabitEthernet0/1
```

```
  Routing Information Sources:
```

```
    Gateway          Distance      Last Update
```

```
    10.1.4.254        90           00:39:11
```

```
    10.1.4.253        90           00:38:55
```

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

EIGRP also advertises only subnets of interfaces that match a network statement. **show ip protocol** provides the matching network statements.

Is the Correct Route Communicated?

EIGRP shares only routes with neighbors—devices with which it has exchanged hellos. Verify that connected devices are neighbors using **show ip eigrp neighbors**. **debug ip eigrp packets** should show hellos and updates if devices are connected, and **debug ip eigrp** should show details about the contained routing information communicated.

EIGRP neighborship requires bidirectional communication, authentication, that the AS be the same, and that timers are close to the same. EIGRP also sends only hellos over interfaces that match a network statement. If a router hasn't identified a link as an EIGRP link in this way, it will not send hellos and it will not form neighborship. EIGRP values, such as timers, and a list of EIGRP interfaces is available through **show ip eigrp interfaces**:

```
Hickory-rtr01#show ip eigrp neighbors
```

```
IP-EIGRP neighbors for process 100
```

H	Address	Interface	Hold Uptime (sec)	SRTT (ms)	RTO	Q	Seq Cnt Num
1	10.1.4.253	Gi0/0	14 2w0d	1	200	0	1797
0	10.1.4.254	Gi0/0	14 2w0d	1	200	0	729

```
Hickory-rtr01#show ip eigrp interface
```

```
IP-EIGRP interfaces for process 100
```

Interface	Peers	Xmit Queue Un/Reliable	Mean SRTT	Pacing Time Un/Reliable	Multicast Flow Timer	Pending Routes
Gi0/0	2	0/0	1	0/1	50	0
Lo0	0	0/0	0	0/1	0	0

Troubleshooting Routing

If the devices are neighbors, routes could be blocked using distribution lists or route-maps. Distribution lists would be listed in **show ip protocol**.

Is There a More Desirable Path?

Finally, if the route is not in the routing table, use **show ip eigrp topology** to see if the route is known to EIGRP. It could be that the route is known, but there is a more desirable path. **show ip route** shows only the selected EIGRP route. To see all known EIGRP routes, use **show ip eigrp topology**.

OSPF

Three OSPF tables can be reviewed in troubleshooting. A fourth—the Routing Information Base—is used to store SPF calculations but is largely unavailable to the administrator.

Table	Command
Interface table: Lists OSPF-enabled interfaces	Show ip ospf interface
Neighbor table: Lists discovered neighbors	Show ip ospf neighbors
Link State Database: LSAs received	Show ip ospf database

If a routing problem exists in OSPF, follow the same basic steps to troubleshooting.

Is the Correct Route Advertised?

Verify that the router attached to the destination subnet is advertising the route. Advertised subnets are visible using either **show running-config | section ospf** or by reviewing **show ip protocol**.

OSPF also limits advertisements to the subnets of interfaces that match network statements. **show ip protocol** provides the matching network statements. **show ip ospf statistics** can also help by showing how often SPF is running, potentially showing network instability.

Is the Correct Route Communicated?

OSPF shares routes with neighbors. Verify that connected devices are neighbors using **show ip ospf neighbors**. **show ip ospf database** displays the link state information. **debug ip ospf adj** should show issues preventing neighborship.

OSPF neighborship requires six parameters to agree:

- Bidirectional communication.
- Equal timer values.
- Matching AS number.
- Routers must agree on the type of their common area.
- Routers must agree on the prefix of their common subnet.
- Authentication, if used, must agree on type and password.

OSPF sends only Hellos over interfaces that match a network statement. If a link does not match a network entry, no Hellos will be transmitted and no neighbors will form over the link. OSPF protocol values can be seen using **show ip ospf interfaces**.

If the devices are neighbors, routes could be blocked at boundary routers using distribution lists or route-maps. Distribution lists would be listed in **show ip protocol**.

Is There a More Desirable Path?

It is possible that OSPF has chosen an unexpected path to a destination. It could also be that routes from other routing protocols are present with a lower administrative distance or that an intermediate system has a static route. Checking routing tables along the expected path is the best way to reveal this.

BGP

BGP maintains two tables outside of the routing table, one for neighbors and one for BGP routing information.

Table Command

Neighbor table: Lists neighbors **Show ip bgp neighbors**

BGP table: Contains all received BGP prefixes and associated attributes, as well as showing the BGP best path **Show ip bgp**

BGP troubleshooting can also follow the three basic steps.

Is the Correct Route Advertised?

Verify that the router attached to the destination subnet is advertising the route. This can be seen from the running configuration (**show running-config | section bgp**) or the BGP table (**show ip bgp**—self-originated routes have a next hop of 0.0.0.0).

BGP advertises only explicitly identified prefixes for which there is a matching route from another source (like a connected route).

Is the Correct Route Communicated?

BGP communicates prefixes with administratively defined neighbors. Verify that defined neighbors are reachable using ping and that they are neighbors by reviewing **show ip bgp neighbors**. A partial output from this is shown next—**show ip bgp neighbors** includes considerable detail. **debug ip bgp updates** should show hellos and advertisements, and **debug ip bgp** should show details about the contained routing information being communicated:

```
Hickory-rtr01#show ip bgp neighbor
```

```
BGP neighbor is 10.1.255.5, remote AS 4800, external link
```

```
  BGP version 4, remote router ID 59.43.0.71
```

```
  BGP state = Established, up for 2w0d
```

Troubleshooting Routing

```

Last read 00:00:15, last write 00:00:17, hold time is 90, keepalive interval is 30 seconds
Neighbor capabilities:
  Route refresh: advertised and received(old & new)
  Address family IPv4 Unicast: advertised and received
Message statistics:
  InQ depth is 0
  OutQ depth is 0

                Sent      Rcvd
Opens:           1         1
Notifications:   0         0
Updates:         2        1162
Keepalives:      40808     40817
Route Refresh:   0         0
Total:           40811     41980
Default minimum time between advertisement runs is 30 seconds

```

...

BGP neighborship requires bidirectional communication, authentication, and that the AS match the expected AS. BGP values, such as timers and AS, are available through **show ip bgp**.

If the devices are neighbors, routes could be blocked using distribution lists or route-maps. Distribution lists would be listed in **show ip protocol**.

Is There a More Desirable Path?

If the route is not in the routing table, use **show ip bgp** to see if the route is known and valid. Routes can be invalidated if the BGP next hop is unreachable; if so routing to this address must be recursively troubleshoot. The following partial example shows several routes that are **valid** and **best**, shown by the preceding ***>**.

Troubleshooting Routing

```

ahk-rtr01#sh ip bgp
BGP table version is 17312, local router ID is 10.254.254.12
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

```

Network	Next Hop	Metric	LocPrf	Weight	Path
*> 0.0.0.0	182.225.207.13	0	65000	65097	i
*> 10.43.0.0/24	182.225.207.13	0	65000	65086 65042	i
*> 10.43.0.0/22	182.225.207.13	0	65000	65086 65042	i
*> 10.45.128.0/24	182.225.207.13	0	65000	65100 65044	i
*> 10.49.0.0/22	182.225.207.13	0	65000	65086 65300	i
*> 10.61.0.0/16	182.225.207.13	0	65000	65060	i
*> 10.63.0.0/20	182.225.207.13	0	65000	65062	i
*> 10.65.0.0/19	182.225.207.13	0	65000	65064	i
*> 10.71.0.0/16	182.225.207.13	0	65000	65086 65302	i
*> 10.87.0.0/16	182.225.207.13	0	65000	65086	i
...					

Route Redistribution

Organization sometimes must support more than one routing protocol. For example, a business might use EIGRP within a campus and BGP over the MPLS WAN. Routing information is passed between the protocols using redistribution. Redistributed routes are treated as external in the receiving protocol.

Redistribution extracts routes from the routing table, so only routes that appear in the routing table will be exported. If routes are not present, confirm the routes are present in the routing table at the redistribution point. You need to identify and understand the interaction of all redistribution points. Creating a routing loop through multiple redistribution points is quite possible.

Because routing protocols use different metrics, redistributed routes lose routing information. Distance Vector routing protocols, including EIGRP, assume that the metric for imported routes should be infinity unless another value is specified. When redistributing into EIGRP, a default metric must be set or no routes will be imported! OSPF will import only classful routes unless **redistribute subnets** is used, so this is also a point to review in troubleshooting.

In addition to protocol specific commands, **debug ip routing** can show routes as they are added or withdrawn from the routing table.

If **ip route profile** is added to the config, the **show ip route profile** command shows routing table changes over consecutive 5-second intervals. This is particularly helpful to show that routes are flapping—being added and withdrawn continuously.

Router Performance

Routing protocol performance can be symptomatic of general router problems. Routing protocol problems can be seen if the router CPU is overburdened or memory is fully utilized.

Transient events, such as SNMP communication or a heavy traffic load, can temporarily spike the CPU. High CPU utilization is a concern when it becomes on-going. Signs of CPU oversubscription include dropped packets, increased latency, slow response to telnet and console, and when the router skips routing updates.

Show process cpu can identify processes that are consuming CPU cycles. The ARP Input process consumes more cycles if the router has to generate a large number of ARPs, for instance in response to malicious traffic. Net Background is used to manage buffer space. IP Background is used whenever an interface changes state, utilization here could indicate a flapping interface.

Show process cpu history displays the overall utilization as a bar graph. This is a nifty way to see if the current load is an aberration or the norm.

Troubleshooting Routing

A second general router issue is the router switching mode. There are three common modes:

- Process switching uses the CPU to process each packet. Process switching is CPU-intensive and reduces throughput and increases jitter. It is turned on by using **no ip route-cache**.
- Fast switching uses the CPU to process an initial packet but then caches the result. It is less CPU-intensive, but utilization still tracks the traffic load. It is turned on using **ip route-cache**, and the cache can be reviewed using **show ip cache**.
- Cisco Express Forwarding (CEF) is the default switching mode. CEF is resilient to traffic load. It is turned on using **ip cef**, and CEF entries can be seen by using **show ip cef** and **show adjacency**. CEF is required for some IOS features, such as NBAR, WRED, and AutoQoS.

The interface switching mode is shown from the **show ip interface** command.

A third general router issue is router memory utilization. Memory is over-used when there is no available system memory or when the memory is too fragmented to be useful.

One easy, but not pleasant, way to see a memory problem is to load a version of IOS that requires more RAM than is present on the router. Memory can also be depleted by a memory leak—a bug that assigns memory to processes but does not clean up when the process is complete. Memory leaks can be recognized over time using **show memory allocating-process totals** and **show memory dead** and by researching known bugs within CCO. If found, the only solution is to move to a known good version of IOS.

Memory leaks sometimes appear on interfaces as buffer leaks. Buffer leaks can be seen using **show interface**, where the “input queue” shows buffer utilization. **Show buffer** also shows a buffer leak, here by looking at the number of free buffers.

Finally, memory leaks are sometimes seen in BGP, which is a heavy consumer of memory in the best of times, so a memory leak here can quickly bloom into a larger issue. **show process memory | include bgp** shows the memory utilization of the four BGP processes. **show diag** can be used to evaluate memory used on the line cards.

Chapter 6

Troubleshooting Security Features

Network security has been seen as a separate function, but security has evolved to be a pervasive element. Routers are both potential targets for attacks and platforms that can offer security services.

Network devices have three types of functions and traffic, all of which are affected by security concerns:

- **Management plane:** The functions involved in management, such as device access, configuration, and telemetry.
- **Control plane:** The functions spoken between network devices, such as routing protocols.
- **Data plane:** Packet forwarding functionality.

Security for the management plane means controlling all the means of accessing the device and making configuration changes. Common security steps for various protocols include

- **Console:** Physically secure access to the device and set reasonable time-outs. Use password protected modems for out-of-band access, and control authentication centrally with RADIUS or TACACS+ to regularly change passwords.
- **Telnet/SSH:** Limit use of telnet because it transmits usernames and passwords in the clear. Limit telnet access using access-lists to predefined IPs. Use SSH instead.
- **HTTP/HTTPS/SNMP:** Centralize authentication and limit access to predefined IPs. Disable if not used.

Many control plane protocols, such as EIGRP, OSPF, HSRP, and GLBP, include peer authentication based on MD5 hashing. Vulnerabilities in ARP and DHCP can be addressed with switch capabilities to inspect and deal with maliciousness. DHCP snooping observes responses to ensure they come from the server, whereas Dynamic ARP Inspection looks

Troubleshooting Security Features

for and blocks spoofed ARP responses. Likewise, spanning-tree protection is available based on an understanding of the topology using technologies such as root guard and BPDU guard. The router can also protect against maliciousness by performing reverse path checking—making sure that packets arrive on the interface that would be used to route the reply.

The data plane is secured by controlling access, visibility, and flow. Keeping unauthorized users off the network is the role of network access control and 802.1x. Encryption and VLANs can be used to isolate traffic and prevent interception. Finally, traffic flows can be limited and inspected using access-list, flexible packet matching, IOS Firewall, and Intrusion Prevention Systems. IP source tracker allows for an easier, scalable solution to tracking DoS attacks compared to the traditional ACL. Zone-based security firewalls permit you to get granular in inspection and well-defined interface-based zone pairings to specify what traffic is permitted.

The IOS Firewall is easy to set up. An access-list is used to block all nonapproved traffic. Context-based access control(CBAC) is then used to modify the access-list, as replies to all outbound connections are allowed:

```
Ip access-list extended block
  Deny ip any any
Ip inspect name CBACInt f0/0
  Ip access-group block in
  Ip inspect CBAC out
```

Troubleshooting Security Features

The key issue with security features is that they limit traffic to create a security policy. This can work against the natural flow of troubleshooting, where the focus is on allowing communication. The issue is to recognize how the security policy compares to troubleshooting steps and to always work within the organizations change control system.

Troubleshooting Security Features

Troubleshooting the management plane, specifically authentication, can be tricky because it is possible to lock yourself out. The best approach is to have a backup plan to access the router—out-of-band access, a user to reset power, or a second authentication method. If no one is onsite, use the **reload in 10** command to schedule a reboot in 10 minutes before beginning work. It is also a good idea to allow local authentication (shown next) so that if access-list changes block access to RADIUS or TACACS+ there is still a way to login:

```
Aaa authentication default group tacacs+ local
Username brent password denise
```

SNMP uses UDP 161, and access-list blocking can be tested using extended traceroute on that port. SNMP can also be set up with access-lists and authentication to control access. Temporarily lifting these might also provide insight into any problems.

Troubleshooting the control plane comes down to neighbors. If a routing protocol doesn't see a directly connected peer, the problem is either a protocol issue or a firewalling issue. To verify that protocol traffic is passing, consider using **debug** to witness hellos (**debug ip eigrp packets**), or use the router as a protocol analyzer by using **debug ip packet access-list**. (The access list limits **debug** to just the traffic of interest.) The following example shows this done to analyze BGP traffic:

```
(config)#Ip access-list 101 permit tcp any any eq 179
Debug ip packet 101
```

The data plane includes support for user applications. Testing access can be accomplished with traceroute and telnet. Traffic is usually controlled using access-lists, so another way to troubleshoot connections is to log access-list matches. Access-list logging forces traffic to be processor switched and should be used in a limited manner. (Matches can be limited by narrowly crafting permit statements or though the established keyword, for instance.). ACL matches are forwarded to Syslog with this option, so used sparingly it is a good way to understand which line in the access-list is disposing of traffic. To set up logging, add the keyword **log** onto a ACL line. To see the denied traffic at the end of a list, for instance, add the following line to your ACL:

```
Deny ip any any log
```

CCNP TSHOOT 642-832

Quick Reference

Brent Stewart

Copyright © 2010 Pearson Education, Inc.

Published by:

Cisco Press
800 East 96th Street
Indianapolis, Indiana 46240 USA

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

First Digital Edition January 2010

ISBN-10: 1-58714-012-8

ISBN-13: 978-1-58714-012-9

Warning and Disclaimer

This ebook is designed to provide information about networking. Every effort has been made to make this ebook as complete and accurate as possible, but no warranty or fitness is implied.

The information is provided on an “as is” basis. The authors, Cisco Press, and Cisco Systems, Inc. shall have neither liability nor responsibility to any person or entity with respect to any loss or damages arising from the information contained in this ebook.

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

Trademark Acknowledgments

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

Feedback Information

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

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

We greatly appreciate your assistance.

Corporate and Government Sales

The publisher offers excellent discounts on this ebook when ordered in quantity for bulk purchases or special sales, which may include electronic versions and/or custom covers and content particular to your business, training goals, marketing focus, and branding interests. For more information, please contact: **U.S. Corporate and Government Sales** 1-800-382-3419 corpsales@pearsontechgroup.com.

For sales outside the United States please contact: **International Sales** international@pearsoned.com



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

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

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

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

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

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