

Realtime
publishers

"Leading the Conversation"

The Definitive Guide[™] To

Converged Network Management



Ken Camp

Chapter 9: Effective Security Management	185
FCAPS and Security Management	186
Identifying Risks	187
Inconsequential Risk Classes	188
Significant Risk Classes	188
Introducing Incident Response Planning: CSIRT Basics	193
Building Processes for Defense in Depth	194
User Authentication and the Gold Standard	196
Authentication	196
Authorization	198
Auditing	198
Prepare to Respond	199
Policies, Procedures, and Awareness	199
Physical Access and Environment Controls	200
Perimeter: First Line of Defense Between the Internet and Internal Networks	200
Network: Internal Network Layer	201
Systems: Server and Client OS Hardening Practices	203
Application: Application Hardening Practices	204
Data: Protection of Customer and Private Information	204
Summarizing Defense in Depth	204
Prevention	205
Protect the Protection	206
Detection	206
Reaction	206
Proactive and Reactive Strategies	207
Reactive Strategies	209
Proactive Strategies	209
Testing the Strategy	211
Summary	211

Copyright Statement

© 2007 Realtimepublishers.com, Inc. All rights reserved. This site contains materials that have been created, developed, or commissioned by, and published with the permission of, Realtimepublishers.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 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.

[**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 9: Effective Security Management

The convergence of voice, video, and data networks has been evolving and gaining momentum for the past several years. Many organizations have undertaken VoIP implementation to converge networks for cost reduction. Others work to achieve the competitive advantage of integrated services. Whatever the reason for network service integration, you cannot overlook the security risks that arise as technologies converge. VoIP implementers often focus on issues of voice quality, and interoperability. These are truly important factors in the delivery of voice services. In many ways inside the converged service network, voice security needs to be treated as data security. And data security needs to be treated as voice security. Both technologies bring issues and management techniques that benefit the other. This chapter will highlight security management issues facing enterprise deployments today and identify common industry best practices for creating an effective and comprehensive security plan that balances securing the network against the VoIP requirements for availability, reliability, and performance.

Security methods can adversely impact network performance. Firewalls induce delay by inspecting each packet in the data stream. This will add delay to packet delivery. Congestion at the firewall can lead to variable processing time within the firewall. This will increase the problem of jitter. A systematic and holistic approach to managing integrated network performance and security includes working with vendors, services providers, and trusted business partners to ensure a comprehensive approach to security is followed.

As previous chapters illustrate, successful operations are driven by knowledge and information. The more you know about the network, the better you're able to analyze problems. Solid knowledge and understanding of the network leads to an approach that balances all aspects of network management. Building this base knowledge helps you effectively manage the entire life cycle of network services and applications to ensure you're delivering the services needed today and able to meet the needs of tomorrow.

FCAPS and Security Management

Recent reading online brought my attention to an article from Professor Eugene Spafford on the Center for Education and Research in Information Assurance and Security (CERIAS) at Purdue University. CERIAS is viewed as one of the world's leading centers for research and education in areas of information security that are crucial to the protection of critical computing and communication infrastructure. In regard to best practices, Professor Spafford said:

In the practice of security we have accumulated a number of “rules of thumb” that many people accept without careful consideration. Some of these get included in policies, and thus may get propagated to environments they were not meant to address. It is also the case that as technology changes, the underlying (and unstated) assumptions underlying these bits of conventional wisdom also change. The result is a stale policy that may no longer be effective...or possibly even dangerous.

Policies requiring regular password changes (e.g., monthly) are an example of exactly this form of infosec folk wisdom.

*From a high-level perspective, let me observe that one problem with any widespread change policy is that it fails to take into account the various threats and other defenses that may be in place. Policies should always be based on a sound understanding of risks, vulnerabilities, and defenses. **“Best practice” is intended as a default policy for those who don’t have the necessary data or training to do a reasonable risk assessment.***

In small and midsized businesses, and even large enterprises, often the resources aren't available to perform comprehensive or even reasonable risk assessment. Perhaps more importantly, we create deliberate change control processes that don't provide the luxury of time to reassess risk each time we make a change in the dynamic networked environment we support. To perform a reasonable risk assessment each time your environment changes would stall all forward progress in technological growth.

Industry best practices are a tool employed because they afford us the collective institutional knowledge of the entire technology sector in adopting processes that have been proven to work. Best practices leverage, to your benefit, the common body of knowledge that IT managers share based on experience.

The key to the FCAPS model is in establishing processes for managing each of the five service areas—fault management, configuration management, accounting management, performance management, and security management. This chapter will delve into security management processes to protect and sustain the converged video, voice, and data service network.

Identifying Risks

Best practices provide a tool to support operations, yet some basic risk assessment is always prudent. Risk is simply something that is outside normal operations. When voice, video, and data networks converge, every risk associated with each service is incorporated into the newly converged network. This additive risk compounds the necessity for building a layered defense to protect not just the network infrastructure but also each of the services running on the network. Figure 9.1 shows a simple risk model called Jacobson's Window that is often used in considering the range of risks to which networks are exposed.

Jacobson's Window present a simple, two-factor risk assessment model, distinguished by the frequency or probability of an event occurring coupled with the impact or consequences of the event. This two-factor matrix simplifies all risk into one of four categories.

		Consequences	
		LOW	HIGH
Rate of Occurrence	LOW	DISREGARD	
	HIGH		DISREGARD

Figure 9.1: The Jacobson's Window risk model.

The four risk classes identified using this approach are low-low, high-low, low-high, and high-high. This approach is further simplified by breaking risks into two broad classes: inconsequential or significant.

Inconsequential Risk Classes

One risk that is considered inconsequential is the low-low class. It can generally be ignored because it doesn't matter statistically. The likelihood of occurrence is low and the consequences of these risks are deemed low, so they represent minimal impact on the organization. A risk that occurs once a year and has the impact of costing \$1 is just too trivial to worry about.

The opposite end of the spectrum, high-high risks, can also be ignored for the most part. In his writing, Jacobson suggests these risks just don't ever occur in the real world. He uses the example of a 50-ton meteorite crashing into your computer room on a daily basis. Although this example is extreme, it demonstrates why the high-high risk class is inconsequential. If this event occurred, you wouldn't build your computer rooms in the first place. In the real world, the high-probability and high-loss risks are generally immediately mitigated or we simply couldn't conduct business. These risks demonstrate lessons that we have already learned and for which we've developed remediation strategies. A new high-high risk is typically spawned out of some disruptive, radical shift either in technology or business practice. They're typically addressed as they're created.

Significant Risk Classes

If you're looking for a simplified approach to risk management, you've just eliminated half the classes in the entire risk spectrum. That only leaves two categories or classes of risk you really need to be concerned with: high-low and low-high. Spam email represents a perfect example of a high-low risk. There is a very high likelihood of occurrence. Current studies indicate that more than 50% of the email received is spam. In most cases, the resulting loss is a loss of productivity as people delete the spam messages.

Jacobson uses a fire destroying a telephone company central office as an example of a low-high risk. It demonstrates an event that has low probability of occurrence but a high consequential loss.



On May 9, 1988, an electrical fault started a fire in the Hinsdale, IL central office operated by Illinois Bell. Early during the fire, telephone services failed. The fire department didn't arrive on site for 45 minutes. Because of the dense, black smoke, firefighters had difficulty entering the building and locating the source of the fire. Emergency power was automatically provided by generators and batteries and could not be shut off easily. Neither standard dry chemical nor halon extinguishers were effective against the fire. Water had to be used, which exposed the firefighters to electrical shock danger. It took firefighters more than 2 hours to shut down power, enabling them to control and extinguish the fire. It was more than 6 hours after the first fire crews had arrived on the scene that the fire was declared under control.

This fire was confined to an area roughly 30 feet by 40 feet on the ground floor. Cables were burned to various degrees and smoke residue covered most of the ground and parts of the first floor. The most severe damage away from the fire was caused, not by flames, but by corrosive gases in the smoke. These corrosives damaged the equipment that survived the fire. Although the existing equipment was cleaned and used to provide interim service, it was deemed unreliable. This equipment all had to be replaced over time after the fire.

There is a broad spectrum of real-world risk that ranges from low-high to high-low. Human nature drives another facet of low-occurrence risks. Jacobson postulates that “People tend to dismiss risks that they have not experienced themselves within the last 30 years.” Although 30 years may seem arbitrary, psychologists believe it’s accurate and related to the range of human life experience. We generally remember events that happened during our lifetime. There are plenty of examples of Jacobson’s “30 Year Law” at work, but they’re easily identified. How often have you heard senior managers or public officials say, “How could we expect anyone to anticipate this event? We’ve now taken appropriate measures to ensure that this will never happen again.” Human nature shows us that some lessons may only be learned through direct experience.

When providing security for network services, managers need to visualize every possible risk that might occur in the enterprise. You’re often called on to consider events you haven’t personally experienced. It’s often useful to pull together incident response teams and think well outside the box of routine business. Natural disasters, cyber security incidents, and unrelated networking problems, such as a cut fiber, all impact service delivery. Sometimes a good planning session involves discussion of “all hazards” risk. This approach, undertaken with a broader set of staff members and managers, can help bring other physical and misconfiguration incidents into clearer view as potential problems. The next step is to identify optimal mitigation strategies for each risk.

Some organizations lean in favor of ROI-based risk management. This approach is often met with limited success. The ROI-based approach can work well with events that are high-likelihood, low-consequence because managers believe the risk exists, and therefore, it should be addressed. Low-probability, high-consequence risks present a different problem. It’s difficult to quantify the likelihood of occurrence, making it difficult to achieve management concern that the risk is tangible.

Every approach to risk assessment brings benefits and each has drawbacks. ROI-based decisions aren’t appropriate for every situation. When implementing security measures, it’s often wise to take a step back and evaluate why the security measure is necessary in the first place. Generally, security measures are implemented due to one of four reasons:

- The value of a security measure is significant but the cost is trivial. For example, the lock on the front door of an office costs little yet provides the very first measure of security. If the door is left unlocked, the consequences may be quite high but the cost of implementing a policy to lock the door is quite small.
- In many cases, the potential reduction in future losses will more than offset the cost of security. In this case, the security measure has a demonstrable positive ROI. This is the justification widely used for many security measures related to high-low risks in business. Card key access controls to buildings, password change policies, and basic security training for employees are good examples.

- For many enterprises, legislation or regulatory requirements necessitate specific security measures. The Health Insurance Portability and Accountability Act (HIPAA), Graham-Leach-Bliley Act (GLBA), Sarbanes-Oxley Act (SOX), and compliance with ISO17799 standards are good examples of this justification for security measures.
- In many cases, security measures are implemented to address a low-high risk that has an unacceptable consequence for any one single occurrence. This single occurrence loss (SOL) might be something that exceeds the owner equity or net worth of the company. These measures represent protection against risks that bring devastating consequences.

Low-high risks often require senior management involvement in decision making. ROI analysis may not justify implementing protective measures. Sometimes good judgment is required to decide what's acceptable to feel "safe enough" for the needs of the business.

SOLs and Reducing the Risk

There are two tactics to reduce overall risk using this two-factor approach to risk assessment. Because they offer tangible risks with a heightened sense of urgency, we have a natural tendency to focus on risks that generate a large SOL. You can typically mitigate these risks by:

- Reducing your vulnerability—Enterprises have often thought in terms of disaster recovery—how you might restore operations after a major event. There has been a shift in recent years to think more in terms of business continuity, or continuity of operations. Robust enterprise business continuity and resumption plans reduce any adverse consequences by shifting some aspect of the consequences to a resilient business continuity environment.
- Spread the risk across the enterprise—Distributing work across multiple geographic locations or business units can help spread the impact of any event across Help desk, training, and other organizations. The risk to each business area might be smaller and more easily addressed than a large risk to the entire enterprise.
- Transfer the risk—One solution may be obtaining insurance. The insurance premiums will factor in the exposure against a deductible amount, so the impact might be lowered to a tolerable level. A \$50 million dollar risk reduced to a \$5 million dollar potential single loss through insurance may provide an acceptable alternative.

After the potential risks have been identified, each enterprise will need to make unique decisions relative to their own business needs. High-low risks may be analyzed using an ROI-based calculation. Low-high risks often need to be reviewed in a management strategy session. It's wise to use the team approach to estimate the SOL potential and gain a collective sense of the likelihood of occurrence. Management will make a decision and draw a line somewhere along the spectrum. Usually this line is drawn based on the largest acceptable SOL.

Risks in the enterprise tend to spread out in a graph like the one shown in Figure 9.2. Viewed in this context, a conscious decision can be made about each risk and each type of risk. Low-low risks might be deemed inconsequential and addressed in another venue. As the risk increases, some risks may be transferred. Higher risks that occur infrequently might be areas for insurance considerations. Toward the middle of the low-to-high consequence range may be risks that can be addressed through policy and training efforts.

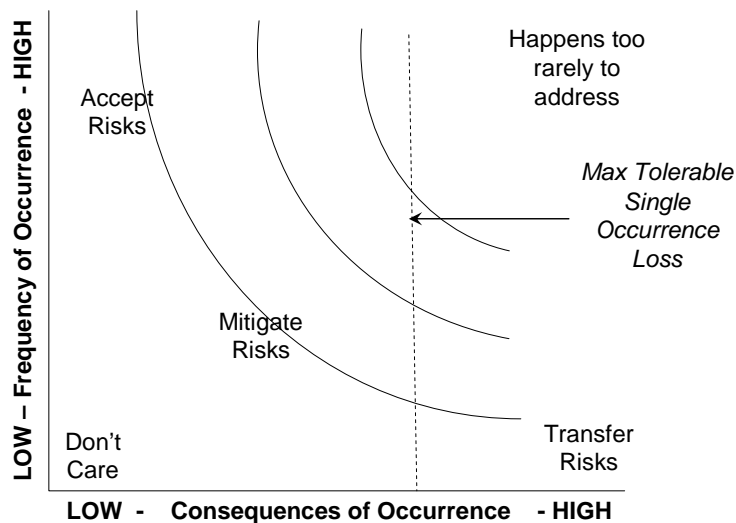


Figure 9.2: Modeling risks.

Although this four-quadrant view of risk provides a simple tool, for most businesses, the risk mitigation focus falls toward the middle area of the graph. The lower left quadrant represents things that happen so rarely and cost so little you might not need to factor them into daily operations. The upper left quadrant represents things that occur more frequently but cost little. These are often viewed as acceptable risks. They're typically treated as a factor of the cost of doing business. High-occurrence risks that rise in consequence show up in the lower right. These may be transferred to other parts of the enterprise, insurance carriers, and when appropriate, outsourced to a business partner. The upper right portion represents that area of high-high risk and consequence events that, at least in theory, have been addressed before you started doing business.

Jacobson's Window isn't always the best approach to risk assessment for an enterprise. It's a simple tool for quantifying and analyzing risk to help make sound business decisions that is easy to use and provides some method for basic risk assessment.

Risk Management Life Cycle

Like every process touched on throughout this guide, risk management is a never-ending process. Figure 9.3 shows the ongoing cycle of identifying, analyzing and assessing, mitigating, or transferring risks. Throughout the risk management life cycle, there are four key questions to consider:

- What could happen?
- If it does happen, what is the impact?
- How often could it happen?
- How accurate are the answers to the first three questions (measure of uncertainty)?

Risk management includes risk assessment and mitigation as two elements of a much broader set of activities. Other elements include establishing focal point for management, implementing controls and policies, promoting awareness, and continuously monitoring effectiveness.

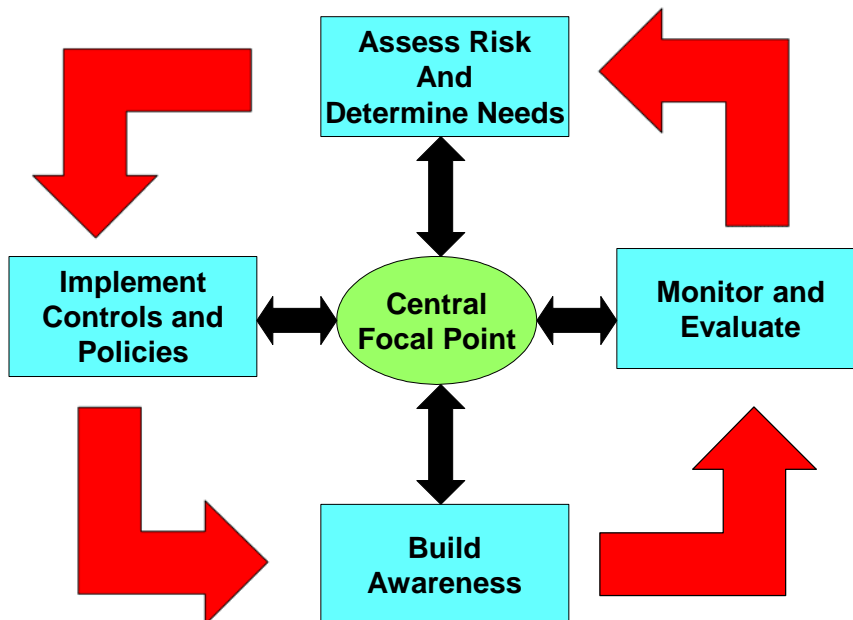


Figure 9.3: The risk management life cycle.

Risk assessment provides the foundation to all the other components of the risk management life cycle. It provides a basis for enterprise policies and helps maintain a focus on cost-effective approaches to implementing these policies. Risk assessment is an ongoing process that never ends. New services, video, voice, and data are constantly being incorporated into the service network. New vulnerabilities and exploits in operating systems (OSs) and applications are discovered every day. Reconnaissance attacks through firewall probing, port scanning, worms, and viruses are an ever-increasing nuisance. Threats and risks evolve over time and require vigilant monitoring.

Enterprise information security needs a management focus. This central focus point comes about through policy recommendations, business controls, awareness, and other areas that shape the overall corporate culture of stewardship toward business intelligence. Security awareness supports both daily operations and the risk management life cycle. Educating users on acceptable use and security risks is a key factor in mitigating security problems.

Introducing Incident Response Planning: CSIRT Basics

One vital component of how any enterprise responds to a security incident is the maturity and structure of the incident response team. In network security, there are three key management factors: detection, prevention, and response. Although detection and prevention are the ideals, you cannot detect and prevent every incident. Problems will arise. Security incidents will occur. Detection is crucial; it provides the situational awareness as to what is happening in the service network. Your reaction determines how you mitigate problems and restore business to normal operations.

Most large enterprises have some form of Cyber Incident Response Team (CIRT) or Computer Security Incident Response Team (CSIRT). For many companies, these working groups emerge from a *de facto* team that engages when an incident occurs. Organizations that don't think ahead and fail to build a team find themselves caught in the throes of *ad hoc* response. This causes delays in mitigation, often leaving critical services impaired for an extended period of time. A well-organized and trained incident response team can provide a range of services that help maintain service delivery capabilities. Typical roles include:

- Notification of the organization's management when an incident occurs
- Activation of a documented response plan to resolve the problem
- Record keeping of the steps taken to reach resolution
- Problem containment
- Evidence collection and gathering, including making backups
- Problem mitigation and service restoration to maintain continuity of operations
- Coordination of incident post-mortem reviews

These steps collectively gather lessons learned from every incident that occurs in the enterprise. These lessons learned feed process improvement plans, which leading to a continually rising ability to provide services with a minimum of disruption.

Many enterprises treat incident response as a value-added service within the organization, coupling prevention, detection, and response services to internal measure or service level agreements (SLAs). As Figure 9.4 shows, the CSIRT team provides incident response services.


 Some organizations are now shifting to a model of building a cross-divisional response teams that can provide a wider variety of services to multiple customer organizations inside the enterprise. Later this chapter will address some of these services in more depth.



Figure 9.4: Reactive, proactive, and quality management services in CSIRT.

Building Processes for Defense in Depth

A primary driver for a defense-in-depth, or layered, defense strategy is to create a holistic framework for end-to-end security across the enterprise. This comprehensive approach encompasses layers of security and multiple points of enforcement throughout the architecture. If the worst happens and there is a serious security breach, the organization must detect, contain, and correct the problem. Defense in depth not only reduces the likelihood of a successful attack; it significantly increases the chances of problem detection.

Confidentiality, integrity, and availability (CIA) of corporate data and network resources are paramount for any organization. This CIA acronym has become a watchword for the security industry. One goal of a layered defense is to build a service network infrastructure that can withstand attacks or disasters and protect the enterprise from unacceptable losses.

One approach that is often successful is to implement a series of routine technologies and procedures than embrace a complex and elaborate scheme that might contain single failure points. The simpler the infrastructure, the easier it is to protect. It's also easier to maintain. Service networks are large and complex enough without making them more so. And although simplicity isn't always a realistic goal in a large enterprise, it's important to build a service network infrastructure that is easily understood and maintained.

Any single control point can too easily become a single point of failure. Every layer of a strategy for defense in depth supports the other layers. When breaches occur at the perimeter, this layered approach limits the organization's exposure. Security through layers is often compared to an onion, as Figure 9.5 illustrates. There is no single layer that can be breached to get to the heart of the onion. Several layers must be either penetrated or stripped away. In network security, each layer provides another barrier, protecting the valuable business resources.

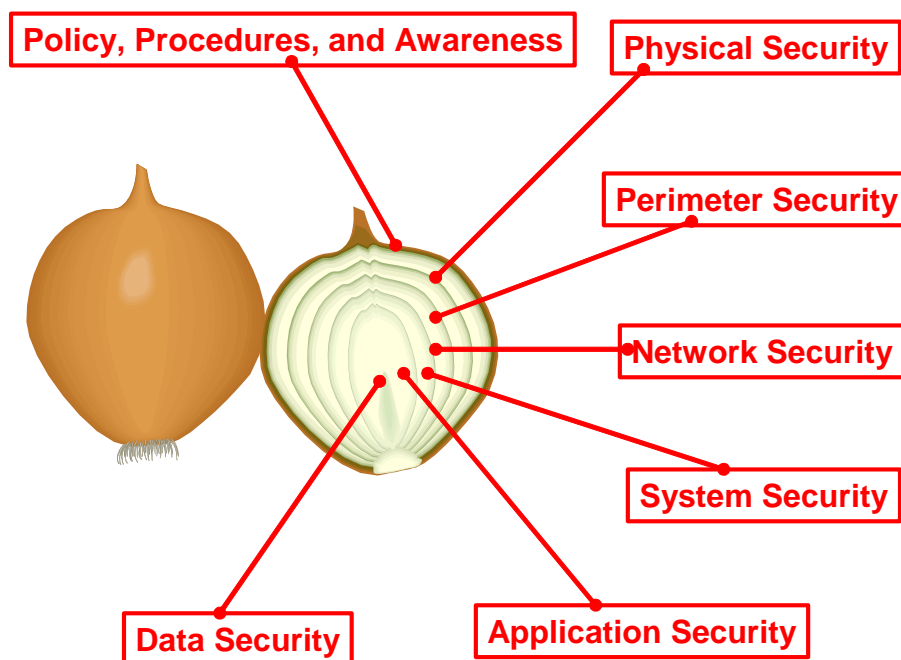


Figure 9.5: Layered security; a defense-in-depth strategy.

The best strategies provide effective protection of critical assets while allowing business to flourish. The layers or concentric rings of a layered defense allow for controlled and monitored access at each layer.

User Authentication and the Gold Standard

The primary focus of enterprise security efforts is to protect not only the services network infrastructure but also corporate data. Corporate data is a form of business intelligence, and as such, it's a vital asset to any business. Au is the symbol representing gold in the periodic table of elements. This reference to gold provides an easy mnemonic tool to help remember three of the most critical elements of network and data access security. Figure 9.6 shows these three “golden rules” in data access protection.

- **Authentication – Verify that users are who they say they are**
- **Authorization – Ensure users only access things they have permission to access**
- **Auditing – Know who did what, and when they did it**



Figure 9.6: The golden rules of data access protection.

Authentication

Authentication is simply the process of users proving that they are who they say they are. It's important to authenticate all users connecting to the network through some form of credential. User authentication may occur many times throughout the work day. The most common form of user authentication is simply an ID and password, but much stronger mechanisms exist and are in widespread use. There are some considerations to be mindful of with regard to the authentication process itself and how you deploy the technologies used:

- Do you employ a single authentication data store or are multiple resources involved?
- Are IDs and passwords encrypted during transmission or sent in the clear?
- Is simple user ID and password authentication sufficient or should you use some form of strong or two-factor authentication?
- How secure is the authentication and user validation process?

Authentication is simply the process of verification that users are who they claim to be. It's important that you authenticate users when they connect to the network and when they access business information that is restricted in any way. Many networks use simple password authentication technologies based on either Password Authentication Protocol (PAP) or Challenge Handshake Authentication Protocol (CHAP). These methods require only that the user know a user name and password. This is single-factor authentication.

Two-factor, or strong, authentication is typically based on a combination of “something you have” and “something you know.” One widely recognized example of this is the key fob technology that uses a time-based token to generate a number (something you have) coupled with a PIN (something you know).

Biometrics may be used to introduce a third authentication factor: “something you are.” This addition of a fingerprint/palm scan, facial recognition, voice recognition, or retinal scan adds another layer of depth to the authentication process. Three-factor authentication is common today online in the highest security environments and many research facilities. Although this technology works, the cost is generally still prohibitive to deploy in the typical business network.

Two-Factor Authentication Risk Analysis

Strong, or two-factor, authentication cannot be considered an inconsequential risk under Jacobson's Window analysis. Access control cannot be treated as a low-low inconsequential risk due to its high likelihood of occurrence. In a recent US-CERT Analysis Report, unauthorized access attempts rate in the highest category of network probes. This category encompasses 85.6% of attacks on networks in the latest quarterly evaluation.

Access control cannot be treated as a high-high inconsequential risk either. These exposures occur daily, constantly in the stream of questionable activity inbound to state networks. Unauthorized access attempts inbound to state government resources number in the millions each month. The rate of occurrence is high, but the success rate is very low. The initial consequence of breach to a user account may well be considered low.

Two-factor authentication concerns clearly fall into the category of significant risk. Significant risks are those that have either a high rate of occurrence or present a high consequence risk. Given millions of attempts per month against state government network resources, the rate of occurrence for login attempts through password dictionary attacks and other methods is clearly a high occurrence risk.

Whether the consequences of a breach are high is dependent on which system(s) are breached. The danger is that one breach easily leads to another, allowing deeper penetration into the network. These cascading attacks are among the most dangerous, as the attacker gains increasing ability to masquerade as an authorized user with each layer of penetration.

Single user passwords have long been proven as the weakest form of security in widespread use. In a recent survey, 70% of users will trade their password for a chocolate bar (see <http://www.securityfocus.com/columnists/245> and <http://www.securityfocus.com/columnists/245>).

There is research underway on many fronts to shift the technology from two-factor and three-factor authentication to a newer mindset of continuous authentication. This might be accomplished through a variety of heuristic information, such as keyboard typing and other behavioral patterns. As networking technologies have improved, so have security options, and many networks implement a variety of these technologies.

Authorization

After a user provides his or her credentials and has authenticated, the next step is authorization. This is simply a matter of matching the authenticated user against a stored profile and applying the appropriate access permissions for that individual. For some users, network access may be the only action permitted, perhaps to the entire network, but in some cases only granting access to a small subset. Different employees will typically have permission to work with different sets of data. This is commonly referred to as role-based authorization. (for example, Human Resources, sales, and engineering staff all need access to different resources but not to each another's resources).

Other authorization examples in a unified communications environment include:

- The ability to establish a remote VPN tunnel. This is generally predefined by user group and establishes another set of user controls based on rules in a VPN concentrator or firewall.
- The ability to program speed dial lists in the VoIP system may be restricted to only users with specific privileges.
- Conference call reservation and management is often a restricted service.
- Access to log and reporting servers is often employed to ensure that appropriate management or support staff are the only employees with access privileges.

Auditing


Auditing has become an important factor for enterprise business. In the traditional telephony provider, this focus was on accounting to support customer billing. As enterprise services grow and mature, auditing becomes a more necessary tool.

Auditing is often used in testing security configurations. It's also important for troubleshooting. Audit records will let you see whether the network permits multiple, simultaneous logins from the same user. Some organizations permit this by design, while others deny this capability. Multiple logins from different locations might provide an indicator that an employee's account has been breached or their password has been stolen.

A comprehensive audit trail provides a record of what every user did, and when. It can provide tremendous value in event correlation analysis when some incident occurs. Knowing when a VoIP call was placed, when a router configuration was changed, or when a server patch was applied can all provide crucial business intelligence about the service network.

Prepare to Respond

You and your network are a target, whether you think you are or not. Security incidents will occur regardless of whether you plan for them. The odds are high that when an incident does occur, it will be broader in scope and more deleterious to your network services than you've planned for. The time to prepare is now; and tomorrow; and the next day. When your network is under attack and services are failing, you will not have the luxury of thinking about a response plan. It's wise to take preparatory steps in anticipation of the inevitable.

 Treating the network like a doctor treats a patient, managing overall health care, provides a robust and sustainable network service delivery environment. Gathering documentation and consolidating knowledge about the unified communications service delivery environment is vital to understanding the healthy profile of the network. Strong, industry-adopted security practices, like good managed health care, require monitoring the health of the network on a continuous basis. Quick response to abnormal conditions will ensure early diagnosis and treatment of problems. Prevention, detection, and timely response are cornerstones to maintaining the good operating health of the service delivery network.

Policies, Procedures, and Awareness

The most secure system in the world is only as good as the people administering it and the users. People are the foundation for any good security model, and historically they are the weakest link. Policies and procedures provide the foundation for any layered defense strategy. People form the first layer of defense. A corporate culture of stewardship for protecting the company's proprietary information is fundamental to good network security.


It's not enough to put security policies in place. Employees need to understand them. Security policies should be reviewed with all employees annually. Every employee should be trained and understand their individual role in protecting network services and corporate information.

These processes and procedures are more important than security technologies. It's absolutely vital that you account for the human factor. People are the weakest link, and it's widely accepted across the information security community that the greatest threat to security lies within. People can all too easily make mistakes or become careless. A disgruntled, unhappy, or careless employee can be a threat to the security of the entire enterprise. Employee awareness and training are essential to building a sense of ownership into the corporate culture.

Physical Access and Environment Controls

The physical layer security of a services network seems easy to protect. Grant appropriate physical access only to systems providing services. Often, physical access grants administrative access via console ports and the like. Environmental controls, such as power and HVAC, should also implement controlled access mechanisms. Drastic changes in the operating environment, like an overheated server farm room, can result in damage and unexpected system behavior.

Physical access controls also mean controlling access to server farms, wiring closets, and other services equipment locations. You build layers of defense in the network. It's important to build layers of defense in your physical security as well. You use monitoring systems to monitor network performance and security. For many businesses, physical access monitoring may also be appropriate.

 There is a simple security review you can easily conduct that often yields some surprising results. Simply walk around your enterprise imagining that you're a burglar "casing the joint." Pay close attention to the server farm rooms and data center areas. Make sure you visit every room containing critical systems—mainframes, servers, voicemail systems, telephony systems, routers, and switches.

Check out the door hinges. Are they on the inside or the outside? Often doors are built to swing outward for specific fire or building code logistics. But a door that swings outward often leaves the hinges exposed on the outside. How easy would it be for an intruder to gain entry by removing the pins from the hinges on a locked door? Wiring closets housing equipment are often built with doors that swing outward because they're built in shallow recesses. They're particularly vulnerable to this problem. It's worth the time and effort to conduct a complete physical inventory and identify wiring closets that should have reinforced hinges and locks installed.

Perimeter: First Line of Defense Between the Internet and Internal Networks

The network perimeter used to be easy to identify. It was a single point of network ingress and egress, typically protected by a firewall. It's certainly true that the enterprise connection to the Internet is one perimeter point. But every extranet connection to a business partner is a perimeter point. Every VPN tunnel linking to another network is another perimeter point. *De-perimeterization* of the network has become a large issue for security managers to address.

Not only do you have external perimeter points but there are numerous internal perimeters within any large organization. Sales and marketing need to be cordoned off from research and development groups. Human Resources operations require special access controls to protect confidential personnel records. The network management resources need to be isolated from everyone except authorized users.

It's important to recognize that the goal of security perimeters isn't to keep everyone out. If that was the goal, you would simply disconnect the network and isolate it, eliminating all danger. Security mechanisms are implemented to support business need and allow appropriate access permissions to users who are trusted to do certain things in the network.

Enterprise networks today provide a layered defense system. Firewalls provide a hard outer shell, but without access permissions to demilitarized zones (DMZs) in the network and other resources as required, business would grind to a halt. The network services infrastructure and network operations center represent vital organs that sustain the health and continuous operation of the corporate organism.

The external hard perimeter is most often a firewall or series of firewalls. You implement firewalls to protect corporate assets from the malicious traffic on the Internet. They are vital to network safety but they also must allow business to take place. Traffic must get in and out of the network while you reduce the threats to an acceptable level. Some tactics used include:

- Using internal and external firewalls in combination
- Logging, analyzing, and reporting all access to the network
- Eliminating use of plaintext or unencrypted data that might expose internal systems to external threats—FTP and Telnet are cleartext protocols that can expose vital business information to anyone who might be capturing traffic; SIP is a widespread standard protocol in VoIP services today that is also a cleartext protocol that could potentially expose information about the enterprise network
- Encrypting incoming network traffic that is directed toward secure internal systems with VPN technologies, Secure Sockets Layer (SSL for Web services), and Secure Shell (SSH) instead of Telnet can provide an added measure of protection
- Proxy servers can provide traffic aggregation and funnel services through a single point to limit exposure; reverse proxy technology can hide all the inner workings of the corporate network
- Checking the identity of everyone accessing the network with strong user authentication as described earlier will help ensure that only known users gain access to corporate resources
- DMZs and filtered networks can be used to constrain external traffic to specified areas or segments of the network
- Routine penetration testing and vulnerability assessments can help identify weaknesses in the enterprise security posture

Network: Internal Network Layer

Layered defense can map conveniently to the OSI Reference Model in many ways. At the network layer are tools you can bring to bear to deliver security. You can use network monitoring tools as well as segmentation in the network design to help provide another layer of security.

Intrusion detection systems (IDSs) can provide alerts and warnings to assist the security team in quickly responding to incidents as they're detected. The faster a problem or security incident can be identified, the more quickly it can be mitigated. Quick problem resolution helps keep security incidents from cascading into a larger chain of events. In short, early detection improves reaction time.

Intrusion prevention systems (IPSs) go a step further. An IPS may be configured so that it can reconfigure other network elements in an automated response to some trigger event. In early 2004, the Bagle and Netsky worms spread widely and caused major network problems for many companies. Today, years later, both Bagle and Netsky variants are still being distributed daily. For most companies, some antivirus engine, either running on the mail gateway server or individual workstations frequently mitigates this problem. But Bagle and Netsky are both easily recognized by their digital signature. Intrusion detection devices routinely identify this malicious traffic. These worms have been around for so long that they are easily recognized. An IPS might not simply discard the packets at network ingress. The IPS might reconfigure border routers to shun all traffic from the sender. This might be a rule that is implemented automatically with a pre-set period of time. It might also be configured to make a permanent change to reduce risk and eliminate the CPU processing and bandwidth consumption that follows from allowing these malicious packets onto the enterprise network.

Traffic-shaping tools are used in large networks to provide trending information about network traffic. For example, assume that your services network traffic is normally made up of 30% Web traffic, 40% VoIP/video traffic, 20% email traffic, and 10% other mixed traffic. Traffic-shaping tools can provide indicators of pattern changes in the traffic. Although patterns evolve and change over time, making this useful planning information, a sudden or unexpected change in traffic patterns is often indicative of a network security problem. A new worm propagating via the corporate email system might cause email traffic to spike to 40% of all network traffic or more. Monitoring traffic patterns can provide an early indicator of a problem of this sort.

Access control lists (ACLs) can help block traffic and allow only specific IP addresses to communicate across network segments and subnets. Although they are basic and simple, ACLs can help protect critical VoIP systems by only allowing access to authorized system administrators from predefined network management workstations.

Internet Engineering Task Force (IETF) RFC1918 (<http://rfc.net/rfc1918.html>) defines non-routable addresses. A well-developed network strategy for deploying RFC1918 addresses can increase the services network security and further segregate traffic. In the VoIP environment, a design consideration might include using only RFC1918 addresses for VoIP phones and constraining voice service traffic to VLANs and MPLS domains dedicated to voice. Simply not allowing voice traffic on data VLANs and subnets provides a degree of granularity that offers some level of service protection by preventing intermingling of voice and data traffic.

Implementing a network-based approach to antivirus technologies may help eradicate viruses and worms before they spread to affect other layers. One layered approach to antivirus solutions might be to deploy one vendor's solution in the network, at the email gateway, and another vendor's solution on the desktop workstations. This two-fold approach might also circumvent any single vendor being slower than another to release updated virus signature files.


Conducting routine vulnerability assessment and testing of network services helps provide insight from the attackers' point of view. During normal operations, network personnel always strive to follow best practices but the demands of business and heavy workloads can lead to potential errors. In some cases, a solution is implemented quickly just to "get the job done." Regularly scheduled assessments can help ensure that any suboptimal techniques used to meet business needs today don't become the services network security vulnerabilities of tomorrow.

Systems: Server and Client OS Hardening Practices

Chapter 8 looked at configuration management, but server and client OS hardening goes beyond the network routers and switches to include all systems. You need to build processes for standardized installation and configuration of OSs. Services that aren't specifically required should be disabled on servers and workstations. Programs that aren't needed should be removed. All configuration changes to enterprise systems should be traced through audit logs. Change management procedures are necessary to ensure that only authorized changes are made to production systems.

Establishing a standardized set of configurations, across all platforms, is a widely accepted practice. Software programs that have been approved and accepted for use within the enterprise should be installed in standardized configurations. Some divisions or workgroups within any company may have specific software requirements, but only tested, approved, and authorized software should be installed. In many organizations, users are not allowed administrative rights to their local desktop machine. This approach can prevent the installation of unknown software and help enforce the use of standard enterprise configurations.

Strong user authentication and password technologies need to be implemented. When a system is prepared to move into network services production, special precautions should be taken to eliminate all guest accounts and vendor default accounts and passwords. Keep in mind that default password information is easily obtained over the Internet directly from vendor documentation.

 When setting up new systems in the pre-production environment, it is vital that all OS patches be applied and the system scanned prior to connecting it to the production network. Systems built from a 6-month-old CD will be built from an image that isn't at current OS patch levels. It's a sure bet that these older patch versions have vulnerabilities that are known. It's highly likely that there are exploits "in the wild" to take advantage of those weaknesses. These systems, once connected to the network, can be infected in minutes; well before OS updates from the vendor can be applied. Several industry studies have shown that an unpatched machine is typically compromised within 4 to 6 minutes of being connected to the Internet. A rigorous process of building, patching, scanning, and then revalidating patch levels will help ensure that new systems are not infected during the process of being brought online in production.

Application: Application Hardening Practices

The applications running on hosts often have control of how corporate data is handled. It's important to see that corporate applications don't introduce security vulnerabilities. A poorly written or malfunctioning application might allow for unauthorized access or undesirable manipulation of data. Secure coding techniques and stringent quality control processes are vital for any in-house application development.

You implement change control mechanisms in the IP network. You must take the same steps with your business applications. If a vendor partner releases a software patch for a VoIP gateway in the services network, diligent testing and documentation should be completed as the patch is applied in a methodical fashion.

Just as network hardware has a life cycle, so do business applications. Part of the application life cycle should include code auditing and peer review of the programming. Doing so will help ensure the integrity of enterprise applications. Authentication, authorization, and audit capabilities at the application level should be a consideration in all enterprise software, whether purchased commercially or developed in-house.

Data: Protection of Customer and Private Information

The single most valuable asset within the services network is the data. This data provides business intelligence information about the enterprise. It provides transaction records and history with customers. It provided details about products, services, employees, and finances. Every layer of a strong defense-in-depth strategy is implemented with one primary goal—to protect the data.

One common approach in businesses is to follow traditional military strategies and classify data in terms of the intended audience for distribution. When customer data is involved, privacy laws and regulations to protect personally identifying information will often dictate the protective measures your company must take.

 Chapter 10 will delve more deeply into regulatory compliance issues.

Summarizing Defense in Depth

The basics of a defense-in-depth strategy needn't be difficult to remember. Policies, procedures, and awareness provide the foundation of a secure network. Physical security ensures that only authorized personnel can touch systems. Perimeter security segments traffic to protect the boundaries between untrusted, semi-trusted, and trusted areas of the network. Network security adds a layer of intrusion detection and prevention solutions and provides traffic trend analysis to identify changes in patterns. Systems security ensures standardized configurations that adhere to company policies. Application security ensures best practices in coding and configuration of applications are followed. Data protection safeguards vital business intelligence information.

Prevention

Networks are vulnerable to a number of known problems. Perhaps the most feared, classic problem is a denial of service (DoS) attack. When under a DoS, the network is reduced to a state in which it can no longer carry legitimate users' traffic. Attacking the routers and flooding the network with extraneous traffic have been the historical approach malicious attackers have used to accomplish this. As malware has increased, today a worm attempting to replicate across the corporate network may present the greatest DoS problem as it consumes all available network bandwidth. Another classic problem is IP address spoofing. To protect the network against increasingly sophisticated attacks, you need to build stronger layers of defense. Just as layers of clothing are the best protection from the cold, layers of defense eliminate single points of failure and provide broad network security.

Digital Common Sense

In an article published in Computer World in 2002, Thornton May said "There can be little argument that the digital world would be much improved if all senior executives were required to enroll in some kind of information protection program. And we must hold each employee responsible for protecting intellectual property. But this will be difficult because many executives lack a digital common sense."

Although business executives understand that protecting intellectual property is a key protective measure, perhaps their greatest responsibility is to lead by example and establish the climate of the corporate culture within any company. What you take as common sense in the real world often translates poorly to the digital networked world. It's vital that the enterprise leadership team foster a culture of awareness, stewardship for corporate information, and heightened *digital common sense*. You must develop and hone your digital common sense so that you can implement smart, cost-effective approaches that protect the safety of your corporate assets in the networked digital world.

In the integrated services network, each service you provide may have its own security requirements. These requirements often vary based on the intended use of the service. Services that are only intended for use inside the corporate network are likely to require very different protection methods than services provided for external use. For example, if VoIP is used only for internal calls, completely blocking all external access to the server may be adequate.

Servers that provide internal service to the organization shouldn't normally be co-located on the same server hardware that is used to provide external services. One common best practice is to isolate external traffic to a set of outside-facing DMZ segments while using another inside-facing set of segments for access from inside. These segments are often protected by implementing firewalls to allow only authorized communications between the inside and outside.

Internal communications services still require design review considerations:

- Will every employee use VoIP services?
- Do you allow vendors, contractors, or visitors to connect to the internal trusted network?
- Will these vendors, contractors, or visitors ever be authorized VoIP users?

If not, special care must be taken to ensure that only authorized VoIP users can gain access to servers, gateways, and other VoIP service delivery elements.

Protect the Protection

Perhaps the most critical component in establishing a network security posture is protection of the network services security and management platforms themselves. It's a fairly common network design practice to implement a dedicated management segment within the network. Because this management segment has oversight of the whole network, systems on this segment may require unfettered access to the entire network. The reverse is not true. This segment represents in some ways, the "keys to the kingdom" of the enterprise services network. Management servers should be accessible only to trusted employees in authorized work groups. Only authorized staff should be permitted to view log data for analysis. Intrusion detection and prevention systems shouldn't be seen or touched by anyone outside the network security staff.

It is vital that network management and security protection services themselves be protected. Firewall rules and ACLs can define authorized network locations by IP addresses that are granted access. Strong user authentication techniques like those described earlier can ensure that only authorized users have permission. Discovery protocols such as Simple Network Management Protocol (SNMP) and Internet Control Message Protocol (ICMP) should be disabled if they aren't necessary.

Detection

Although prevention tools are valuable aids to securing the integrated services network, you can't prevent everything bad that might occur. Tools help with prevention, but ultimately your survival depends on detection. When prevention fails, and it will fail, you must have strong detection and notification mechanisms. You need to know as much as possible about every event in the network. Comprehensive detection tools can shorten response time and lead to quicker remediation. Reliable incident management techniques help shift the security posture from a reactive mode to a proactive one. These measures make pre-emptive security changes part of the corporate culture. Like every process, good incident management involves continuous review and assessment.

It's important to take measures that will allow detection of not only situations in which information has been damaged, altered, or stolen, but also about how it was damaged, altered, or stolen, and who has caused the damage. This can be accomplished through procedures and implementing technologies discussed earlier that can detect intrusions, damage or alterations, and viruses.

Reaction

Just as every facet of network service delivery has an accompanying life cycle, everything you do in network security relies on layers of protection. This chapter has looked at some of the steps you can take to prevent security breaches, but you know breaches will happen. You establish procedures and implement tools to detect security breaches as they occur. You do so to enable a quick reaction. By building comprehensive processes into your day-to-day operations, you can be primarily either reactive or proactive. How an organization responds to an incident is driven by how well prepared everyone is.

Proactive and Reactive Strategies

A comprehensive security response plan should include both proactive and reactive strategies. The proactive strategy is really a pre-attack strategy. It documents a series of steps taken to reduce any existing security policy vulnerabilities. It also involves developing contingency plans. Determining the damage that an attack can cause, along with the vulnerabilities exploited during an attack, will aid in building a complete proactive strategy.

The reactive strategy is a post-attack strategy. It provides tools for assessing any damage caused by the attack and identifies steps to repair the damage. A reactive strategy also engages any contingency plans developed as part of the proactive strategy. During reaction, you restore normal business operations, then document and learn from the experience.

Incident management is really a matter of addressing a few simple questions:

- What happened?
- Where did it originate?
- Who did the incident impact?
- What steps are needed to mitigate the problem?
- How can it be prevented in the future?

Reacting to a Security Incident

We spend a lot of time trying to prevent security incidents from occurring. What sometimes gets lost in all of this preparation is plans for dealing with an incident should it actually occur. This is a very brief synopsis of several incident-handling guides that provides a high-level framework for dealing with either the realization that a system has been compromised or the recognition that a system is under active attack. The SANS Reading Room has a large set of papers about incident response at http://www.sans.org/reading_room/.

Remain Calm

To successfully handle any perceived emergency situation, you must remain calm so that you can assess what is going on around you and react in a methodical manner. A compromised system/network or an attacker on the loose demands well-thought out action; and frankly, the bad guys have probably been in your computer for days or maybe even longer, and another few minutes won't make much difference. You're probably going to have to rebuild your compromised servers anyway.

Notify Your Organization's Management and Activate Your Response Plan to Get Help

Your security policies should identify the pecking order of who gets called when if there's a security event. Individuals with particular responsibility for the affected server(s) and/or network(s) should be notified as well as any information security personnel. The severity of the incident (and your own policies) will dictate who else is brought in—your Internet Service Provider (ISP), department head, corporate officers, the press, law enforcement, consultants, response centers, and so on. Notify whoever is necessary to assess the situation and get it under control, but it is generally best to maintain a "need to know" stance and communicate, at least initially, with only the necessary parties.

Whenever possible, use telephones and faxes during a computer security incident. If the attackers have full access to your computer, they can possibly (probably?) read your mail. If you use your computer, this allows them to know when you report the incident and what response you got. There is a real possibility that other systems at your site have also been compromised and one or more packet sniffers are running on your network. Thus, if you absolutely must use a computer to communicate and you are fairly certain it can't be intercepted, use a different system and/or dial-up ISP access if possible.

Take Good Notes

This cannot be stressed enough—document, document, document!! Maintain a log of everything you see and do, everyone you speak with, and the team working with you. This will not only help you in criminal cases (and in remembering the events at a trial that might take place a year or two down the road) but also in the investigation/forensics process, post-event analysis, and as an educational/intelligence gathering vehicle for others in the InfoSec community. Notes should be detailed, organized, and complete, and should reflect the basic who, what, where, when, and how ("why" might be left for later on). Keep copies of any altered files before restoring your system.

Contain the Problem

Take any necessary steps to keep the problem contained and prevent it from spreading to other systems and/or networks. This may well involve disconnecting the compromised system from your network and/or disconnecting your network from the Internet. Containment may require a physical disconnect or might be accomplished while you clean up and recover; circumstances, including whether you are dealing with an active attack or the aftermath of an attack, will dictate what is a prudent action. Note that the latter approach (containing the problem while still online) might well leave you vulnerable to additional attacks.

Gather Evidence and Make Backups

For purposes of learning what happened and to have evidence for future analysis (and possible prosecution), make backups of OS and file system information as well as any state and network information (for example, output from netstat or route). Keep a detailed history of this activity if you have even the least suspicion that this information will be used in a criminal or civil trial; digital signatures and file timestamps are part of the procedures you should follow to maintain the custody chain. If possible, coordinate your evidence gathering with that of a second source, such as an ISP or another network (if you detect another network that is involved in this incident). Finally, as you make the backups, consider where they are going and who will be using them; if possible, make multiple copies and secure one for historical purposes while analyzing/sharing the other(s).

Get Rid of the Problem and Maintain your Business

This step might be easier said than done. If your server has been compromised, you should totally rebuild it from scratch unless you are 100% certain what the entire problem is. Before you can eliminate the problem, you need to be sure that you understand the cause of the incident. What vulnerability did the intruder use to gain access and what have you done to prevent another attack? You should rebuild the server and applications from original media.

The next issue, of course, is that of re-installing content. Note that some files that were exploited might have been on your system for some time already and, therefore, might have been backed up as part of your regular operations. So, although you've rebuilt the OS and applications software, you might very well reinstall files that can be exploited over and over. Again, it is imperative that you understand how an incident happened to avoid this eventuality.

Finally, business continuity is a major issue. Get rid of the problem and get your server back online as soon as possible.

Perform a Post Mortem

Once the situation is resolved and you're back in operation, get all relevant parties together to review the incident and the response. Review your security policies and operational procedures to see what changes, if any, are required. To the extent possible, contact appropriate incident response agencies, such as US-CERT, and share your knowledge with them.

Hold this meeting a day or two after the incident is deemed "over," when everyone is rested and has had time to reflect on what happened and why, and what went well and what didn't go so well. Don't do it immediately while people are still tired and don't wait weeks when people will forget and will have moved onto other things.

Reactive Strategies

A reactive strategy is implemented when the proactive strategy for the attack has failed. The reactive strategy defines the steps that must be taken after or during an attack. It helps to identify the damage that was caused and the vulnerabilities that were exploited in the attack, determine why it took place, repair the damage that was caused by it, and implement a contingency plan if one exists. Both the reactive and proactive strategies work together to develop security policies and controls to minimize attacks and the damage caused during them.

The enterprise incident response team should be included in the steps taken before, during, and after any attack to help assess the incident, document what happened, and to learn from the event.

Reactive strategies include responding to calls for help and reporting of incidents, threats, or attacks. These steps might be manually driven processes or triggered through technology in IDS and firewall monitoring systems.

From an operational perspective, incident handling involves taking reports, dealing with information requests, assisting with triage, and analyzing incidents. The incident response team can coordinate triage efforts between different divisions or work groups within the company and help share appropriate mitigation strategies. This team can help monitor for malicious activity in other parts of the network. They may be directly involved in rebuilding systems, applying patches, and developing any workarounds. The incident response team provides communications support, notifying work groups across the enterprise about vulnerabilities and sharing information on remediation tactics.

Vulnerability response might also involve researching to find the necessary patches, hotfixes, or workarounds. Response also involves notifying other workgroups of the mitigation strategy through advisories or alerts. In some enterprises, this service will include patch management procedures to install appropriate patches, fixes, or workarounds.

Proactive Strategies

Proactive strategies are designed to avoid incidents in the first place. They also help control the scope of impact when incidents inevitably do occur. The technologies used by the incident response team to handle incidents are often used to watch new technical developments, intruder activities, and trends that might indicate future threats. Members of the incident response team need to be afforded the time to read security mailing lists, security Web sites, and current news articles to mine relevant information from the broader security community. This intelligence-gathering function can be combined with past lessons learned to provide a powerful defensive resource to the enterprise.

Ongoing, regularly scheduled security assessments can provide detailed review and analysis of an organization's security infrastructure. These assessments might be comprehensive system audits or simply desktop reviews of security practices. There are several types that can be performed:

- Reviews of hardware and software configurations, routers, firewalls, switches, servers, and workstations. These reviews can confirm that systems all meet enterprise standards or follow industry best practices.
- Interviews with employees can provide insight into how well actual security practices match the enterprise security policy.
- Vulnerability or virus scans can be performed against network segments to identify vulnerable systems and networks.
- Penetration tests can be conducted in a controlled, methodical environment, including social engineering, physical, and network attacks. This information can be leveraged to provide a stronger security posture before a malicious outsider takes advantage of any weaknesses that are discovered.

Security tools vary in scope and functionality. Because of their focused experience and understanding of security issues, the incident response team may provide valuable guidance on how to harden systems and develop a set of approved security tools.

Incident Management In a Nutshell

Assess the Damage

Determine the damage that was caused during the attack. This should be done as swiftly as possible so that restore operations can begin. If it is not possible to assess the damage in a timely manner, a contingency plan should be implemented so that normal business operations and productivity can continue.

Determine the Cause of the Damage

To determine the cause of the damage, it is necessary to understand what resources the attack was aimed at and what vulnerabilities were exploited to gain access or disrupt services. Review system logs, audit logs, and audit trails. These reviews often help in discovering where the attack originated in the system and what other resources were affected.

Repair the Damage

It is very important that the damage be repaired as quickly as possible in order to restore normal business operations and any data lost during the attack. The organization's disaster recovery plans and procedures should cover the restore strategy. The incident response team should also be available to handle the restore and recovery process and to provide guidance on the recovery process.

Document and Learn

It is important that once the attack has taken place, it is documented. Documentation should cover all aspects of the attack that are known, including the damage that is caused (hardware, software, data loss, loss in productivity), the vulnerabilities and weaknesses that were exploited during the attack, the amount of production time lost, and the procedures taken to repair the damage. Documentation will help to modify proactive strategies for preventing future attacks or minimizing damages.

Testing the Strategy

The last element of a security strategy, testing and reviewing the test outcomes, is carried out after the reactive and proactive strategies have been put into place. Performing simulation attacks on a test or lab system makes it possible to assess where the various vulnerabilities exist and adjust security policies and controls accordingly.

These tests should not be performed on a live production system because the outcome could be disastrous. Yet the absence of labs and test computers due to budget restrictions might preclude simulating attacks. To secure the necessary funds for testing, it is important to make management aware of the risks and consequences of an attack as well as the security measures that can be taken to protect the system, including testing procedures. If possible, all attack scenarios should be physically tested and documented to determine the best possible security policies and controls to be implemented.

To be effective, both proactive and reactive strategies need to be tested or exercised. Regular simulation exercise can help maintain a vigilant security posture.

It's altogether too easy to work through an incident that only has minor repercussions, and gloss over the lessons learned. Remember that lessons learned are only learned when they feed into some process improvement loop and a change is actually put into practice. Don't let your organization get trapped into collecting a set of lessons observed but never acted on. Learn the lessons well and incorporate them into the every day work flow to maintain a vigilant security posture.

Summary

In the converged services network, security goals will be largely determined by the following key tradeoffs:

- **Services Offered vs. Security Provided**—Each service offered to users carries its own security risks. For some services, the risk outweighs the benefit of the service and the administrator may choose to eliminate the service rather than try to secure it.
- **Ease of Use vs. Security**—The easiest system to use would allow access to any user and require no passwords; that is, there would be no security. Requiring passwords makes the system a little less convenient, but more secure. Requiring a device-generated one-time password makes the system even more difficult to use, but much more secure.
- **Cost of Security vs. Risk of Loss**—There are many costs to security: monetary (the cost of purchasing security hardware and software such as firewalls and one-time password generators), performance (encryption and decryption take time), and ease of use. There are also many levels of risk: loss of privacy (the reading of information by unauthorized individuals), loss of data (the corruption or erasure of information), and the loss of service (the filling of data storage space, usage of computational resources, and denial of network access). Each type of cost must be weighed against each type of loss.

The threats to the security of your service network environment have reached staggering proportions from both within and without. This chapter has identified a number of widely accepted industry best practices to help secure and closely monitor your environment. The next and final chapter will delve further into compliance, asset management and reporting, and documenting the whole environment.

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>.