Realtime publishers
"Leading the Conversation"

The Shortcut Guide™ To

Network Management for the Mid-Market

sponsored by

SOLARWINDS
NETWORK MANAGEMENT SOLUTIONS

Greg Shields

Introduction to Realtimepublishers

by Don Jones, Series Editor

For several years, now, Realtime has produced dozens and dozens of high-quality books that just happen to be delivered in electronic format—at no cost to you, the reader. We've made this unique publishing model work through the generous support and cooperation of our sponsors, who agree to bear each book's production expenses for the benefit of our readers.

Although we've always offered our publications to you for free, don't think for a moment that quality is anything less than our top priority. My job is to make sure that our books are as good as—and in most cases better than—any printed book that would cost you $40 or more. Our electronic publishing model offers several advantages over printed books: You receive chapters literally as fast as our authors produce them (hence the "realtime" aspect of our model), and we can update chapters to reflect the latest changes in technology.

I want to point out that our books are by no means paid advertisements or white papers. We're an independent publishing company, and an important aspect of my job is to make sure that our authors are free to voice their expertise and opinions without reservation or restriction. We maintain complete editorial control of our publications, and I'm proud that we've produced so many quality books over the past years.

I want to extend an invitation to visit us at http://nexus.realtimepublishers.com, especially if you've received this publication from a friend or colleague. We have a wide variety of additional books on a range of topics, and you're sure to find something that's of interest to you—and it won't cost you a thing. We hope you'll continue to come to Realtime for your educational needs far into the future.

Until then, enjoy.

Don Jones

Foreword

I've written a number of books—including several for Realtimepublishers—focused on network administration. Most of my books, however, have been focused on *enterprise* administration. In today's industry, *enterprise* has become a poorly defined, overused word. I usually use this term to mean *very* large, distributed organizations—ones with tens of thousands of users and annual IT budgets in the millions. Many of the management techniques and tactics I've discussed have been oriented for these larger companies, and might not be appropriate for smaller companies who are trying to work on an IT budget of "merely" a few hundred thousand dollars, or even just a "few" million. It's an important distinction to make: *Most* businesses in the world today *aren't* super-sized enterprises, and *most* businesses don't spend more on IT each year than some small countries have in their entire annual budgets.

That's why I think it's high time we release a book like the one you're now reading. In it, my good friend and colleague Greg Shields tackles network administration from the small- to medium-sized business' point of view. His focus is on the underlying technologies of network administration, and on using tools and technologies that are more practical for a midsized company's operations staff and budget. He'll take a detailed look at options such as open source tools in addition to helping you come up with a good set of specifications for various types of tools—whether you decide to build them yourself, go open source, or evaluate commercial products.

Greg's also going for a lot broader coverage in this short book than I have in other books, where I focused mainly on network device management. He's also looking at network performance monitoring, network analysis, security issues, and much more. That's appropriate, too; the largest companies tend to segregate these tasks across dedicated teams, but midsized companies tend to rely on a handful of "do everything" professionals who may even be responsible for server maintenance and the occasional desktop support call. This book's coverage will reflect those broad and varied responsibilities, although it'll stay firmly focused on the network infrastructure—sorry, no server and desktop support help, here!

The network's been taken for granted for far too long, and it's great to see experienced authors like Greg taking a heightened interest. With Voice over IP (VoIP), intense media streaming, and other new functionality, the network is working harder and harder to deliver the functionality companies require. That means the network administrator—for too long, the guy who was known only for hoarding IP addresses—has to be smarter and more efficient about making the network do its job unfailingly. Management refers to email as the "killer application;" without the network, of course, email would be dead in the water. I think this book takes just the right focus on the network: Keep it running, keep it efficient, and when things *do* go wrong, fix them fast. Greg assures me we're in for a fun ride, so let's jump right in.

Don Jones

Series Editor

Realtimepublisher.com

Copyright Statement

THE MATERIALS ARE PROVIDED “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE AND NON-INFRINGEMENT. The Materials are subject to change without notice and do not represent a commitment on the part of Realtimepublishers.com, Inc or its web site sponsors. In no event shall Realtimepublishers.com, Inc. or its web site sponsors be held liable for technical or editorial errors or omissions contained in the Materials, including without limitation, for any direct, indirect, incidental, special, exemplary or consequential damages whatsoever resulting from the use of any information contained in the Materials.

The Materials (including but not limited to the text, images, audio, and/or video) may not be copied, reproduced, republished, uploaded, posted, transmitted, or distributed in any way, in whole or in part, except that one copy may be downloaded for your personal, non-commercial use on a single computer. In connection with such use, you may not modify or obscure any copyright or other proprietary notice.

The Materials may contain trademarks, services marks and logos that are the property of third parties. You are not permitted to use these trademarks, services marks or logos without prior written consent of such third parties.

Realtimepublishers.com and the Realtimepublishers logo are registered in the US Patent & Trademark Office. All other product or service names are the property of their respective owners.

If you have any questions about these terms, or if you would like information about licensing materials from Realtimepublishers.com, please contact us via e-mail at info@realtimepublishers.com.

Chapter 1: FCAPS, Network Management Fundamentals, and Fault Management

Building an exceptional network involves the proper mix of skilled individuals, an intelligent design, the correct hardware, and the knowledge and experience to put it together correctly. In the largest of networks, that mix regularly produces some of the best networks in the world. However, there are hundreds of ways to design and run a network.

Money is often considered the differentiator between the best-run and the worst-run networks in the world. When businesses have plenty of cash to throw at their network infrastructure, they end up with industry experts who create best-in-class designs that leverage the market's greatest tools, don't they? Maybe, but having all the money in the world doesn't necessarily mean you're spending it well. Plenty of companies throw millions at their network and don't achieve the great things they had hoped. That suggests that wisdom, not just piles of money, plays a big role in making things great.

Businesses that operate in the small to midsize business (SMB) space and within the mid-market don't usually have the luxury to afford industry experts and the most expensive tools. In the mid-market, the people we label as "network engineers" wear multiple hats, doing server administration in the morning, acting as the Help desk in the afternoon, and working throughout the night on network administration. Thus, SMB and mid-market businesses must think wisely when making network infrastructure decisions. This guide is written specifically for you, the harried network administrator, in an attempt to show you some of the wisest tools and techniques to administer, troubleshoot, and automate your network infrastructure.

FCAPS and the Life Cycle of Proactive Management

The hardest part of any administrative activity is determining exactly what to administer. You're given the keys to the network and told, "Make sure it stays up." The relative stability of modern networks have buoyed the expectations of business management to assume the network is a utility function, just like the lights and power that run the building infrastructure. This increase in expectation on the part of business has resulted in a greater expectation of the network administrator.

To help determine both what and how to ensure that the network "stays up," this guide will discuss a number of topics relevant to the heavily burdened network administrator. You may not know the four fields of a Border Gateway Protocol (BGP) packet header or the intricacies of Open Shortest Path First (OSPF) convergence, but you do know that you've got a bunch of servers that need to interconnect with a bunch of workstations and generally never go down.

This guide is broken into four chapters. This chapter, Chapter 1, will discuss the concept of FCAPS, an acronym that describes a networking model used as a discussion framework for the types of things to keep an eye on inside your network. Each letter in FCAPS discusses one component of network management, all of which we'll explore in a minute. We'll discuss the breakdown of FCAPS and the take-aways from the model that you can use to start doing proactive network management. We'll also talk about fault management and some key tools and metrics you'll want to investigate to help identify and resolve network faults.

Chapter 2 expands on the tools discussed in Chapter 1 to encompass performance management. We'll talk about some key measurements that describe performance characteristics and validate a well-oiled network. The business impact of performance management will also be discussed, giving you the business drivers that justify monitoring and managing performance. And we'll finish with a few key tools and concepts for documenting and reporting on performance.

Chapter 3 continues to build on the core concepts, discussing configuration and security management. As we move from the concept of an ad-hoc network to a fully managed network, we'll focus on the items on your network necessary to bring under management to ensure a stable and secure configuration. In this chapter, we'll discuss some technologies that support a secure configuration and the techniques used by the smart networks to validate security.

Chapter 4 introduces the concepts of network troubleshooting and diagnostics. Often the most difficult part of network engineering is troubleshooting devices when they become problems. This chapter will outline a few tool suites for identifying problem devices, discuss the benefits and tools associated with IP address management and DNS problems, and outline network engineering applications that assist with the troubleshooting process.

What Is FCAPS and How Can I Leverage it in my Environment?

Before discussing effective network management practices, it helps to frame the conversation within an easily understood model. The FCAPS model was originally designed by International Telecommunication Union (ITU-T). This organization dates back to 1865 with original responsibilities of ensuring efficient and on-time production of high-quality recommendations covering all fields of telecommunications. In 1996, the ITU-T created the concept of the Telecommunications Management Framework (TMN), which was an architecture intended to describe service delivery models for telecommunication service providers based on four layers: business management, service management, fault and performance management, and element and configuration management.

Although the TMN was a valid initial mechanism for aligning telecommunications assets to business goals, the ITU-T refined the model in 1997 to include the concept of FCAPS. FCAPS expanded the TMN model to focus on the five functionally different types of tasks handled by network management systems: fault management, configuration management, accounting management, performance management, and security management.

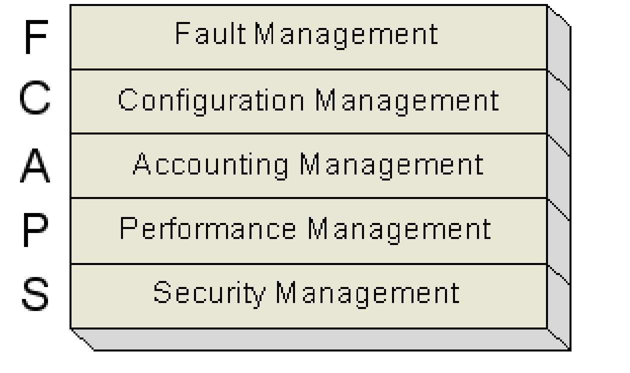

Figure 1.1: The FCAPS model explains the five functional layers of network management.

Although the initial release of the FCAPS model by the ITU-T was intended for telecommunications networks, it was the incorporation of FCAPS into the International Standards Organization (ISO) Open Systems Interconnect (OSI) model that highlighted its usability in describing the necessary functions of network management.

We'll discuss each of these layers in a minute. Within the FCAPS framework, we'll analyze the role of the network administrator and the necessary tools and techniques of network management that support it.

Moving from Reactive to Proactive

If this is the first you've seen or heard about the FCAPS model, you're probably wondering, "How does this affect my ability to better manage my network?" Knowing how to interconnect a series of routers and switches to create a functioning network involves one set of skills. However, as network complexity increases geometrically with an increase in the number of connections, properly managing that network as it grows and scales involves a whole new set of skills.

A little later, this chapter discusses a few technical concepts that enable network management and monitoring, such as the Simple Network Management Protocol (SNMP), SNMP traps, Management Information Bases (MIBs), and the use of Syslog. The use and functionality of these core technologies, among others that later chapters will discuss, are primarily based on the recommended network needs identified by the FCAPS model.

As information networks are growing in size and reliability and as business expects greater uptime and performance from their network backbone, there is an increased need to proactively monitor activity, notify when or before problems occur, dynamically reroute based on conditions, and provide alerting to administrators. To ensure this, it is imperative that the network administrator move from a ***reactive*** approach to network management to a ***proactive*** approach. The tools we will discuss over the course of the next few chapters will enable the network administrator to achieve this goal.

Ad-Hoc Management

Most networks begin small. And in those small networks, it is often the work of a small group of trusting individuals to complete the network build and administration tasks. Formalized mechanisms of change control often don't exist in the smallest of networks as the cost of administrative overhead required to support them does not outweigh their benefits.

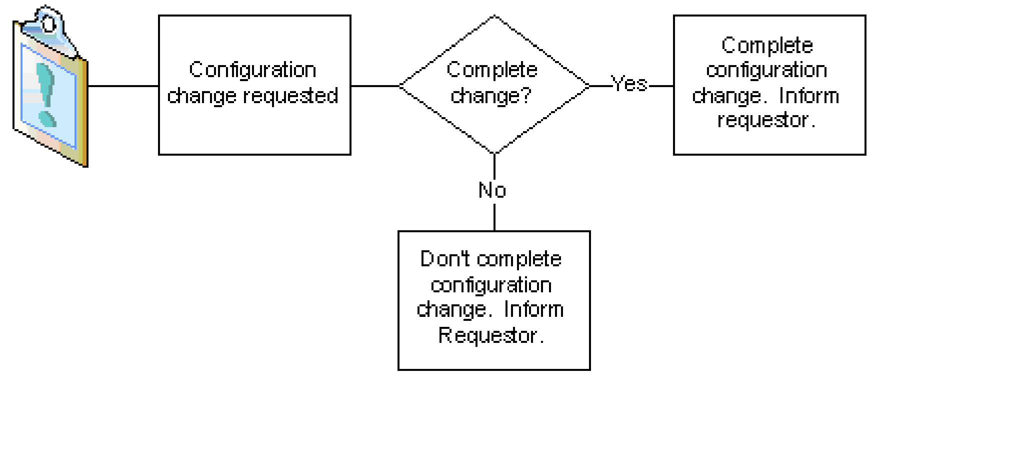
Configuration change requested
Complete change?
Yes
Complete configuration change. Inform requestor.
No
Don't complete configuration change. Inform Requestor.

Figure 1.2: In the ad-hoc network, the decision to complete a change lies completely with the administrator.

However, as the size and complexity of the network increases, the number of configurations and the number of *configurators* increase as well. Adding to this complexity is the highly text-based interface of most business-class network devices.

```
firewall multiple-vlan-interfaces
firewall module 1 vlan-group 1
firewall vlan-group 1 200-202
interface vlan 200
ip address 10.0.1.3 255.255.255.0
standby 200 ip 10.0.1.4
standby 200 priority 110
standby 200 preempt
standby 200 timers 5 15
standby 200 authentication Secret
no shutdown
interface vlan 201
ip address 10.0.2.3 255.255.255.0
standby 200 ip 10.0.2.4
standby 200 priority 110
standby 200 preempt
standby 200 timers 5 15
standby 200 authentication Secret
no shutdown
```

Listing 1.1: Implementing even a simple switch configuration involves numerous text-based lines of code.

As most companies grow their business and correspondingly grow their networks, the number of configuration items grows geometrically with the size of the networks. Businesses must retreat from ad-hoc management to some mechanism of configuration control to ensure network stability. The FCAPS model can outline the necessary management tasks.

The next five sections will discuss the five functionalities of FCAPS and their associated management tasks. As you move your network from ad-hoc management to full-management, consider the incorporation of network tools that facilitate these management tasks.

Fault Management

The F in FCAPS discusses the management tasks associated with fault management. Among other traits, an effective fault management system will recognize that a problem has occurred, alarm the administrator when that recognition occurs, and provide information as to the location and manner of the fault. Twelve management tasks are identified by the FCAPS model as necessary for a successful fault management system:

- Fault detection
- Fault correction
- Fault isolation
- Network recovery
- Alarm handling
- Alarm filtering
- Alarm generation
- Clear correlation
- Diagnostic test
- Error logging
- Error handling
- Error statistics

Configuration Management

Configuration management is the concept of ensuring a consistent, repeatable, and auditable configuration on all devices that make up the network. Effective configuration management ensures that devices contain the right configuration based on policy. It also provides a mechanism for rapid return to operations of faulted devices, as an effective configuration management tool will store each device's configuration in a separate searchable repository. In today's environment of stringent compliance regulations, only through effective configuration compliance will a network pass an auditor's review. Consider the following management tasks when choosing any configuration management system:

- Resource initialization
- Network provisioning
- Auto-discovery
- Backup and restore
- Resource shut down
- Change management
- Pre-provisioning
- Inventory/asset management
- Copy configuration
- Remote configuration
- Automated software distribution
- Job initiation, tracking, and execution

> 🕮 A detailed discussion of configuration management will continue in Chapter 3.

Accounting Management

Of the five items in FCAPS, accounting management tasks are potentially the least relevant for many networks. Although some network backbone architectures and organizations honoring Service-Oriented Architectures (SOAs) may incorporate charge backs and cost-based servicing, most networks in the mid-market likely don't incorporate these accounting systems. However, some components of accounting management—such as the need to track resource use for metrics generation—should be a component of any mature network. The following list highlights the eight considerations for tools that enable accounting management:

- Track service/resource use
- Cost for services
- Accounting limit
- Usage quotas
- Audits
- Fraud reporting
- Combine costs from multiple resources
- Support for different accounting modes

Performance Management

Performance management involves the effective monitoring of a network's response time and the proactive management of needed upgrades to support its users. Performance management expands upon simply answering the question, "Why is the network so slow?" It involves proactively analyzing a network's activity and making informed business decisions about expansion before performance becomes critical. Businesses who engage in monitoring and taking action based on performance management can recognize substantial return on investment (ROI) based on prevention of loss of worker efficiency due to network conditions. When looking at performance management systems, look for the following traits:

- Utilization and error rates
- Performance data collection
- Consistent performance level
- Performance data analysis
- Problem reporting
- Capacity planning
- Performance report generation
- Maintaining and examining historical logs

Security Management

Aligned with the needs of configuration management, the tenants of security management ensure the integrity and reliability of the network. Many network devices by default enable security through a shared password concept, which can be a violation of established security policies. Enabling successful security management means segregating the roles and responsibilities of administrators and users, logging their activity, and ensuring the privacy of data on the network. An effective security management system will provide mechanisms for security administrators to easily record network activity and parse that activity for anomalies. Consider the following activities as critical for an effective security management system:

- Selective resource access
- Access logs
- Data privacy
- User access rights checking
- Security audit trail log
- Security alarm/event reporting
- Take care of security breaches and attempts
- Security-related information distributions

Choosing the Right Suite of Tools

It should be obvious that no single management system can likely handle each and every one of these network management activities. However, a suite of tools can provide for the necessary subset of capabilities needed by your business. When considering a network management suite of tools, your business should consider a trade study process for identifying your requirements and separating requirements from those items considered “nice to have.” The process of completing a trade study can run from an informal intra-group decision-making process to the use of a formalized trade study framework.

Product Scoring		**Product 1**		**Product 2**		**Product 3**	
	Weight	**Score**	**Weighted Score**	**Score**	**Weighted Score**	**Score**	**Weighted Score**
Fault Detection	9.0	8	72.0	10	90.0	9	81.0
Fault Correction	9.3	4	37.3	4	37.3	5	46.7
Diagnostic test	10.3	6	61.8	8	82.4	2	20.6
Error statistics	4.3	8	34.4	2	8.6	5	21.5
Network recovery	5.3	4	21.3	8	42.7	3	16.0
Total		30.0	***226.9***	32.0	***261.0***	24.0	***185.8***

Table 1.1: An example of a formalized trade study that weighs the need for features and compares each product’s capability to deliver on that feature.

Network Management Fundamentals

Four fundamental technologies are enablers for much of the characteristics discussed in the FCAPS model. These technologies are SNMP and the related concept of SNMP traps, MIBs, and Syslog. These four technologies interoperate to provide monitoring capabilities for devices, notification to administrators when preconfigured conditions occur, and storage of device and log information in searchable formats. Throughout the rest of this guide, we will refer back to these core technologies as we discuss additional tools available to network administrators. Most important, these technologies are key in moving a network from an ad-hoc and reactive mode to one that is fully and proactively managed.

SNMP

SNMP describes a network protocol as well as an information framework used to provide remote monitoring and configuration capabilities for network devices. The main purpose of SNMP is to enable the centralization of network device management and monitoring through a common language across all devices.

SNMP-enabled devices are said to be SNMP agents. Enabling SNMP on these devices provides the capability of being remotely interrogated by an administrator through a Network Management Server. The NMS incorporates high-level management software that typically has the capability of storing agent device configurations within a management database. The NMS additionally has the capability of receiving notifications from devices when preconfigured conditions occur. This notification is called an SNMP trap and will be discussed in the next section.

The utility of the SNMP protocol is in its extensibility. As the protocol and framework are device-independent as well as NMS-independent, this allows for the interconnection of all SNMP-capable devices of any vendor into a single domain of management. The NMS chosen for management of the devices is not reliant on the device vendor and multiple NMS' can be used to manage the same devices.

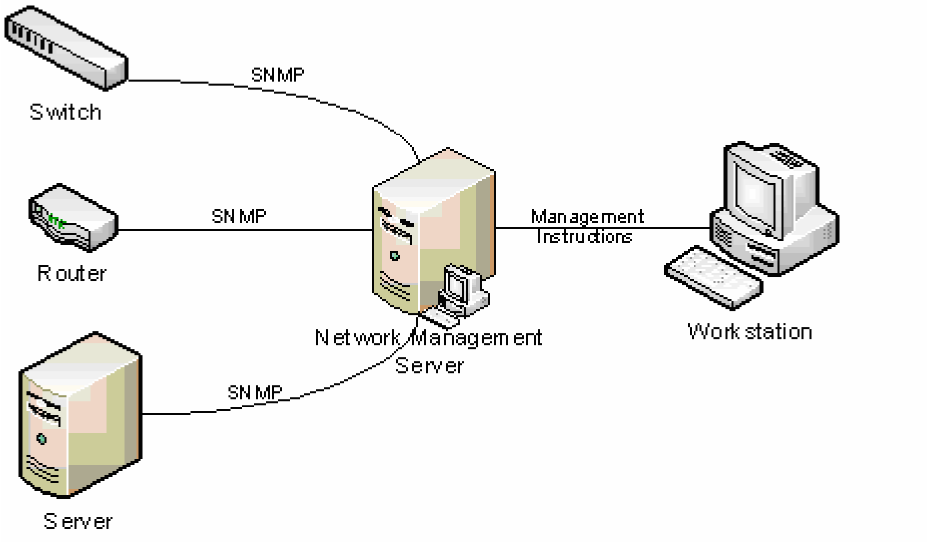

Figure 1.3: SNMP devices report information to the NMS. Administrators instruct the NMS to collect and update information on managed devices.

SNMP can be configured to allow for writing and/or updating of configurations on managed devices as well. This capability, leveraged through the NMS, greatly enhances the configuration control capabilities within a managed network. Using SNMP to manage configurations means that configuration files can be centrally stored within the NMS database. The NMS can be used as a gatekeeper for devices, ensuring that only the approved and agreed-upon configuration is enabled on the device. The configuration database can additionally be backed up using typical network backup devices, ensuring that configurations can be rapidly restored in the case of a disaster event.

Listing 1.2 shows the commands needed for a basic SNMP configuration on three different types of Cisco devices. The first is a Cisco IOS-based router; the second is a Cisco CatOS-based switch; the third is a Cisco PIX firewall.

```
router1#config terminal
router1(config)#snmp-server contact admin@abccorp.com
router1(config)#snmp-server location Downtown Office 2nd Floor Router
router1(config)#snmp-server chassis-id 123456
router1(config)#snmp-server community public ro
router1(config)#snmp-server community private rw

Console> (enable) set snmp community read-only public
Console> (enable) set snmp community read-write private

firewall1#config terminal
firewall1(config)#snmp-server contact admin@abccorp.com
firewall1(config)#snmp-server location Downtown Office 2nd Floor Router
firewall1(config)#snmp-server community public ro
firewall1(config)#snmp-server community private rw
```

Listing 1.2: Commands for a basic SNMP configuration on three different types of Cisco devices.

MIBs

MIBs are collections of device characteristics that are available for reading and writing via SNMP. Think of a MIB as a small database that contains all the characteristics of the device. Individual device settings within a MIB are called *MIB variables*. Within a device, all MIB variables are collected into a single document object called a *MIB module*. It is within these MIB modules that each managed attribute of a device is described and its interfaces are discussed. MIB modules are specific to each particular device and must be preloaded into the NMS for the NMS to be able to manage the object.

MIB modules are typically delivered with the network device or can be downloaded from the device manufacturer’s Web site. As the number of MIB modules for a complex network can get extremely large, effective management of these MIB modules is one of the primary features of a good NMS.

Although MIBs are not addressed by their location on a network, there is a hierarchy to MIB modules, called the *MIB tree*, which allows for each MIB across all device types to be uniquely addressed. Managed by the Internet Assigned Numbers Authority (IANA), this unique addressing ensures that two manufacturers cannot use the same MIB module information. The unique address for any MIB module is called its Object ID (OID) and is represented by a long string of dot-separated numbers.

This long string of numbers references the object's location on the tree. The first number in the OID references its position at the top of the tree—the most general descriptor—and each subsequent number further defines the object in relation to its position on the tree.

> ✎ For example, the OID for a Synoptics 3000 concentrator is 1.3.6.1.4.1.45.1.3.2. The first five numbers of this OID reference iso(1).org(3).dod(6).internet(1).private(4). As you can see, the tree is exceptionally general at the top levels, not even reaching the Internet until the fourth level. Substantial numbers of objects not typically thought of as network devices can be managed by SNMP. Some examples are water-level indicators, air quality measuring tools, and entryway monitors.

Using the OID as the addressing for that device configuration, four types of communication transactions can occur between the NMS and an agent. The following list highlights these four types:

- Get—The Get operation is initiated by the NMS to retrieve information from a managed device. Because Get operations require exact addressing, the NMS will provide a complete OID to the agent to locate the characteristic of interest. The agent will respond with the value of the requested MIB variable and the NMS will store this information and/or notify the administrator of the result.
- Getnext—Similar to a Get operation, the Getnext operation is used when additional information is needed. Where the Getnext operation is different is that it does not need to provide exact addressing information for the requested characteristic. Instead, Getnext will request the next characteristic in the tree.
- Set—SNMP can be used for both read and write operations. The Set operation is used when the administrator wants to update or change the value of the requested characteristic.
- Trap—An SNMP trap is a way for an agent to notify the NMS that a preconfigured condition has occurred. Traps are the main component of the notification piece of SNMP.

SNMP Traps

As stated earlier, an SNMP trap describes the ability for an agent to notify the NMS that a preconfigured condition has occurred. The trap is a unidirectional notification from agent to NMS that includes the OID of the characteristic of interest and its associated MIB value. Trap characteristics must be configured within the NMS and relayed to the agent prior to the condition occurring. The agent must be preconfigured with a network location to send traps. Multiple locations can usually be configured for redundancy purposes.

> For example, let's assume the agent on a managed device is configured to watch the internal temperature conditions on that device. The administrator sets a threshold of 100° for those conditions. When the temperature exceeds 100°, the device will initiate a trap to the NMS notifying it (and the administrator) that the condition has occurred. Typical NMSs provide the capability of setting trap characteristics such as value, time exceeded, amount of time exceeded prior to trapping, and return-to-normal thresholds. These added trap characteristics prevent situations occurring in which, for example, the temperature rapidly bounces back and forth between 99° and 101° and the administrator is repeatedly notified.

Where SNMP traps make the biggest contribution in moving a network from reactive to proactive is in the ability to link trap information to administrator alerts. Typical NMSs can be configured to notify administrators via email, page, or SMS message when a trap condition occurs.

Remember that sending a trap does not accomplish anything towards fixing the condition on the device. It is still up to the administrator to resolve the issue once the notification of the issue has been raised. In some situations, however, if SNMP set commands are enabled, the NMS can be preconfigured to perform an action when a trap occurs. That action can be to change a MIB value or even shut down the system.

These action capabilities available to the administrator are based on the intrinsic capabilities built-in to the device and its SNMP interface as well as the feature sets available on the NMS. The decision to purchase an NMS may include a determination whether that NMS has these sorts of automatic capabilities built-in to the system.

Often, these capabilities require some code development. Although it is relatively easy to set up and configure an NMS to handle SNMP and SNMP traps, the greatest portion of any NMS installation is usually in the individual tuning of alerts.

> Many an administrator has lost nights of sleep due to an overly sensitive NMS alerting that wakes the administrator in the middle of the night for a non-critical event. Consideration should be made for quality of life issues for any administrator whose pager is linked to an overly sensitive NMS!

Syslog

Syslog is to centralized system logging as SNMP is to centralized system configuration control. Syslog is a mechanism for sending event messages from managed devices to a centralized server that runs the Syslog service. Although there are some minor differences between the various forms of Syslog that handle Windows, Cisco, UNIX, and other device messages, many native and commercially available Syslog applications can handle event messages from all these types of devices.

The Syslog service is configured on a server within the network environment to accept messages. The nature of the Syslog service is such that the transmission of message information is unidirectional from the sending device to the Syslog server. There is no acknowledgement of receipt. The service stores these messages in a searchable database that can be queried by a management workstation. Typical types of Syslog messages include system error messages, statistics, system warnings, security and access notifications, and access denied notifications, among others.

Syslog is particularly useful for ensuring security in networks due to the external nature of the log collection. If a network device is compromised by an intruder as part of an external attack, one of the tasks usually completed by that intruder is to remove any record of their presence on the device. This process of wiping clear the logs can prevent a security administrator from detecting the compromise of the device and hinders the ability to track the entry point and location of the intruder. By enforcing that all system logs are immediately copied both to the local device as well as to the Syslog server, there is a greater chance of the actions of the intruder being detected.

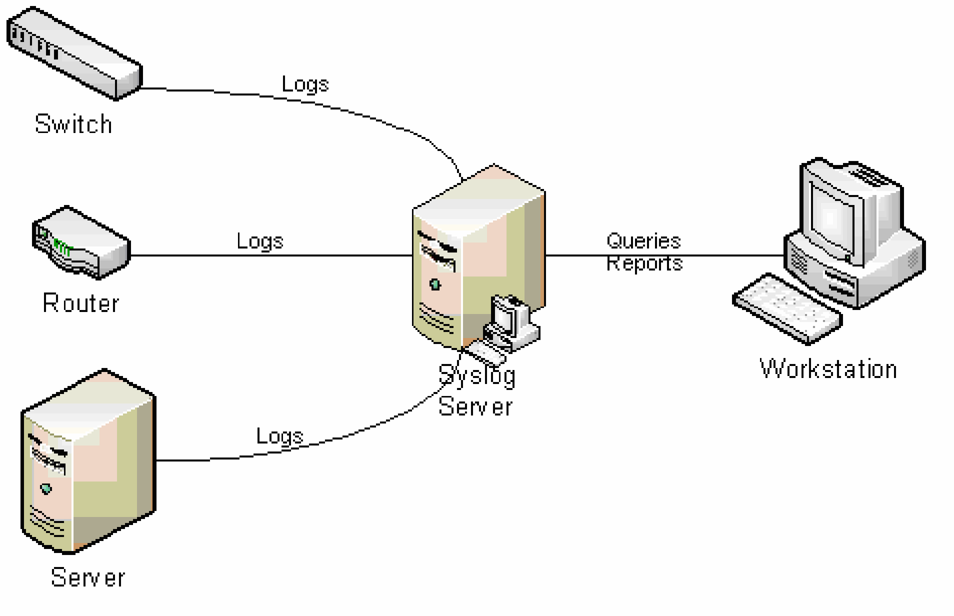

Figure 1.4: Similar in architecture to a network's SNMP infrastructure, Syslog centralizes event log information.

Although there is no standard that defines the content of a Syslog message, most Syslog messages consist of three parts:

- PRI—The PRI part of a Syslog message determines the message's priority. This priority is a numerical value that is a mathematical combination of the facility generating the message and the severity of the message.
- HEADER—The HEADER part of a Syslog message contains the timestamp denoting the creation of the message along with either the IP address or the hostname of the sending device. This part is used to identify the date, time, and origin of the message.
- MSG—The MSG part of a Syslog message contains the text of the message itself. The MSG part can typically have two fields, called the TAG field and the CONTENT field. The TAG field typically stores the name of the application or process that generated the message. The CONTENT field typically contains the message itself.

At first blush, the lack of standardization in the Syslog format can appear to be a weakness in its implementation. However, it is within this lack of standardization that Syslog gains its popularity and its near universal acceptance. Syslog requires only priority and origin information for each message, so it becomes exceptionally easy to use for all types of devices from Cisco to Microsoft Windows to all flavors of UNIX to network hardware appliances. This universal acceptance of the Syslog format means that Syslog can be used in the collection of nearly all device information on the network.

Often, problems on the network involve more than one device and correlating the event information for each of these devices can be cumbersome. But leveraging a centralized mechanism for storing event information for all devices means that a longitudinal timeline across all devices can be easily queried. This gives the administrator a more holistic view of the network and provides better detail into identifying the problem.

When choosing a Syslog system, it is important to choose one that has the capability to handle Syslog messages from each of the types of devices connected to it. As Listing 1.3 shows, different device types can have slightly different log formats. The onus for reading these device formats is on the application hosting the Syslog service.

```
*Mar 6 22:48:34.452 UTC: %LINEPROTO-5-UPDOWN: Line protocol on
Interface Loopback0,changed state to up

2000 Feb 21 12:00:27 PST -07:00 %SYS-4-SYS_HITRFC: 65% traffic detected
on switching bus

<110> Oct 16 08:58:07 64.103.114.149 CisACS_13_AdminAudit 18729fp11 1 0
AAA Server=tfurman-w2k,admin-username=local_login,browser-
ip=127.0.0.1,text-message=Administra
tion session finished,
```

Listing 1.3: Examples of Syslog messages from Cisco IOS and CatOS devices. Note the differences in their format.

There are multiple applications hosted on Microsoft Windows as well as various flavors of UNIX that can support these formats. Those hosted on Microsoft Windows are typically non-native applications layered on top of the operating system (OS). For UNIX, the native syslogd daemon is the standard. Tradeoffs in usability, searchability, a graphical interface, and cost exist between a native UNIX solution and a non-native Microsoft Windows solution. During your decision-making process, consider these tradeoffs, the types of connected devices, and the querying and reporting capabilities you will find necessary in completing your log analysis tasks.

Listing 1.4 shows the commands needed for a basic Syslog configuration on three different types of Cisco devices. The first is a Cisco IOS-based router; the second is a Cisco CatOS-based switch; the third is a Cisco PIX firewall.

```
router1#config terminal
router1(config)#logging 192.168.0.200
router1(config)#service timestamps log datetime localtime show-timezone
msec
router1(config)#logging facility local1
router1(config)#logging trap warning

Console> (enable) set logging timestamp enable
Console> (enable) set logging server 192.168.0.200
Console> (enable) set logging server 192.168.0.200
Console> (enable) set logging server facility local1
Console> (enable) set logging server severity 4
Console> (enable) set logging server enable

firewall1#config terminal
firewall1(config)#logging timestamp
firewall1(config)#logging host 192.168.0.200
firewall1(config)#logging facility 17
firewall1(config)#logging trap 4
firewall1(config)#logging on
```

Listing 1.4: Commands for a basic Syslog configuration on three different types of Cisco devices.

Key Steps in Identifying and Correcting Faults

Fault management is the first item in the FCAPS model, dealing specifically with the identification, isolation, and correction of network faults. Although proactive management is always a consideration with administering a network, fault management specifically deals with the issue of post-incident remediation.

Typical fault management systems use the four steps of fault detection, event/alarm generation, fault isolation, and fault correction to break down the complicated task of identifying and resolving the fault.

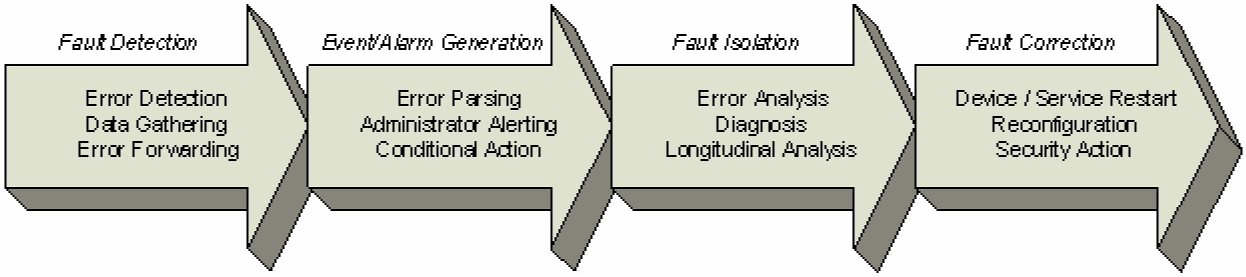

Figure 1.5: The four phases of fault management.

Fault Detection

As Figure 1.5 illustrates, multiple actions can occur at each phase when a fault occurs. The major difference between reactive network management and proactive network management occurs in the first two phases. Business networks with small numbers of devices likely do not have an NMS in place to notify administrators when a fault occurs. In these cases, often the mechanism for fault detection is when a user determines a change in the state of their connection. They contact the appropriate personnel in the IT organization—often a Help desk—and notify them that they have noticed an anomaly on the network.

For the smallest of networks, this ad-hoc mechanism of notifying IT personnel of a network fault is sufficient. In these networks, uptime requirements and user expectations are low for the health of the network.

Ad-hoc notification, however, does have the impact of creating a high number of false positives. As non-technical users are relied upon for the notification of faults, they can overuse their responsibility, notifying the IT department during non-fault situations. This has the effect of causing additional workload on IT to track down phantom problems. Ad-hoc notification also works predominantly with up/down situations as more complicated performance, stability, and security calculations are difficult.

The proactive network leverages technologies such as SNMP and Syslog to provide device notification of fault conditions. Two types of fault management, active and passive, can be configured within the NMS for notification:

- Passive fault management—In passive fault management, SNMP-enabled devices notify the NMS when a preconfigured condition has occurred. Passive fault management can track when conditions vary from nominal, but it is reliant on the administrator to preconfigure SNMP agents with the business' definition of nominal activity. This type of fault management is highly successful for identifying problems not associated with a device outage: performance issues, individual service or interface problems, out-of-boundary changes in network traffic flow, and so on. However, passive fault management suffers from an inability to notify during a complete device outage. After all, if the device is non-functional, the SNMP agent cannot raise a notification to the NMS.
- Active fault management—Active fault management is a mechanism for notifying when those outage conditions occur. Active fault management typically runs concurrent with passive fault management and usually from the same NMS. In active fault management, the NMS sends a network PING command to each managed device on a regular basis and listens for the reply. If the device does not reply after a preconfigured interval, the active monitoring will notify the administrator of a device outage. For this reason, active fault management is often also referred to as "up/down monitoring."

When choosing a fault management system, consider one that can provide both active and passive management. Well-designed fault management systems can also provide some internal logic that enables the network administrator to pre-generate a series of if/then statements associated with the recognition of faults. Six categories of logical conditions and actions can make up a typical fault management system.

- Time of day—Most Service Level Agreements (SLAs) for business networks assign a period of network operability. This period can occur from typical business hours (8:00am to 5:00pm, Monday through Friday), to full 24/7/365 operations. In either case, there are usually times in which alerting is nonsensical considering the actions taking place on the network, such as during maintenance windows. Time of day logic allows the administrator to assign times when fault detection will be disabled or enabled but lacking notification.
- Trigger condition—The trigger condition for a fault is the administrator-configurable variable that identifies faults of interest. Device MIBs can notify on dozens or hundreds of possible conditions but only a few will be relevant per device for each particular network. Effective NMSs incorporate robust logic for trigger condition creation.

- Reset condition—Reset conditions are the opposite of trigger conditions. When a device has gone over a threshold and a fault has been identified, there must be some mechanism of automatically resetting the notification of the fault to nominal. Similar to trigger conditions, effective NMSs incorporate similar robust condition logic to enable creation of reset conditions.
- Alert suppression—In some situations, the administrator might want conditional-based suppression of fault notification. In those cases, alert suppression logic allows for the creation of values that prevent fault identification. This may be based on known issues within the network, such as known problems with particular network devices.
- Trigger and reset actions—Robust NMS systems should provide for a suite of actions to occur when a fault is recognized and again when that fault has been reset. The next section will discuss examples of these actions.

In all cases, these data gathering and error forwarding actions are components of the managed device. The managed device will forward the error, usually through SNMP or Syslog, to the NMS for processing.

Event/Alarm Generation

When a fault occurs and the fault information has been forwarded to the NMS, the NMS has a few actions of its own to accomplish. First, that error must be parsed and compared with the pre-generated logic as explained in the previous section. When alerting matches have occurred, the NMS will trigger a notification to the administrator through one of many mechanisms. Some of these mechanisms for notification include:

- Sending an email message
- Sending an SMS message to a cell phone or pager
- Playing a sound or recorded message on the management workstation
- Logging the alert to the Network Event log
- Logging to a text file
- Sending a Syslog message
- Sending an SNMP trap
- Logging the alert to a Microsoft Windows event log
- Sending a Microsoft Windows Net-Message
- Executing an external program
- Executing a script
- Speaking an alert message using a text-to-speech engine

As you can see, numerous capabilities for notification exist. A robust suite of notification actions on fault triggering and fault resetting ensures that administrators are properly notified no matter where they may be and no matter what they are using as their primary notification device. Additionally, robust external program and script execution further enhances an NMS' proactive capabilities, enabling common triggers to be automatically resolved by the system through execution of an application of a program or a script.

> ✎ For example, assume that a router has a specific networking problem that regularly occurs and the known resolution is to flush the ARP cache. In that case, in which both the problem and the resolution are known, a trigger condition can be set up to run a script that automatically flushes the ARP cache when the condition occurs.

> Be careful with automatic actions! Configuring scripts to automatically fire when conditions occur can sometimes exacerbate a problem or band-aid it without actually resolving it. What if that router problem is actually an external hacker attempting to infiltrate the network? Setting the action to automatic may prevent the administrator from noticing the attempt.

Fault Isolation

Fault isolation is a step involved with complicated problems where identifying the root cause is difficult and where the problem may span multiple devices. In these cases, a root cause analysis is often the troubleshooting approach to tracking down the problem.

For complicated problems, fault isolation can involve substantial error analysis and deep diagnosis of the problem. These sorts of deep-dive problems can occur in large networks with multiple administrators who manage devices within separate but connected domains of management.

In these cases, the task of fault isolation can involve review of logs across multiple devices and across multiple management domains. Aggregation of Syslog and SNMP data across multiple devices and management domains may be necessary to track down the problem. A *longitudinal* or *cross-device* approach to problem isolation can significantly improve results. In a longitudinal approach, logs are correlated across multiple devices and listed in time order across all devices. Using this approach, it is possible for the administrator to get a "big picture" view of the network and its interconnected devices.

One other approach is the use of network mapping to provide a graphical representation of the interconnected devices and show their status and interrelation. Many network management tool suites have the capability of automatically discovering devices on the network and auto-generating network maps of those devices and their connections.

Fault Correction

Fault correction is the obvious last step in the process of fault management. With fault correction, the administrator has identified the fault—often the most difficult and time-intensive part of the process, recognized the corrective actions through analysis, and is ready to commit the corrective action.

Three items are identified as potential actions involved with fault correction:

- Device/service restart
- Reconfiguration
- Security action

All three of these actions involve a potential change on the part of the device. Approval and personnel notification of these changes should be disseminated through some form of change management. Additionally, tools and technology exist to ensure that the change is logged into a configuration management database.

> 🕮 Chapter 3 will discuss change management fully and the tools that exist to enable this storage of configuration information.

Key Metrics for Fault Management

The last sections of this chapter discuss useful business metrics for measuring a network and the faults that may occur within that network. These metrics are useful for trending purposes so that businesses can understand the health of the network and recognize over the long-term when that health begins to degrade. This trending analysis is especially necessary during periods of business growth and the resulting network growth that accompanies it. As additional nodes on a network come online, the trending of metrics in fault management can provide the information needed to make informed decisions regarding purchasing and network expansion.

Mean-Time Between Failures

Mean-Time Between Failures (MTBF) is a hardware-based metric provided by manufacturers to customers to denote the average amount of time that occurs between failures on a particular device.

> ✎ For example, when you purchase light bulbs, you are given the option of the “regular” brand that last for 1 year or you can pay a premium for “long life” bulbs that may last for 10 years.

Although some manufacturers will no longer release their MTBF statistics in today’s marketplace, analyzing industry studies on MTBF metrics for the products in your environment will help you make informed purchasing decisions. MTBF statistics directly relate to the network uptime metric and all relate to loss of worker productivity associated with network failure.

Mean-Time To Restore

Mean-Time To Restore (MTTR) is typically a metric internally defined by the business. Some businesses can sustain a multiple-day outage with little affect on operations. Some networks cannot survive more than a few minutes of outage before the dollars-per-minute of downtime grows critically expensive.

It is important for businesses to very early determine their MTTR metric for each network service based on the pain associated with the loss of that service. This is critically necessary as disaster recovery and fault prevention technologies come in many shapes, sizes, and costs. There is an inverse geometric relation between the sensitivity to outage (and the related shortened MTTR) and the cost to implement preventative mechanisms.

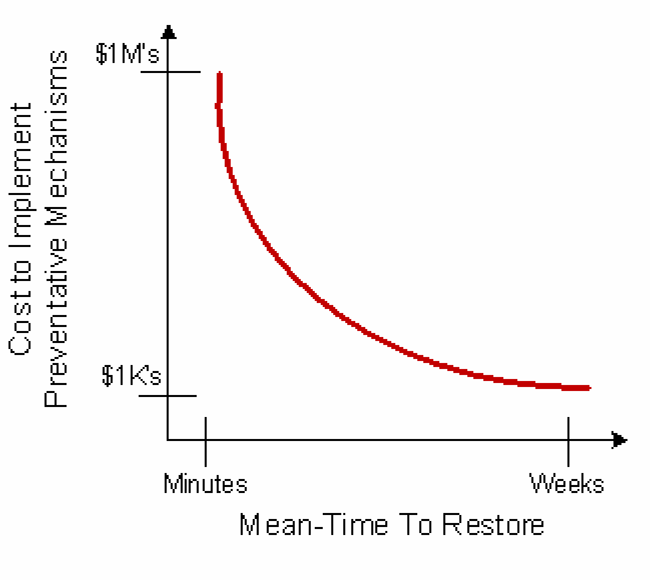

Figure 1.6: The less downtime the business can handle, the more expensive it is to implement technology to prevent it.

Network Uptime

Network uptime is the final metric associated with fault management. Network uptime can be defined as an integer number associated with the amount of time between individual device outages. More often, however, it is related to the amount of time that users on the network are not able to accomplish their daily tasks due to a network problem. This is more of a holistic approach to total system availability.

This number is different than MTBF because MTBF numbers typically relate to an individual device and its outage potential. In business networks, redundancy is typically incorporated into the network design that lessens the effect of a device outage on the ability for workers to accomplish work. Network uptime is a useful metric for reporting to management based on the ability for the network administrator to ensure worker productivity.

Relating to Your Business

Obviously, all the metrics in the world matter little if they mean nothing to your business. Depending on the needs of your business, the process by which you identify and resolve faults and the means by which you measure your success will be driven by internal needs. What is important to take away is that there are applications available in the marketplace that can move your network from reactively dealing with faults to proactive notification when faults occur and automated response to deal with them.

The next chapter will take what we've explored thus far about network management and relate it to another item in the FCAPS model—the concept of performance management. Chapter 2 will talk about the important measurements you can employ to determine the usability of your network infrastructure as well as the related business metrics. We'll go over some performance management concepts for documenting your existing environment and planning for expansion and discuss tools that can help you ensure a well-oiled network.

Download Additional eBooks from Realtime Nexus!

Realtime Nexus—The Digital Library provides world-class expert resources that IT professionals depend on to learn about the newest technologies. If you found this eBook to be informative, we encourage you to download more of our industry-leading technology eBooks and video guides at Realtime Nexus. Please visit http://nexus.realtimepublishers.com.