

Realtime
publishers

"Leading the Conversation"

The Shortcut Guide[™] To



**Network
Management
for the Mid-Market**

sponsored by



Greg Shields

Chapter 3: Configuration Management and Security	41
Key Steps in Managing an Environment Configuration.....	42
Establish a Baseline	43
Document Configuration	43
Control Change	44
Audit Environment.....	45
Ad-Hoc/Manual Configuration vs. Managed Configuration	46
Configuration Standardization	46
Configuration Backup and Archival	47
Post-Incident Restoration.....	48
Policy-Based Configuration.....	48
Inventory and Mapping.....	49
Rogue Device Identification and Adjudication.....	49
Provisioning	50
Deprovisioning.....	50
User Access.....	50
Business Drivers for Configuration Management	51
MTTR Reduction	51
Loss of Business Revenue	51
Security	52
Regulation and Compliance	52
Auditing Requirements	53
Personnel Turnover.....	54
Supporting Technologies	54
Understanding Security Management.....	56
Practicing Good Network Security	56
SNMP Community Strings and SNMP Weaknesses	57
Port Scanning and Port Minimization.....	57
Penetration Testing	57
Vulnerabilities, Exploits, and Patches	58
Configuration and Security Management Provide Measurable Benefit	58

Copyright Statement

© 2007 Realtimedpublishers.com, Inc. All rights reserved. This site contains materials that have been created, developed, or commissioned by, and published with the permission of, Realtimedpublishers.com, Inc. (the "Materials") and this site and any such Materials are protected by international copyright and trademark laws.

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 Realtimedpublishers.com, Inc or its web site sponsors. In no event shall Realtimedpublishers.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.

Realtimedpublishers.com and the Realtimedpublishers 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 Realtimedpublishers.com, please contact us via e-mail at info@realtimedpublishers.com.

[**Editor's Note:** This eBook was downloaded from Realtime Nexus—The Digital Library. All leading technology guides from Realtimepublishers can be found at <http://nexus.realtimepublishers.com>.]

Chapter 3: Configuration Management and Security

The average mid-size company of 250 employees typically serves the same number of workstations with about 25 servers. Those nodes on the network are interconnected by about nine network devices, through firewalls, switches, and routers. For a network of that size, the average network device configuration contains about 300 lines per device. Multiplying those two numbers, you get the potential for more than 2700 individual configurations, just to connect a relatively small number of devices.

The big question is this: In a critical situation, could you rebuild those 2700 lines purely from memory? It's the implementation of configuration management into your network environment that helps you answer that question in the affirmative.

Chapter 2 discussed performance management in relation to the FCAPS model of network management. The chapter discussed how managing performance in a network can be virtually impossible without a baseline to measure it by, and talked about how to use your NMS to measure the changes in performance from your baseline and how those changes in performance can trace back to configuration inconsistencies or other underlying problems. The chapter also brought forward some good technical and business metrics that illustrate network performance and validate it to your business leaders.

This chapter will move away from the P in FCAPS and focus on the C and the S—configuration management and security. This chapter will discuss how you can use a good NMS to consistently manage, store, and audit the configuration of devices on your network. We'll explore the four steps in establishing and maintaining an environment configuration and relate those to the underlying financial reasons why configuration management has business relevance. The chapter will go over a set of features that an effective NMS should incorporate to assist with this task, and will finish up with a short discussion about how good device configuration dovetails into good device security.

Key Steps in Managing an Environment Configuration

As businesses elevate through the growth cycle, they typically go through phases of network control. When the typical business starts, it usually hosts few employees, zero to few IT personnel, and a very small network presence. Because the network presence in these small businesses incorporates few devices, the management of those devices is relatively simple. Documentation and change control of these small networks is usually ad hoc. This is not because of laziness on the part of IT but has more to do with the priorities of the typical business startup.

As it grows, a business' internal processes become more refined. The business requires more network services to support its internal processes. As the number of network services mounts, network infrastructure improvements become necessary, and this natural increase in network infrastructure geometrically increases the number of network configurations. For nearly all businesses, priorities eventually shift away from the early ad hoc mode to the need for more stability and a slower change rate. It is this changeover from a request-driven mode to an architecture-driven mode that can cause stresses on burgeoning networks. Incorporating good configuration management at the right time is the critical success factor in preventing those stresses from negatively impacting the flow of the growing business.

Configuration management can take on many forms, and a number of frameworks such as the IT Information Library (ITIL) and FCAPS have been created by the industry to provide a tangible outline to work within. Distilling those frameworks into something useable by SMBs and those in the mid-market can be the most daunting part of incorporating formal change control. Let's make the process easy by breaking it down into four easy-to-understand steps. Similar to the graphic used in each of the previous chapters, let's use a variation of it to discuss the steps needed for implementing good configuration management as well as the tasks associated with those steps.

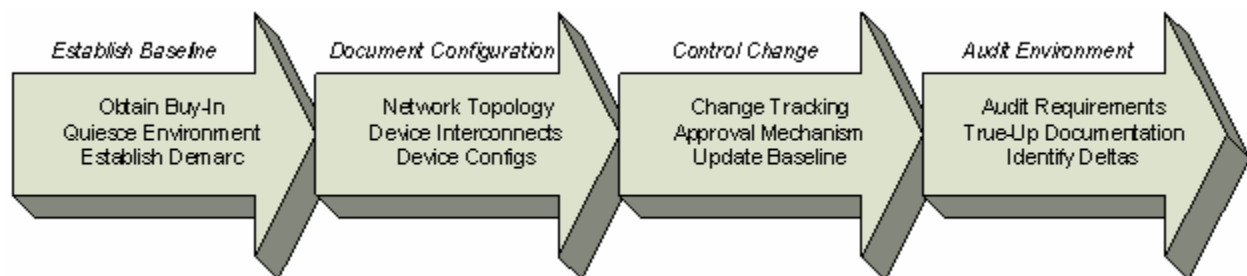


Figure 3.1: Network configuration management can be broken down into four easy-to-understand steps.


To take your network from one of ad hoc change control to one that incorporates good change management practices, you really only need these four basic steps. You'll first need to establish a configuration baseline, document that configuration, establish a mechanism for changing and updating the configuration, and put in place an auditing process to verify it. Let's talk about each of these steps in turn.

Establish a Baseline

Assuming your network is one that has not yet been baselined, the first step in implementing effective change control is to establish a baseline of your environment. This baseline is an understanding of the current configuration on the network and the interrelations between the devices that make up that configuration.

Even the smallest of networks involve a lot of change, so the first step in establishing a baseline is to quiesce, or quiet, the environment. This is not so much a technical step as one of an organizational “snapshot” of the environment past which all changes are logged until the documentation is completed and the change control processes established. This process of snapshotting the environment is a procedural one where a deadline date is established and all changes past that date are logged. This “line of demarcation” separates the previous ad hoc procedures from the new controlled ones.

Because creating the line of demarcation can have the potential of interrupting the normal operations of the network, it is imperative that business management buy-in is obtained prior to the activity. Their buy-in can be easily obtained by identifying the ROI associated with a healthily documented and controlled network. We’ll discuss some of the business metrics for establishing that ROI later in this chapter.

 Of the four steps in the configuration management process, establishing the baseline involves little to no technology solution. Where your NMS can provide documentation automation is in the next step, where you document the configuration.

Document Configuration

Step one in this process is almost completely a procedural step within the company and does not often involve the use of technology. Once you’ve obtained buy-in from management on the desire to move towards proactive network management, you can incorporate the features of your technology solution to assist with the documentation of that configuration.

Looking back to the initial question of this chapter, network devices typically have hundreds of lines of code that define their configuration. This text-based method for configuring network devices makes them cumbersome to configure manually but makes them excellent for automatic configuration storage and archiving. If each of the network devices on your network is configured with little more than a text file, and that text file can be transferred over the network, your NMS should easily have the ability to store the configuration into a central database. Doing so should be the largest component of your configuration documentation.

A good NMS should also be able to inventory the network for individual devices and their interconnections. This inventory should drive the generation of a map of those devices and their connections. The combination of each device’s configuration file along with the map of its interconnections should fully document your network configuration. It is reasonable to assume that with a good NMS, each of these processes should be at least partially automated. An excellent NMS will incorporate full automation into this process, making this documentation step very easy.

Although historically much of this configuration capture has occurred through automated tools that utilize Telnet or SSH to log into the device, newer technologies now allow for full SNMP-based management of all devices. Moving from command-line technologies to SNMP allows for all facets of device management to be done through a single protocol, making administration much easier.

Control Change

The third step in managing the configuration of your environment is the establishment of a process to formally request, approve, and document changes to your configuration. This process of change control gives you a formal mechanism to ensure that changes are done correctly, that others in the IT organization and the business are notified of the changes, and that the changes get correctly reflected in your baseline. Whatever organization process you choose should ensure that the IT organization and/or business leaders of your environment are informed of the need for change and have a reasonable ability to either approve or reject that change.

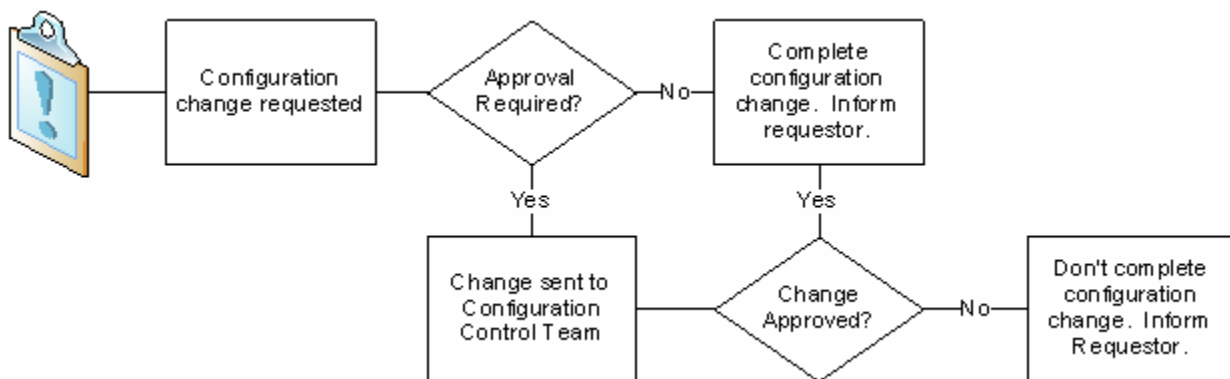


Figure 3.2: A simplistic formalized change control procedure.

From the perspective of your NMS, this change process should integrate with the configuration storage and archival process as well. Once you've stored your device configuration into your NMS, you need a process to update it after the change occurs. Some excellent NMS systems allow for the change to be made within the NMS and "pushed" out to the device through an update process. This feature allows the NMS to confirm the change and relate it to the configuration for other devices on the network and helps to reduce the chance that that change will negatively impact other devices on the network.

If the change involves an environment topology change, it must also integrate with the interconnection mapping capability of the NMS. This can be either through a re-scan of the devices on the network or an integrated update of the topology within the NMS prior to deployment.


Either solution ensures that change within the network is reviewed by others within the business, goes through an approval process, and is captured by the configuration management tool to ensure a one-to-one mapping between the actual configuration and the documented one.

Audit Environment

Lastly, any process for configuration management needs to include a process whereby that environment can be audited against its baseline. Whether network changes are done within the NMS automatically or done outside and synchronized with the NMS database, there are times when unapproved or inappropriate changes make their way into the network. It is the process of network auditing that validates your network's configuration and ensures that nothing inappropriate or unapproved has been done.

The auditing process is essentially the very same as the automatic documenting process discussed in step two, with the exception that your NMS should notify you when a network configuration doesn't match what is expected. This can either be done automatically and at regularly scheduled intervals, or it can be done as a manual or partially manual activity on a semi-regular basis.

Depending on the industry in which you do business, there may be one or more compliance regulations that require this auditing step to occur. Your ability to show successful compliance to network security regulations and prove your configuration can prevent you from expensive and damaging litigation.

 Later, this chapter will discuss more about compliance and compliance regulations.

Compliance Regulation	Industry
Sarbanes-Oxley Act (SOX)	Publicly-traded institutions
Gramm-Leach-Bliley Act (GLBA)	Financial institutions and those that handle personal financial information
Payment Card Industry Data Security Standard (PCI or PCI DSS)	Institutions that accept payment cards
Health Insurance Portability and Accountability Act (HIPAA)	Medical institutions

Table 3.1: Some compliance regulations that may have auditing requirements and the industries that must follow those regulations.

Ad-Hoc/Manual Configuration vs. Managed Configuration

Throughout, this guide has talked about how you can leverage a good NMS to move your style of network administration from reactive to proactive. Nowhere can you get more “bang for the buck” than in using your NMS for configuration management. This ability to standardize configurations, back them up in case of disaster, automatically restore them should they get corrupted, and automate many other network administration tasks eliminates much of the difficult and manual parts of being a network administrator. Removing these manual components allows administrators to focus time on more productive and value-added activities—such as performance management and reducing downtime.

Incorporating an effective NMS into your network will grant you these abilities. An effective NMS will have some, if not all, of the following feature sets to assist with this automation activity.

Configuration Standardization

With an average of 300 lines of code needed to configure a typical network device, it is reasonable to assume that over time small differences between devices can manifest. Ever tried to line up two files and analyze them line-by-line to verify they are identical? The process is painstaking and fraught with error. Standardizing on a framework for device configuration helps to reduce that error. Using an NMS to provide the framework and to automatically notify you when deltas occur between devices goes even further.

```
!
version 12.1
no service pad
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname switch1
!
!
ip subnet-zero
!
vtp domain [smartports]
vtp mode transparent
!
spanning-tree mode pvst
spanning-tree extend system-id
!
vlan 2
 name VLAN_2
!
vlan 3
 name VLAN_3
```

Listing 3.1: Network devices typically have very similar configurations; the differences appear within individual port network assignments. Standardizing on configurations reduces their complexity.

A good NMS will provide a suite of standardized device configuration templates that incorporate best practices for performance, readability, and security. By providing these templates, you need only “fill out the form” with the individual configuration. A good NMS will also provide the ability to do a side-by-side and line-by-line comparison between two devices to validate their similarities and differences. Standardizing on a single template for each network device type ensures that multiple administrators can manage the network with a minimum of administrator “personality” in coding and configuration style.

Configuration Backup and Archival

A configuration is only useful if it’s loaded onto the device and performing its function. This always holds true in a device’s running configuration but never holds true if the firmware of the device wipes clean. In those critical situations, it is often difficult to rebuild the configuration by hand. The stress of needy users and irritated business leaders can result in mistakes being made when they are needed the least.




Figure 3.3: An NMS enables automated configuration backup, archival, and retrieval should a configuration become corrupted or problematic.

A good NMS will provide the ability to back up each device’s configuration files over the network and store those configurations into a centralized database. This database should have the ability to show and store multiple configurations over periods of time. This allows the administrator to rollback to a previous configuration if the situation warrants. The archival of configurations is additionally necessary to ensure that long-term backup of device configuration can be done should historical research be necessary by either administrators or outside auditors in the case of a compliance audit.

Post-Incident Restoration

Backing up your configurations is handy only if you have a mechanism for easily and quickly restoring them onto the correct device or its replacement. Once an event has occurred and the initial triage determining the source of the problem is complete, the best course of action in many critical down events is to “turn back the clock” and revert the device to its last-saved configuration.

In many cases—especially when good configuration control is in place—that last-saved configuration is equal to the configuration that was on the device prior to the outage. Restoration of the configuration onto the device will return it to service.

In some cases, however, the device may have been in a maintenance period where changes were being implemented or functionality testing was being performed. In those cases, either an administrative misconfiguration or a configuration mismatch or corruption caused the outage. Reverting the device to its last-saved configuration will bring the device back to a state where its functionality is known. The speed and ease of post-incident restoration is a critical determinant in choosing a good NMS for configuration management.

Policy-Based Configuration

As the number of devices on your network increases, you will begin to find that many devices are configured nearly the same across the network. It is this similarity in configuration that allows for the device configuration framework to be guided by corporate, network, and security policies.

The process to create a policy involves business and IT working together along with necessary compliance and security regulations. Often, an NMS can provide a known best practice as a starting point for creating device policy, such as one that automatically complies with SOX, HIPPA, or other regulations. This device policy becomes the framework in which all devices are ultimately configured.

By configuring devices according to an agreed-upon policy framework, a network now has the ability to quickly update all device configurations should that policy change. An effective NMS will offer the functionality to provide best practices in device policies and policy frameworks, and most importantly, incorporate the ability to deploy changes to that framework across all devices through an automated manner.

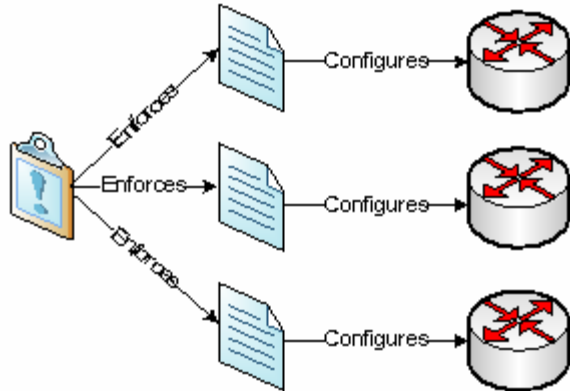


Figure 3.4: With policy-based configuration, a single corporate policy can define and enforce multiple device configurations.

As an example, presume that for performance reasons your business wants to prevent traffic associated with media streaming from traversing your network. A policy-based configuration will allow that denial to be enabled on all routing and/or switching devices in the environment. Should a change to corporate policy be later incorporated allowing streaming media traffic, that configuration can easily be updated on all devices vis-à-vis the policy. A configuration based on a policy framework means that those exclusions can be rolled out to the pertinent devices automatically with a minimum of manual intervention.

Inventory and Mapping

SNMP is a great protocol for identifying devices on the network. Combining SNMP with the routing and switching protocols that enable network devices to identify their closest neighbor enables a smart NMS to quickly build inventory and mapping data. Among other items, network inventory shows the administrator devices enabled on the network, their make and model, firmware version, network location, and connected hosts. As discussed earlier, the mapping component of your NMS should enable a graphical representation of the interconnections between devices on the network.

Rogue Device Identification and Adjudication

As a security function, your toolset should also allow for the capability to scan your internal network for known and rogue devices. Having the knowledge of the known devices on your network is good only for verifying their availability. However, being able to know when rogue or inappropriate devices appear on the network is critical for network security. Your NMS should include the ability to scan network ranges to look for devices that both should and should not be there as well as provide some limited information about the rogue devices and their configuration.

In some cases, the introduction of a rogue device onto the network is a harmful activity that should be immediately auctioned upon. Depending on the severity of the business security policy, a good NMS can be configured to either notify when that rogue device comes on the network or even go as far as to shut down the network port of its attached neighbors, isolating the rogue device from the rest of the network.

Provisioning

The ability to rapidly provision new devices onto the network improves the efficiency of the network administrator. Though small networks may not often add new devices into the network, there tends to be a related increase in device count as the number of new employees increases. Dovetailing with policy-based and standardized configuration capabilities, your NMS should have the ability to “cookie-cutter” those configurations while preserving individual device uniqueness.

This combination of standardized configuration along with device personalization ensures that each device on the network can be provisioned automatically. Just hook up the device to the network, enable its SNMP community strings, and instruct the NMS to download the correct configuration and begin regular configuration backup, monitoring, alerting, and auditing functions.

Deprovisioning

Although the provisioning activity is usually well done by good NMSs, where many fall short is in the ability to quickly deprovision them. The process of deprovisioning involves removing the configuration from the device and removing the device from the network. It also involves decisions associated with maintaining that device’s configuration in the NMS database and/or removal.

Where critical features exist are in device upgrade situations. When a device is upgraded either with a brand-new replacement or with an internal improvement or augmentation, a provisioning activity and a deprovisioning activity are rolled together. Your NMS should incorporate the ability to roll the configuration from the previous device to the new device while still maintaining the uniqueness characteristics of the new device. In the case of a device augmentation, the NMS should have the ability to roll the changes to the device IOS into its configuration.

User Access

Lastly, a granular and auditable user access policy should be a feature set on your NMS. Compliance regulations specify that users and user logons are to be segregated in such a way that users log on with individual accounts. They also specify that activity should be tracked into a database that cannot be manipulated by users other than the top-level administrator. This segregation of roles prevents collusion between administrators and reduces the chance that a single administrator can take down the entire network through his or her actions alone.



Through all the aforementioned features, you’ll notice a centralization of control to the NMS. This centralization of control has the effect of “putting all your eggs into one basket.” Thus, if not properly secured, it could be hacked by an aggravated administrator with a desire to damage the network. Every network device can be managed and controlled by the NMS, so the administrator’s malicious intent could indeed do major damage to each network device.

An effective NMS will have the ability to segregate administrators into roles and assign tasks to those roles. The lead administrator that is ultimately in charge of the entire system can assign users to roles and tasks, preventing any one user from gaining too much access into the system. Not having this distribution of work share can often cause problems on a compliance audit.

Business Drivers for Configuration Management

Moving your network administration style from reactive to proactive provides a lot of benefit to the business as well. It allows the administrator to better understand changes in the network prior to implementing them. And reducing firefighting often frees up enough time that an otherwise harried administrator can spend more time focusing on ways to prevent problems from ever happening.

Some of the business metrics that can be improved by the movement from reactive to proactive are described in the following sections. Each directly benefits by the changeover to formalized and automated network configuration management.



Metrics that define an IT person's worth are often driven by Number of Closed Tickets or Number of Resolved Problems. Unfortunately, driving an administrator's measurement of success by the number of problems fixed is kind of like paying a programmer by the number of lines of code they can write. A lazy programmer incentivized in this way will spend less time writing efficient algorithms and more time pushing out unoptimized code. An unscrupulous one may even artificially inflate the algorithm line count if they're behind on the house payment. Finding efficient and correct incentives for network administrators is just as critical.

MTTR Reduction

Chapter 1 discusses MTTR, which is a metric associated with the Mean-Time To Restore a particular system. If a company's MTTR requirements are measured in minutes, an administrator has only a very small amount of time to get a failed system up and operational. If network device configurations are backed up and can be quickly restored to failed devices or replacement devices, this significantly reduces the MTTR during a failure event.

Your MTTR requirement should be driven by your business needs. A short MTTR will drive the change to a tightly controlled environment.

Loss of Business Revenue

The MTTR conversation links directly to loss of business revenue. When a network device goes down during the business day, customers are unable to do business with the company and employee productivity reduces to zero. That reduction in productivity has a direct bearing on the company's bottom line.

Implementation of an NMS that can automatically triage a down device, notify an administrator, and provide suggestions for remediation all increase the speed of recovery and reduce the accompanied loss of revenue associated with an outage event.

Security

Security is a critical component of all networks. However, in most networks, the majority of the security exists at the perimeter. Inside the LAN, security controls can be lax. This has the tendency of creating a security profile with a solid outer shell and a soft inside. This typically happens in networks because of the complexity of enabling and implementing centralized security controls, access control, and a cohesive security policy.

With the incorporation of policy-based configurations on network devices, the security profile for a network can be enhanced. This is especially so when the policies that drive each device configuration are generated from known security best practices. When considering the ROI associated with an NMS purchase, consider one that provides such best practices right out of the box to augment the security of your soft internal network.

Regulation and Compliance

Depending on the industry in which you do business, there may be governmental or other regulations that drive a certain security and auditing stance within your organization. For most of those regulations, the concepts of role separation, policy enforcement and follow-up auditing, and securing of data during storage and transport are of great importance.

For an example of how that can impact network operations, consider the PCI DSS standards used by vendors or service providers that handle, transmit, store, or process information using payment cards. These standards require a predetermined level of internal security for any networks that transmit personal financial information. Depending on the number of payment card transactions your business does in a year, that level of security increases. Your business can incur costly and damaging litigation should an information disclosure event on your network result in the loss of that personal financial information to an outside party. By ensuring a stable and auditable configuration and security profile, you fulfill the PCI DSS along with other compliance regulations while reducing your liability.

PCI DSS High-Level Requirement	Affected Through Effective Network Device Change Management Leveraging a Good NMS
Install and maintain a firewall configuration to protect cardholder data	Yes
Do not use vendor-supplied defaults for system passwords and other security parameters	Yes, for network devices
Protect stored cardholder data	Yes, upon data transmission
Encrypt transmission of cardholder data across open, public networks	Yes
Use and regularly update antivirus software	Yes, on device IOS
Develop and maintain secure systems and applications	Yes, for network devices
Restrict access to cardholder data by business need-to-know	Yes, through process
Assign a unique ID to each person with computer access	Yes, for network devices
Restrict physical access to cardholder data	No
Track and monitor all access to network resources and cardholder data	Yes, for network devices
Regularly test security systems and processes	Partially, through auditing capabilities
Maintain a policy that addresses information security	Partially, through process

Table 3.2: The PCI DSS security regulation and how effective network device change management leveraging a good NMS can affect compliance with that regulation.

Auditing Requirements

Hand-in-hand with compliance regulations is the requirement to validate your configuration. Any change process that incorporates manual steps can introduce error into the result due to data entry errors, missed steps, or incorrect documentation. Because of this, a regular and out-of-band auditing activity needs to occur on your network. Depending on the features associated with your NMS, that activity can either be highly manual or highly automated.

With highly manual activities, additional error is introduced as the administrator reviews the logs and configurations and completes the true-up between the documented and the actual configuration. Additionally, the cost in actual dollars and administrator time to complete the activity can be prohibitive or even prevent it from actually occurring.

With a highly automated activity, the process of comparing the actual and archived configurations can be done with very little effort by the administrator. Thus, what may have normally been an annual activity can be configured to run daily or hourly. This enhances the network's security profile and ensures that inappropriate or unapproved changes are immediately identified and adjudicated.

Personnel Turnover

In terms of loss associated with inadequate security, it is well known that the people inside the company are the biggest risk. Personnel turnover, especially in cases of a terminated employee, have the chance of interrupting network operations should that employee leave behind backdoors into the network or cause damage prior to departure.

In an environment that incorporates centralization of user accounts and passwords along with role-based security attached to each username, it is less likely that a terminated employee can cause this sort of damage. When all network devices look to a centralized location for their identification, authentication, and access privileges, the elimination of privileges can be done very easily and from a single point rather than requiring an emergency password change on all network devices. This has the effect of eliminating or reducing the problem associated with employee turnover.

Supporting Technologies

This chapter has discussed several feature sets that make up a good NMS. There are some additional supporting technologies that may be incorporated either within your NMS or through associated tool sets that help enable those features. Three that are discussed in this section are configuration analysis and comparison tools, task scheduling tools, and RADIUS/TACACS.

Configuration Analysis and Comparison Tools

Network devices are notoriously command-line driven due to their highly optimized architectures and a long-standing tradition towards text-based configuration. Doing configuration via command line is exceptional for experts and for doing batch manipulation of many devices, but it can be very cumbersome for non-experts or cross-network updates. In those cases, a graphical interface that retains the text-based nature of the text file yet enhances it with additional features can be a key tool.

These graphical “skins” for configuration files make an excellent addition to either your NMS or your network administrator’s tool set. Some features that these graphical tools provide are:

- Color-coding to differentiate configuration sections
- Side-by-side comparisons, also with relevant color-coding
- Support for multiple device vendors
- Real-time configuration change detection and alerting
- Bulk-update capability
- Management and administrator reporting
- Itemized configuration change history
- Device grouping
- Individual port status

Task Scheduling Tools

Many of the tools discussed are handy but only so when the administrator is sitting in front of his management workstation. Adding to these tools the capability to schedule reports, device backups and restoration, and configuration updates mean that administrators can move activities to known times in the future. The ability to configure when these tasks occur allow for regular activities to take place without needing to remind the administrator to accomplish the task. Consider the ability to schedule tasks a necessary tool for your tool set.

RADIUS/TACACS

Much of this discussion regarding security lies on the centralization of identification and authorization information away from the individual network device. Most network devices have the capability of setting up access based on username on the device. This works well when the number of devices is small and the users using those devices trust each other. However, in larger organizations with many devices and many administrators, the need for centralization of account authority grows with device count.

Remote Authentication Dial-In User Service (RADIUS) and Terminal Access Controller Access-Control System (TACACS) are two technologies that enable this centralization of access. Rather than creating “enable” passwords on each device, RADIUS/TACACS allows users to use individual accounts for login and enable rights.

As stated before, many compliance regulations require unique usernames and passwords for all administrators for tracking purposes. Thus, using the default per-device authentication mechanism may violate certain compliance regulations. Most network devices have the capability to authenticate against one of these external authentication databases.

```

!
aaa new-model
aaa authentication login default group RADIUS local
aaa authentication login CONSOLE local
username root privilege 15 secret MyP@ssword
!
no enable password
no enable secret
!
!
ip radius source-interface Fa0/0
!
radius-server host 10.0.0.1 auth-port 1645 acct-port 1646 key P@ssword2
!
Line console 0
  logging synchronous
  login authentication CONSOLE
!
line vty 0 15
  Privilege level 15
  login authentication default
!

```

Listing 3.2: An example of a device connection to a RADIUS server for centralized administration.

Understanding Security Management

Moving the conversation away from the C in FCAPS and focusing it on the S for a moment, let's briefly explore security management. Security management is a critical part of any network, and doing it correctly makes the difference between a hacked network and one that can survive external attack.

The focus of this discussion is on how your NMS can enable good security management in your network. An effective NMS will help you practice good security as well as notify when you when commonly known best practices are not being followed for the devices on your network. This section will focus on five places where your network architecture and your NMS can assist with creating a good security posture for your environment.

Practicing Good Network Security

Obviously, the most important facet of security is simply doing it correctly. A correct security posture for your network will drive fewer infiltration successes. Much of this starts with generating a cohesive security policy for your network. This is of critical importance because lacking a good security policy prevents enforcement of security procedures when situations do occur. That security policy should take into account factors such as:

- Password policies
- Acceptable use policies
- Lockdown and access policies
- Mobile device access and lockdown policies
- Business data encryption policies
- Antivirus, anti-spam, anti-malware, and anti-spyware policies
- Security policy violation adjudication procedures

The incorporation of a sound and cohesive security policy in your organization will drive the creation and enforcement of other policies and procedures that fulfill the need for security. From a very high level, the need for security should drive answers to these questions:

- Does network traffic route to its appropriate destination and nowhere else?
- Does only authorized traffic pass into and around the network?
- Does that traffic pass at a recognized and acceptable volume?
- Does data arrive intact without corruption and without interception by inappropriate internal or external entities?
- Does traffic that requires encryption traverse the network encrypted and with the correct level and strength of encryption?
- Are mechanisms in place to ensure malware and potentially unwanted software do not enter or traverse the network?

SNMP Community Strings and SNMP Weaknesses


As described earlier, SNMP is the major component of configuration enumeration and manipulation for most network devices. However, SNMP strings themselves have weaknesses that should be taken into account when deploying SNMP-based NMSs to manage them.

Unlike user accounts, which can be segregated to individual persons through the use of RADIUS/TACACS servers, SNMP community strings are often not unique for each device in the network. If you've taken the time to segregate your administrators and their logins, they may still have a backdoor into the network via the shared SNMP community string on each device. An SNMP community string is effectively a password into the system, especially when used to update device characteristics, so extra care must be incorporated to ensure their security.

One feature of a good NMS is the capability of rapidly changing SNMP community strings for all devices on the network during events such as administrator termination or suspected infiltration. SNMP community strings should be changed as part of a regular password change cycle. And, if possible, each device should use a different string. As each device can have as few as two strings for reading and writing data and potentially more, a manual update of these strings can be costly in terms of time. Automating this process ensures it is done regularly, further enhancing the security profile of a network.

Port Scanning and Port Minimization

Although port scanning from the Internet is usually considered a bad thing, port scanning by an approved administrator with good intentions in mind is a good thing. By completing port scans on internal hosts—and especially those connected to the Internet—the network administrator can ensure that only the necessary services are enabled and listening on a particular host.


 This is especially a problem on Microsoft Windows servers, which tend to listen on numerous ports, some of which are of very high risk in untrusted network environments.

Port scanning is a feature of many NMSs and NMS tool sets. The port scan can identify on which ports a server is listening and may also provide some information about which service is listening on that port. Once the host is scanned, the network administrator can work with the systems administrator to shut down any unnecessary services. In a case in which the service cannot be disabled, the network administrator can modify firewall rules to prevent external entities from accessing the high-risk service. This back-and-forth process encapsulates the idea of port minimization.

Penetration Testing

Taking the concept of port scanning even further is the idea of penetration testing. Penetration testing can be done either internally or by an external entity such as a security consulting organization.

No-cost and for-cost tools are available that simulate dozens or hundreds of known exploits against a particular host, whether that host is a network device or a server. Pointing these tools towards the hosts and network devices in the demilitarized zone that separates your internal network from the Internet is an excellent way to get a feeling for your network security posture. For network administrators that lack the experience in penetration testing and network device hacking, these tools can provide a report discussing where security vulnerabilities may exist.

 Many of these tools are excellent for identifying the technical security posture for your organization, but investing in a security assessment by an external organization may provide better results. Additionally, recognize that as technical security controls improve over time, one of the biggest vectors for infiltration today is through “social hacking.” With social hacking, an external entity bypasses technical controls by talking to employees and pretending to be a member of the organization. Any penetration testing should include a look into the feasibility of an external entity to use this method as well.

Vulnerabilities, Exploits, and Patches

Lastly, all computer devices have bugs and vulnerabilities, either in the code or within that system’s architecture. No matter how you architect your external security posture, a problem with the device itself that goes unpatched becomes a vulnerability that can be exploited. Only through constant and vigilant patching of those vulnerabilities can those holes in your security posture be fixed.

As a network administrator, you will want to ensure you are constantly apprised of known vulnerabilities in your network devices. Subscribe to newsletters, visit pertinent Web sites, and keep yourself aware of these problems as they occur. You also will want to identify a regular period of time that critical patching—and the associated downtime—can occur on your network.

Configuration and Security Management Provide Measurable Benefit

This chapter talked about many of the topics associated with configuration and security management in your network. It has also related how implementing those technical tasks can directly benefit the business. Depending on the timeframe of your business’ life cycle, the maturity of its network, and your desire to move from reactive to proactive, implementation of automation may be a critical success factor. If you and your business make the decision to become proactive in network administration, consider the factors discussed in this chapter when choosing an NMS that assists with this automation.

The final chapter will depart from proactive management and talk specifically about the steps and the tools you need to effectively troubleshoot network devices when they have problems. It will explore keys topics—such as IP address management, network engineering applications, and DNS—that will help you when the network shows a problem.

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>.