# Microsoft Solutions for Security

## *The Microsoft Windows NT 4.0 and Windows 98 Threat Mitigation Guide*

Microsoft®

# Table of Contents

# 1

# Introduction

## Executive Summary

Interest and investment in secure computing systems have escalated and changed a great deal over the last ten years. Some factors in this shift might have been predicted, but several—including the widespread use of the Internet, the broad deployment of always-on broadband connections, and the increasing use of personal computer-based hardware in environments as diverse as retail stores, cars, and entertainment devices—would have been extremely difficult to foresee. One such issue is the fact that organizations today are using earlier versions of the Microsoft® Windows® operating system that now face security threats that did not exist—and could not reasonably have been foreseen—when those versions of Windows were designed years ago.

Newer versions of Windows offer significantly increased security because they have been designed to protect against the kinds of threats common in today's computing environments. Microsoft recommends that you upgrade client and server operating systems to these newer versions to improve their security and the security of the networks to which they belong. However, the use of earlier versions of Windows in widely deployed business applications (such as point-of-sale and store management systems and branch office desktop clients) means that not all organizations can easily and quickly upgrade their systems. Those organizations with large numbers of custom applications running on earlier versions of Windows are especially challenged, because the applications themselves need updating—which is not always easy or even possible in some cases.

For organizations that are not able to immediately update all of their older systems, this guidance provides prescriptive information and test plans for strengthening the security of (or *hardening)* clients and servers running earlier versions of Windows. This guidance is designed to protect these systems to the greatest degree possible while migration plans are put into place.

By following the recommendations here, you will be able to reduce risks to deployed systems. It is important to keep in mind that many significant security features and improvements in later versions of Windows will still not be available. This strategy provides improved protection *now*, giving owners of older systems the time to consider the best migration strategy for future deployments. However, this guidance is no substitute for a properly designed migration to more secure versions of Windows.

## The Business Challenges

Many organizations depend absolutely on the integrity and availability of their information systems. Security is not typically considered a core business function, but it is important precisely because the *real* core business functions of an organization depend on it. The key business challenges that enterprises face include:

- Guarding systems and data against targeted or random attacks. Many attacks are more or less random, but that does not make them any less dangerous. Modern business operations have to be able to protect important assets by quantifying what the assets are worth and what threatens them.
- Improving security while maintaining compatibility and controlling cost. IT managers often view the process of improving system security as a battle among compatibility, implementation and maintenance cost, and improved security. For older systems, compatibility often outweighs security. Security changes that break older applications may have an immediate, and serious, impact on business productivity and service quality. However, organizations must counterbalance this against the risk of not adequately securing important assets. To correctly balance these conflicting imperatives, it is important to view security holistically, not just as a goal that can be achieved via a few check boxes or button clicks.
- Providing flexible protection to end users. In many environments, it is difficult to achieve adequate security, because security measures inevitably conflict with user convenience. Highly-managed desktop systems offer good security at the expense of some user flexibility; older desktop operating systems are much less manageable and, hence, harder to secure.
- Justifying security costs. It is very difficult to quantify the actual financial cost of not implementing adequate security until *after* a compromise occurs. In this regard, security expenses can be thought of as insurance. However, as with insurance, the cost has to be carefully weighed against the benefits and level of protection offered, especially because the end date for regular support of these operating systems is rapidly approaching.

## The Business Benefits

Improving the security of your older systems can lead to some direct business benefits. These benefits include:

- Better protection. Microsoft Windows NT® 4.0 and Windows 98 do not support many of the security features developed for Windows 2000, Microsoft Windows Server™ 2003, and Windows XP, but most organizations that have older operating systems deployed are not even making use of all of the protections in the software they already have. Taking full advantage of these features can greatly improve protection when compared to using only the baseline installations of these operating systems.
- Better security for less money. In large part, you can implement the recommendations in this guidance with no additional software or licensing costs. However, they offer a significant increase in security. The cost/benefit ratio for the changes described in this guidance is very high.
- The ability to secure existing systems without breaking anything. Preserving application compatibility is important, but so is strengthening security. Careful application of the recommendations in this guidance will provide improved security *and* continued compatibility.

## Who Should Read This Guidance

The intended audience for this guidance includes architects, IT managers and administrators, technical decision makers, and consultants involved in securing an infrastructure in which Windows NT 4.0 and Windows 98 operating systems are still in production use.

**Note:** This guidance applies only to Windows 98 SR2 and Windows NT 4.0 Workstation and Server. Other versions of Windows (including Windows NT version 3.1 and 3.51, Windows 95, Windows Me, and Windows 3.11) are not covered by the prescriptive material herein.

## Reader Prerequisites

The following knowledge and skills are prerequisite for administrators or architects charged with developing, deploying, and securing installations of Windows NT version 4.0 and Windows 98 in an enterprise:

- MCSE certification for Windows NT, Windows 2000, or Microsoft Windows Server 2003 with two or more years of security–related experience.
- In-depth knowledge of the corporate domain structure (including the Microsoft Active Directory® directory service, if it has been deployed).
- Use of Windows management tools, including the Windows NT system policy editor (Poledit), the Microsoft Management Console (MMC), and the Security Configuration Manager (SCM).
- Experience deploying applications and workstations in enterprise environments.
- Familiarity with applications unique to your individual enterprise environment.

# Guidance Overview

This guidance describes the process of hardening networks and computers in environments with computers that run earlier versions of the Windows operating system. Organizations may have a variety of different combinations of computers running Windows NT 4.0 (Workstation, Server, and Advanced Server) and Windows 98, with or without later versions of Windows clients or servers. This guidance focuses on the protective measures you can apply to Windows NT 4.0 Workstation and Windows 98 clients and Windows NT 4.0 member servers in an Active Directory directory service domain environment to improve their security.

This guidance comprises eight chapters, grouped into two sections. The first section consists of two chapters, Chapter 1, "Introduction," and Chapter 2, "Applying the Security Risk Management Discipline to the Trey Research Scenario," both of which are intended for executives and IT management at all levels.

## Chapter 1: Introduction

This chapter provides an executive summary, introduces the business challenges and benefits surrounding improving the security of older systems, suggests the recommended audience for the guidance, lists the reader prerequisites, and provides an overview of the chapters and solution scenarios in the guidance.

## Chapter 2: Applying the Security Risk Management Discipline to the Trey Research Scenario

This chapter details a company scenario that is used to develop the recommendations in this guidance and explains how an IT generalist would assess the security risks and vulnerabilities of a network infrastructure. Trey Research, the fictitious company in the scenario, has its headquarters in Seattle and field offices in several states throughout the country. The chapter also describes how IT generalists can identify and prioritize their individual organizations' risks and vulnerabilities to generate security requirements that can drive an action plan to mitigate security threats.

The second section of the guidance contains six chapters of prescriptive information for IT administrators and technical managers. Each chapter begins with a discussion of design principles and options before covering the specific hardening measures chosen for the target scenario for this guidance.

## Chapter 3: Network Security and Hardening

This chapter describes network security vulnerabilities and the process of hardening network components (including client and server computers) against these vulnerabilities. The chapter addresses network segmentation, Transmission Control Protocol/Internet Protocol (TCP/IP) stack hardening, and the use of personal firewalls for client protection.

## Chapter 4: Hardening Windows NT 4.0

This chapter explains how to harden Windows NT 4.0 (Workstation and Server) by establishing a baseline for the system and then applying specific hardening measures. It describes the importance and methods of physical security and procedures and methods for applying security policies to file, print, Web, and application servers. The chapter discusses the inherent compromises in various security approaches and concludes by describing in detail the most advantageous hardening policies for Trey Research.

### Chapter 5: Hardening Windows 98

This chapter explains how to harden Windows 98 clients and applications, and describes methods for applying patches, updates, and security policies to computers running Windows 98.

### Chapter 6: Patch Management

One of the main ways to guard against attack is to ensure that your environment is kept up to date with all the necessary security patches. Patches are required at the server and client levels. This chapter shows you how to ensure that you find out about new patches in a timely manner, implement them quickly and reliably throughout your organization, and monitor your systems to ensure that they are deployed everywhere. It describes the compromises with patch management implementations and concludes with a detailed description of Trey Research patch management system.

### Chapter 7: Antivirus Protection

This chapter describes the importance of antivirus software and policies and the security and supportability of client-based and server-based antivirus solutions.

### Chapter 8: Conclusion

This chapter closes out the guidance by providing a brief summary of the hardening processes that have been discussed.

## Solution Scenario

The solution revolves around Trey Research, an environmental engineering firm headquartered in Seattle. Trey has about 55 servers and 500 employees spread among its offices in Georgia, Florida, Arizona, Pennsylvania, and Washington state. The Trey Research servers host a number of custom applications, including ones for controlling laboratory test and measurement equipment, database applications built on Microsoft SQL Server™, and collaboration applications built on Microsoft Exchange version 5.5. Perimeter network security is provided by a hardware firewall/router combination. Antivirus software is deployed on some computers, but not uniformly.

End users use their computers for office productivity applications, Internet access, e-mail, remote access to specialized services (including the Chemical Abstracts Service), and remote control of measurement and laboratory equipment. The custom applications that Trey's employees use range from Web-based Active Server Page (ASP) applications to applications built on top of FoxPro for Microsoft MS-DOS® and Microsoft Access 95. In some cases, the original source code for the application has been lost, so porting the applications would be impossible without completely starting over.

### Network

Figure 1.1 shows a portion of Trey's network, illustrating the Windows Server 2003 domain controllers and selected client and member server resources.

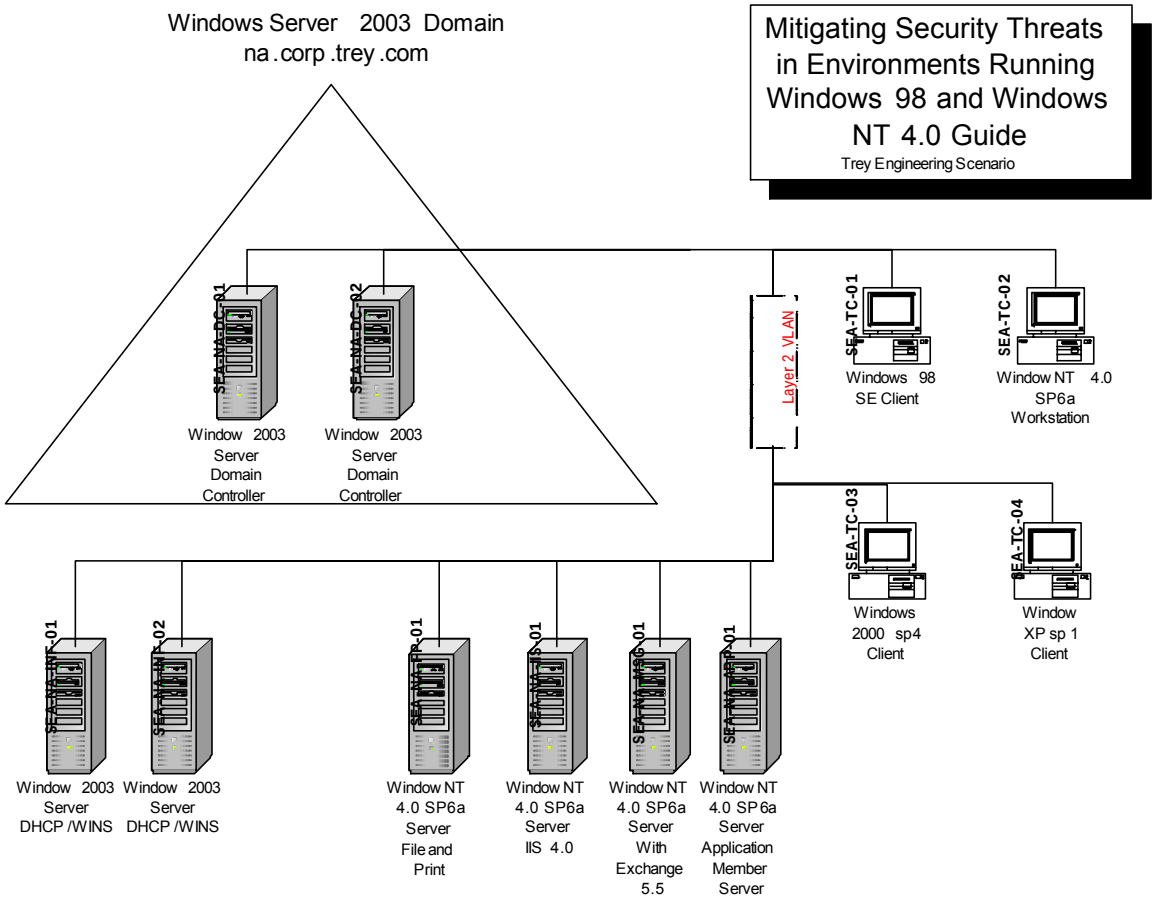


**Figure 1.1**
*A subset of the Trey Research network*

## Active Directory Design

Trey's original network design was built around several Windows NT 4.0 domains, with a mix of Windows 98 SR2 and Windows NT Workstation clients. Over time, Windows 2000 and Windows XP clients were introduced; the company currently has a combination of Windows 98 SR2, Windows NT Workstation 4.0, Windows 2000, and Windows XP.

Although many of the workstations and servers that run Trey's specialized analysis and control applications cannot be upgraded without substantial effort, Trey's IT staff realized that moving its domain infrastructure to Windows Server 2003 would immediately provide them with better security and reliability for domain logon services, Dynamic Host Configuration Protocol (DHCP), Windows Internet Naming Service (WINS), and Domain Name System (DNS), along with the possibility of using Active Directory Group Policy objects (GPOs) for policy application on Windows 2000 and Windows XP clients. Accordingly, Trey's current environment features a single Windows Server 2003 Active Directory domain with Windows Server 2003 infrastructure servers. Other network components remain on Windows NT 4.0 Server. Because Windows Server 2003 domain controllers can emulate Windows NT domain controller functionality, the existing clients are able to function in an Active Directory environment.

### Business Justification for Using Windows NT

Trey Research is currently planning a major re-engineering of its business processes and systems in an effort to become more competitive. This re-engineering involves examining every aspect of the company's business systems, including network design and workstation and server configurations. Until that process is completed, Trey's management is unwilling to invest in upgrading hardware, software, or operating systems except in exceptional cases. The increased security gained from an immediate migration of domain and infrastructure servers to Windows Server 2003 was one such exceptional case, which is why it was approved. For the broader portfolio of applications, though, Trey is working to identify an application migration strategy that fits its business requirements. Because hardware is not being upgraded, many of the benefits that end users might see from upgrading to Windows XP on the desktop cannot be fully realized, slowing demand for upgrades. This process is further complicated by the fact that Trey does not have source code for all of the applications that they use, meaning that Trey is not able to fix all applications that fail when the underlying operating system is upgraded.

## Supporting This Guidance in Your Environment

The scenario presented in this guidance was designed to reflect what many Microsoft customers are using in production. The recommendations presented here were tested in a configuration that replicates the Trey configuration scenario; because your own organizational configuration may differ, you should carefully examine the recommendations to ensure that they are appropriate for your environment. As you prepare to make changes to secure your environment, you should keep the following principles in mind:

- It is very important when making security-related configuration changes to maintain documentation on what settings are being modified by a given script, policy, or template being applied. Before you apply any settings changes to your computers, ensure that you have documented the current state *and* what changes you are applying.
- Applying broad, sweeping security changes to a production Windows network can result in a situation where it is very difficult to determine which configuration changes are the root cause of newly appearing problems. Your deployment strategy should call for applying changes in multiple stages, deploying each new change to small groups of computers. This approach helps limit the scope of problems if they do occur.
- Your deployment plans should not apply a large number of security changes in a single increment. Instead, you should divide the security changes into logical groupings and then deploy these changes discretely using separate policies, scripts, or templates.
- Your deployment plans should include provisions for rolling back changes if they cause problems. Rollback plans should provide for two contingencies: quick recovery by reverting back to the previous settings and reversion of a subset of computers to verify that in fact that the security change did cause the problem. Rollback plans should be tested alongside the deployed changes in the test environment prior to any deployment of the changes to the production network.

# 2

# Applying the Security Risk Management Discipline to the Trey Research Scenario

This chapter describes the fundamentals of how to apply a structured and repeatable method of risk analysis for information systems. A thorough risk management process should be established in all organizations to assess where the expenditure of time and effort to secure systems will provide the best security and return on investment (ROI).

The first step toward improving the security of older applications and clients on business networks is to conduct a thorough analysis of the threats and risks associated with the environment, applications, users, and network. Microsoft recommends that every network be analyzed with a well-defined process such as the Microsoft Security Risk Management Discipline (SRMD). SRMD provides a structured, repeatable process for evaluating the assets that an organization has, risks that threaten them, vulnerabilities that may allow an attacker to steal or damage the assets, and countermeasures that can be applied to mitigate or transfer the risks.

A complete discussion of SRMD is outside the scope of this guidance, but it is sufficient to know that SRMD provides a framework for identifying assets, quantifying their value, and identifying and quantifying threats that those assets face so that organizations can make sound decisions about appropriate and cost-effective security activities.

---

**Note:** For more detailed information on SRMD, see the references in the "More Information" section at the end of this chapter.

---

# Scenario Detail

Trey Research specializes in wastewater analysis, monitoring, and treatment. Trey maintains its headquarters in Seattle and has field offices in Georgia, Florida, Arizona, and Pennsylvania. Trey has a total of just under 500 employees comprised of field workers, lab technicians, and scientists, along with a few administrative personnel.

Trey Research customers include local and state governments who need specialized assessment services (such as measuring mercury levels in groundwater); construction companies who need to perform site tests before, during, and after construction; industrial and manufacturing companies that need ongoing monitoring of their facilities; and others who need emergency environmental monitoring or cleanup. The data that Trey gathers and the analysis that results is often financially or legally sensitive. When Trey engineers are asked to be expert witnesses, there are specific chain-of-evidence requirements that must be met for the life cycle of the data they gather.

Field workers keep paper logs of measurements, which they manually enter when they come back to their offices. A few of the engineers use mobile computers with Microsoft® Windows® 98 to enter data directly while in the field, but this analysis system is specific to a few of Trey's largest customers.

During the last three years, Trey has been growing at approximately 20 percent per year. This growth rate finally prompted the company's CEO to hire an IT director to build and supervise a plan for modernizing their information systems. The IT director began in late 2003, and the first job was to perform a risk analysis to better understand the value of the company's computing assets and the vulnerabilities that they faced. As a result of this analysis, Trey made a number of very rapid changes to its IT environment. The first significant change was upgrading the domain structure to the Microsoft Active Directory® directory service and Microsoft Windows Server™ 2003. This upgrade provided an immediate security increase for domain accounts and made it possible to apply additional Group Policy controls to sensitive computers like those used by executives and their staffs. In addition, the company accelerated its technology modernization plan so that initial deployment of its new analysis and collection system (based on Windows XP and Windows XP Tablet PC Edition) will begin earlier than planned.

However, Trey also chose to invest in hardening its existing systems to reduce the risk of data loss or compromise between now and the time that the new system is fully deployed. The CEO has given the IT director one month to identify, prioritize, and mitigate the most immediate threats and to investigate the upgrading of the company's Microsoft Windows NT® 4.0 systems. Although this is a short time, Trey takes the security threats involved very seriously and is acting aggressively to protect itself as much as possible.

## Network

The recommendations and settings described in this document were tested on a simulation of the Trey network, configured as shown in the following figure:

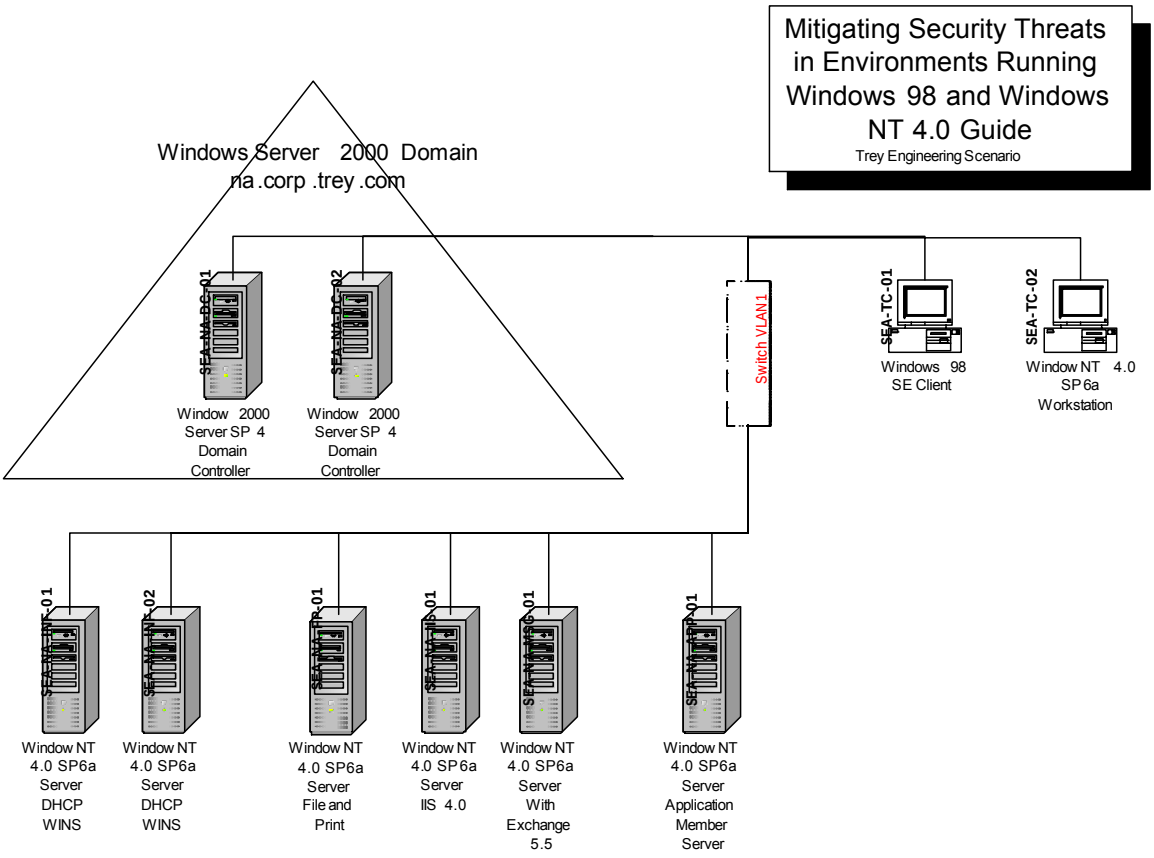


**Figure 2.1**

*Network testing subset of the Trey Research network*

## Active Directory Design

Trey Research maintains a single Active Directory domain for the organization. It chose this structure because it offers ease of maintenance and good control. All field offices are connected to Trey's headquarters through leased private lines, so there is no need to build subdomains for individual branch offices.

## Business Requirements

Trey has five primary business requirements related to the security of the company's systems and networks:

- Maintain the integrity of systems against compromise by outside attackers. This requirement requires hardening the network against penetrations, improving auditing and logging, and reducing systems' vulnerability to widely known exploits.
- Maintain normal business operations after all security measures are in place. Much of the analysis work performed is time-sensitive, so frequent or lasting interruptions would be unacceptable.

- Maintain confidentiality of information where necessary. Some of the information that Trey manages is extremely sensitive, and the company wants to avoid potential liability resulting from disclosure.
- Provide increased protection against malicious code on the network. Trey has a fairly liberal acceptable use policy, so many users are accustomed to downloading and installing software on their own. This situation has led to security and performance problems in the past. One goal of the hardening is to make the company's systems less vulnerable to downloaded *malware* (malicious software).
- Provide an automated method of auditing and distributing security patches.

# Applying the Security Risk Management Discipline

The goal of the SRMD is to provide a way to quantify risks and then reduce those that are within the organization's control to mitigate. To do this, the SRMD defines risk management as an ongoing process with four primary stages as shown in the following figure:

1. Assess Risk. Identify and prioritize risks to the organization. These risks may or may not be associated with specific IT systems or assets.
2. Decision Support. Identify and select control solutions based on a defined cost-benefit analysis process.
3. Implement Controls. Deploy and operate holistic control solutions to reduce risk to the organization.
4. Measure Risk. Determine and report on the effectiveness of deployed controls to manage risk to an acceptable level.


**Figure 2.2**
*The SRMD cycle*

The *Microsoft Security Risk Management Discipline Guide* (to be published later in 2004) describes SRMD in detail. The Trey IT staff reviewed the SRMD material and developed a plan to do the following:

1. Assess the risks, a three-step process that requires Trey to build a plan for evaluating risks, gather data about the actual degree of risk and the organization's vulnerability therein, and prioritize those risks in order of severity and cost.
2. Use the risk assessment to make decisions about specific controls to apply based on the actual degree of risk present.
3. Implement the selected controls. The remaining chapters of this guide are dedicated to discussing controls that can be applied to mitigate specific types of risks.
4. Evaluate the effects of the applied controls on the risk and on the organization's environment.

This chapter will focus on the first two steps and explain how the Trey IT staff adapted the SRMD to its environment to help the company begin the SRMD-driven process of risk management. The remaining chapters will focus on the third step, applying the actual controls.

# Assessing Risks

The first significant step that Trey must take to commence its security hardening process is to assess the risks and threats the company actually faces. This process required Trey to link together several separate steps:

1. Identify the roles and function of each class of computers on the network.
2. Map the communications among different roles. For example, application servers need to communicate with domain controllers and with user workstations. This mapping should pinpoint the protocols, ports, and traffic patterns used for these communications.
3. Identify potential threats that can exploit the computers in various roles.
4. Determine the probability or likelihood that particular threats may apply to a given role.

## Identifying Roles

For most networks, the process of identifying the roles filled by computers on the network is straightforward. By consulting the physical inventory of systems owned by the company, the Trey IT department was able to generate the data in the following table, which lists the key roles in use on its network, the operating systems used for those roles, and the location and hardware types commonly found in those roles. All of this information is pertinent to the threat modeling process.

**Table 2.1: Trey Computer Roles**

| Role | Operating systems used in role | Location | Hardware type |
|---|---|---|---|
| Application / Web server | Windows NT 4.0 | HQ | Conventional server |
| Dynamic Host Configuration Protocol (DHCP) servers | Windows Server 2003 | HQ, field offices | Conventional servers |
| Domain Name System (DNS) servers | Windows Server 2003 | HQ | Conventional servers |
| Domain controller | Windows Server 2003 | HQ, field offices | Conventional server |
| Executive mobile computers | Windows 2000, Windows XP | Mobile | Mobile computers |
| Executive/special purpose workstations | Windows XP | HQ | Conventional desktops |
| Field engineer systems | Windows 98 | Mobile | Mobile computers |
| File/print server | Windows NT 4.0 | HQ, field offices | Conventional server |
| Messaging server | Windows NT 4.0 | HQ | Conventional server |
| Special-purpose control systems | Windows NT 4.0, some Windows 98 | Field offices | Mix of conventional servers and desktops |
| User workstations | Windows 98, some | HQ, field offices | Conventional desktop |
| Windows Internet Name Service (WINS) servers | Windows Server 2003 | HQ, field offices | Conventional servers |

# Mapping Communications

After the computer roles have been identified, it is possible to begin determining what kind of network communications take place among different roles. This determination makes it possible for you to specify which types of traffic should and should not be permitted between your network as a whole and those segments that contain computers running older versions of the Windows operating system.

## Modeling the Network

Modeling the network is very simple. The Trey engineers merely took a diagram of their existing network and used it as the basis for their network map. Such maps should indicate the physical location, network address, and operating system type of each computer on the network. Ideally, they should also visually indicate the location of routers and firewalls and how the network is segmented.

## Adding Data Flow Information

After you have a network map, the next step is to superimpose information about the flow of data on it. This is commonly done by using the Yourdon-DeMarco data flow diagram (DFD) method, which indicates data flowing between systems or objects by directed lines. Each line can be labeled with the port number or protocol in use. The result is a diagram that shows the flow of communications between each set of roles. This diagram makes it very simple to configure firewalls and port/packet filters to allow only specified traffic types. Also, data flow information can easily be used in the future to construct Internet Protocol security (IPsec) filter rules for organizations that have Windows 2000, Windows XP, and Windows Server 2003 deployed.

# Identifying and Modeling Threats

A threat model is an attempt to enhance the security of a distributed system by producing an inventory of all of the threats posed to the system, regardless of origin. The underlying concept is that if you can identify as many threats as possible, knowing what threats exist will make it easier to either mitigate the threats or rule them out. You can base these decisions on whether the mitigation is not possible, is too difficult or expensive, or because the threat is not significant or likely enough to mitigate. A basic concept when creating the threat model is to enumerate all of the realistic threats—*including ones that you know have been protected against already*. Sometimes, discussion of threats that have some current defense may lead to the revelation of similar or tangential attacks.

# Identifying Threats

After the Trey IT engineers built a map of their network indicating the roles present on the network and the communication methods used between computers in and among the various roles, they were prepared to begin identifying and prioritizing specific threats. These threats can be separated into several categories:

- Threats to the physical security or integrity of computers. These threats include fire, flooding, loss of electrical power, accidental or intentional physical damage, and compromises caused by unauthorized physical access.
- Denial-of-service (DoS) attacks involving individual computers, infrastructure services, or the network itself.
- Execution of malicious code, including viruses, worms, and Trojan horses.

- Unauthorized disclosure of sensitive information through network monitoring, account compromise, or other means.
- Compromises caused by loss of control over user or privileged accounts (including those caused by weak passwords, inadequate controls on privileged accounts, failure to follow security procedures, or inadequate auditing).

Each category of threats contains a variety of individual threats, some of which have already been mitigated and some that are very difficult to effectively mitigate in the Trey environment. Each of the following sections describes a class of threats and the measures available to Trey to mitigate them. Note that in many cases, the listed mitigation measures are only partially effective. Only measures available in Windows NT 4.0 or Windows 98 are shown; more effective measures are available in newer releases of Windows.

## Physical Security Threats

The following table shows the key physical security threats that Trey identified for its networks. Most of these threats arise from factors that are outside of the company's controls and can only be effectively mitigated by setting up policies to provide disaster recovery and business continuance processes, which are outside the scope of this guidance.

**Note:** The "Impact and scope" and "Likelihood" columns in the following tables represent the Trey IT department's best estimate of the nature, scope, and probability of each specified threat. The specific values of these columns may vary widely between organizations.

**Table 2.2: Physical Security Threats and Mitigations**

| Threat | Details / attack vector | Impact and scope | Like lihood | Available mitigations |
|---|---|---|---|---|
| Environmental damage | Fire, flood, weather, or other external environmental factors. | High / entire network | Low | Insurance; disaster recovery and business continuance plans. |
| Temporary loss of infrastructure services | Loss of wide area network (WAN) / Internet connectivity, power, cooling, or other critical infrastructure service not provided by Trey. | Medium / entire network | Medium | These outages tend to be short-term. |
| Physical damage to key computers | Accidental or purposeful damage. | Medium / single machine | Low | Backups; physical access controls for sensitive computers. |
| Compromise of single computer | Physical access and compromise of a computer by an attacker. | High / single machine | Low | Physical access controls; boot hardening; strong passwords on local admin accounts; use of Syskey to protect local Security Account Manager database data. |

## Denial-of-Service Threats

DoS threats involve the loss of access to network services or computers because of an intentional attempt to block or overwhelm network equipment or computers with bogus traffic. These threats are normally mitigated at the network perimeter. The following table shows the key DoS threats that Trey identified as significant for its networks.

**Table 2.3: Denial-of-Service Threats and Mitigations**

| Threat | Details / attack vector | Impact and scope | Likelihood | Available mitigations |
|---|---|---|---|---|
| Network traffic tampering or spoofing | Attacker sends inappropriate / malformed messages to hosts. | High / entire network | Low | Network ingress filtering. |
| Tampering with DNS services | Attacker spoofs, pollutes, or blocks DNS traffic. | High / entire network | Low | Monitoring of DNS service quality to quickly detect service problems. |
| Targeted traffic tampering or spoofing | Attacker targets individual computers or assets. | High / single computer | Low | Port and packet filtering; network segmentation; personal firewalls. |
| User account lockout | Attacker exceeds the maximum number of permitted password attempts, triggering the account lockout policy. | Medium / entire network | Low | Deployment of account lockout policy with no lockout count. |
| Service account lockout | Attacker denies access to a service account by exceeding the password retry count. | Medium / entire network | Low | Deployment of account lockout policy with no lockout count. |
| Bandwidth consumption attack | Attacker intentionally consumes bandwidth to target network or device. | Medium / entire network | Low | For perimeter network, ingress filtering and Internet service provider (ISP) monitoring. For internal hosts, Transmission Control. Protocol/Internet Protocol (TCP/IP) stack hardening and ingress filtering. |
| DHCP / WINS services tampered with | Attacker spoofs, pollutes, or blocks DHCP or WINS traffic between clients and infrastructure servers. | Low / single computer | Low | Multiple DHCP and WINS servers to provide overlapping scope. |

## Malicious Code Threats

The following table shows the key physical security threats that Trey identified for its networks. Like physical threats, most of these threats arise from factors that are outside of the company's controls and can only be effectively mitigated by setting up policies to provide disaster recovery and business continuance processes, which are outside the scope of this guidance.

**Table 2.4: Malicious Code Threats and Mitigations**

| Threat | Details/attack vector | Impact and scope | Likelihood | Available mitigations |
|---|---|---|---|---|
| Virus outbreak | Virus spreads after being introduced to Trey network by an internal user. | High / entire network | Medium | Deployment of client and server antivirus software; user education; patch management; segregation of older computers. |
| User execution of malicious code | User downloads and runs malicious code disguised as something innocuous. | High / single computer | Medium | Microsoft Internet Explorer hardening; user education. |
| Worm outbreak | Worm spreads after being introduced from the Internet or through an infected internal computer. | High / entire network | Low | Patch management to reduce exploitable vulnerabilities; segregation of older computers. |

## Information Disclosure Threats

Information disclosure threats include accidental leakage of confidential data, purposeful disclosure by authorized users to unauthorized parties, and targeted attacks to disclose data.

**Table 2.5: Information Disclosure Threats and Mitigations**

| Threat | Details/attack vector | Impact and scope | Likelihood | Available mitigations |
|---|---|---|---|---|
| Network sniffing | Attacker surreptitiously monitors network traffic to capture passwords or other sensitive data. | High / entire network | Medium | Physical access controls for network; Server Message Block (SMB) signing; use of Windows NT LAN Manager version 2 (NTLMv2) instead of NTLM or LM authentication. |
| Theft of data from mobile/laptop computers | Attacker steals computer and recovers data from it | High / entire network | Medium | None |
| Leakage of password data | Attacker steals password hashes from compromised computer or network. | High / entire network | Low | Physical access controls for domain controllers; use of Syskey; NTLMv2. |
| Purposeful information disclosure | Authorized user discloses information to an unauthorized party. | High / single computer | Low | None |

## Account Compromise Threats

Account compromise threats can be separated into two broad categories: compromises that occur because an attacker gains physical access to a computer (and can thus remove the local administrator password, install a keystroke logger, or otherwise tamper with the computer), and network-based attacks. The following table shows the most significant account compromise threats with which Trey is concerned.

**Table 2.6: Account Compromise Threats and Mitigations**

| Threat | Details / attack vector | Impact and scope | Likelihood | Available mitigations |
|---|---|---|---|---|
| Domain administrator account compromise | An attacker obtains the password for a domain administrator account. | High / entire network | Low | Physical access controls. |
| Local administrator account compromise on individual computer | An attacker learns the password for a local administrator account through password cracking or other means. | High / single computer | Low | Physical security controls; NTLMv2 authentication. |
| Local administrator account password reset on individual computer | An attacker gains physical access to a computer and resets its local administrator password. | High / single computer | Low | Physical security controls. |
| User account compromise | An attacker obtains access to an ordinary user account. | Medium / single computer | Low | Physical access controls; SMB signing; NTLMv2 authentication. |

# Making Risk Decisions

After the Trey IT staff identified the most significant risks that the organization faced (as listed and prioritized in the preceding tables), the staff decided which mitigation measures to take based on the potential impact and likelihood of individual threats. Some of the most significant threats cannot effectively be mitigated *at all* on computers running Windows 98 and Windows NT 4.0, which is why Trey decided to migrate its infrastructure systems to Windows Server 2003. Other risks may be mitigated by a combination of operating system-specific steps, network configurations, and policy changes. By examining each of the potential threats and calculating the cost required to defend against them, Trey developed a plan for mitigating as many serious risks as possible. The remaining chapters of this guidance describe the specific measures on which Trey decided.

# Summary

This chapter has described some of the considerations involved with applying the SRMD to a common customer scenario. All information provided for this example is based on actual data; however, this information represents only a fragment of the overall information required for an organization to be able to perform a thorough security risk assessment. Including the entire risk analysis table or every security risk statement would have made the information provided in this chapter more difficult to understand. Instead, relevant examples were highlighted for quick reference and easy comprehension.

The guidelines in this chapter were applied to develop a list of risks that are addressed with specific remediation steps. After the Trey engineers complete the list, they can proceed to identify the steps required to mitigate the risks by securing their systems to protect against the listed vulnerabilities. These remaining chapters in this guide examine these steps in detail.

## More Information

For more information about how to apply SRMD to an enterprise environment, see the following resources:

- The “Understanding the Security Risk Management Discipline” module on Microsoft TechNet at http://www.microsoft.com/technet/security/guidance/secmod134.mspx.
- The “Applying the Security Risk Management Discipline” module at http://www.microsoft.com/technet/security/guidance/secmod135.mspx of "Identifying and Managing Security Risks" on TechNet.

# 3

# Network Hardening and Security

This chapter describes network security vulnerabilities and the process of hardening hosts on the network against these vulnerabilities. It addresses network segmentation, Transmission Control Protocol/Internet Protocol (TCP/IP) stack hardening, and the use of personal firewalls for host protection.

## Background

Older systems are often a target of unwanted attention from attackers because their existence implies some elevated level of trust or interaction with internal applications—this is most likely why they have been retained. When this theoretical value is combined with a perceived vulnerability, older systems become extremely tempting and can be seen as a natural choice for further scrutiny.

When securing older systems, you must consider the place those systems inhabit within your entire environment. By paying attention to the design and configuration of the entire network, you can create logical points in it to restrict the amount of hostile traffic as much as possible before it reaches your older systems. These measures are in addition to the system-specific hardening measures that the subsequent chapters will explore.

Traditionally, the term *perimeter network* refers to an isolated network segment at the point where a corporate network meets the Internet. Services and servers that must interact with the external, unprotected Internet are placed in the perimeter network, also known as a DMZ, demilitarized zone, and screened subnet. This is so that if attackers are able to exploit vulnerabilities in exposed services, the attackers will only be able to take one step toward accessing the trusted interior network. One way to get stronger protection for your entire network is to treat your older systems in a similar manner to the way that you treat your perimeter network— put the older systems on their own network segments and isolate them from other hosts on the network. This approach has two benefits: It lowers the risk that a compromised older system will affect the rest of the network, and it enables more aggressive filtering and blocking of network traffic to and from the older computers.

**Note:** Microsoft recommends that you never expose Microsoft® Windows NT® version 4.0 or Microsoft Windows® 98 systems directly to the Internet, even by placing them in a perimeter network. These systems should be restricted to use on your internal network.

## Network Security Considerations

You should protect the older systems in your environment as well as you treat your perimeter environment. Hardening and securing your network requires you to balance your business needs, budget restrictions, and the following security considerations, which subsequent sections cover in detail:

- Defense in depth
- Perimeter control
- Bidirectional threats
- Dissimilar services separation
- Failure planning and incident response
- Backups
- Time synchronization
- Auditing and monitoring
- Informed awareness

### Defense in Depth

To protect computer systems from today's threats, IT managers should consider a defense-in-depth strategy. A defense-in-depth strategy focuses on a combination of removing factors that increase risk and adding controls to decrease risk. No matter how good your software, hardware, processes, and personnel are, a highly determined attacker may be able to find a way through a single protective layer. The defense-in-depth security model protects key assets by using multiple layers of security throughout the environment to defend against intrusions and security threats. This multilayered approach to system security raises the effort required by an attacker to penetrate an information system, thus reducing the overall risk exposure and probability of compromise.

Rather than depending on only a strong perimeter defense or hardened servers, a defense-in-depth approach to security relies on the aggregation of multiple different defenses against a possible threat. Defense-in-depth does not reduce the need for any other security measures but instead builds on the combined strength of all of the components. Building security in overlapping layers has two key benefits:

- It makes it more difficult for attacks to be successful. The more layers you have, the harder an attacker has to work to affect a successful penetration—and the greater the chances that you will be able to detect the attack in progress.
- It helps mitigate the effect of new vulnerabilities in devices. Each layer protects against a different type of attack or provides duplicate coverage that does not suffer from the same weakness as another layer. As a result, many new attacks can be prevented by having a dependent transaction blocked by a still-intact defense measure, giving you time to address the core deficiencies.

Your business processes need to be adjusted to coordinate changes to multiple layers if they do not already permit this.

## Network Segmentation

Perimeter networks are established to create a boundary that allows separation of traffic between internal and external networks. With this boundary in place, you can categorize, quarantine, and control your network traffic. The ideal theoretical perimeter network passes no traffic to the core system except the absolute minimum required to allow desired interactions. Every additional transaction that traverses the perimeter represents another potential hole in the defense, another possible vector through which attackers may reach for control. Each new service that is enabled increases the threat surface, providing another set of code that can produce vulnerabilities and openings.

Traditional security policy calls for defining a perimeter between hosts on your network that directly communicate with the Internet and those that do not. However, you can gain additional security by treating older systems as though they were part of your perimeter, tightly controlling intercommunication between your “normal” network and those segments that contain older systems. Relegating older systems to their own network segments offers two important benefits:

- It allows you to treat the older systems as you would computers in the perimeter network. Because earlier versions of Windows do not include all of the security features and capabilities of newer versions, they are at greater risk of compromise than newer systems, and they need to be protected accordingly.
- It provides better control over the older systems. By putting these systems into their own perimeter, you segregate them from the network perimeter controls (like firewalls) that allow careful monitoring and control of traffic passing between different networks.

## Bidirectional Threats

Many attacks succeed because they provoke the target system into initiating contact outside the perimeter. This contact is engineered to be established with a hostile system, allowing it to reach back into the target. Another common scenario involves worms and viruses; after the system has been successfully compromised, the malware begins spreading from one system to another, exploiting trust relationships and disturbing external systems in an attempt to spread further. Again, these systems might also reach out to hostile systems to provide a back door for further activity. The perimeter must be designed to limit not only the traffic coming into the protected systems, but also the traffic coming *out* of the protected systems. This design makes it more difficult for your servers to be used against your environment in event of a compromise. But it also makes the attacker's job harder, because your own systems cannot be enlisted to help violate their own layers of defense.

## Dissimilar Services Separation

Concerns about performance often make it tempting to install multiple services on one computer, in order to ensure that expensive hardware is put to full use. However, doing so indiscriminately can make it difficult to properly secure your system. Carefully analyze the services and traffic that your systems host and generate, and ensure that your perimeter measures are adequate to limit traffic to the needed combinations of services and remote systems. The reverse is also true: Grouping together similar systems and services and careful partitioning of the network can make it much easier to provide protection.

## Failure Planning and Incident Response

The process of good security planning and implementation involves asking yourself, “What happens if this measure fails?” It is important to understand the consequences of mistakes, accidents, and other unforeseen events. Identifying them will allow you to design your defenses in such a manner as to mitigate those consequences and to ensure that one failure does not initiate a chain of events resulting in total exploitation of your older systems. For example, every organization should have a predefined plan that describes what to do during a virus or worm outbreak, as well as a plan that describes what to do when a compromise is suspected. In most organizations, incident response teams need to include IT staff, the legal department, and business management; these stakeholders all need to participate in carrying out a coherent response to security breaches. The Microsoft Security Guidance Kit, available at www.microsoft.com/security/guidance, includes information on how to set up and execute your own incident response plans.

---

**Note**: Microsoft describes its internal incident response process and approach in the "Incident Response: Managing Security at Microsoft" white paper, available at www.microsoft.com/technet/itsolutions/msit/security/msirsec.mspx.

---

## Backups

If an attacker is successful in entering your systems, his or her victory can be temporary if you are able to prevent him or her from successfully comprising your key resources and data. A successful backup and restore procedure helps ensure that even if the worst happens, you still have the data you need to rebuild and recover. Ensuring that your data is backed up is just the first step, however. You need to be able to quickly rebuild any compromised or affected system, and if you need to perform forensic analysis on the original hardware (often for purposes of documenting insurance claims or identifying the attack vector), you might need to have spare hardware and software available, as well as tested setup procedures.

## Time Synchronization

The various clues for spotting an attack can be scattered across multiple systems, especially in a perimeter network. Without some way of correlating this data, you might never spot them and put them together. All your systems should have the same clock time to assist in this process. The **net time** command allows workstations and servers to synchronize their time with their domain controllers. Third-party Network Time Protocol (NTP) implementations allow your servers to share time synchronization with other operating systems and network hardware, providing a unified time base across your network.

## Auditing and Monitoring

No matter how good your system defenses are, you still must audit and monitor them regularly. It is crucial that you know what your normal traffic patterns—and attacks and responses—look like. If you develop this sense, you will have indications when something negative happens, because your network traffic rhythms will change. A key area to audit and monitor is authentication. A sudden string of failed authentication attempts can often be your only warning that your system is under a brute-force dictionary attack. A brute-force dictionary attack uses known words or alphanumeric character strings to break simple passwords. Likewise, out-of-pattern authentication successes are a possible indication that your systems have been compromised on at least some level and that an attacker is attempting to leverage the initial exploit into full system access. Regular collection and archiving of event logs, combined with automated and manual analysis, makes a critical difference between failed and successful penetration attempts, in many situations. Automated tools like Microsoft Operations Manager (MOM) make it easier to monitor and analyze logged information.

## Informed Awareness

You cannot know everything, but you can stay alert and aware of the sorts of threats that other administrators are seeing. There are several excellent security resources that are dedicated to providing up-to-date information on current security threats and issues. These resources are listed in the “More Information” section at the end of this chapter.

# Network Security Design

There are several specific measures that you can use to harden your network against internal and external attacks. These measures include to:

- Segregate your older systems into their own perimeter-like network segments and protect them with access control rules, firewalls, and other techniques.
- Deploy firewalls, whether at the network level—through network hardware devices, Microsoft Internet Security and Acceleration (ISA) Server, or other products or on individual workstations.
- Harden the TCP/IP stack by applying more restrictive settings on how the stack processes anomalous packets.
- Use the port and packet filtering features built in to Windows NT 4.0 to provide additional security.

## Prerequisites

The following prerequisites are necessary for this solution:

- Possession of unused network subnets for network segmentation, of sufficient size to contain all of the necessary hosts and network overhead required.
- An understanding of TCP/IP, the ports used within your network, packet filtering, and routing.
- A thorough knowledge of the various TCP/IP options and characteristics of the other devices with which your older systems interact.

## Architecture

Several security considerations apply when designing the network architecture.

### Network Segmentation

Network segmentation controls the flow of traffic between hosts on different segments of a network. A segmented network, when properly designed, improves performance and security by ensuring that only appropriate traffic is forwarded between segments within the network. Moving from hubs to switches can minimize the capability of a hacker to sniff the network for password and other sensitive traffic, but switches do not eliminate the possibility altogether. A compromised system connected to a switch can still be used to gather information from other systems. For that reason, you should consider switches to be an answer to network collisions and performance, not network security.

Port and packet filtering, used with personal firewalls, can also help protect older systems from intrusion and compromise. However, in some situations it is not practical to install and manage firewall software on users' computers because of the administrative overhead. Unless you install firewall software that you can manage and configure remotely through a central server and database, an administrator will need to touch each computer more times after rollout to modify the configuration to address individual user needs. You also increase the likelihood that you will be required to complete additional administrative tasks for each new application you deploy, in response to the additional vulnerabilities and points of failure that the application may introduce.

Some organizations do not have the staff necessary to manage potentially hundreds of personal firewalls. In these situations, organizations can turn to network segmentation as an alternate or additional security mechanism to further protect the network. As discussed previously, network segmentation means quarantining certain servers in a perimeter network. It can also mean dividing the network into discrete segments to provide additional levels of protection for systems that reside in each segment. Segmentation can also provide considerable flexibility for traffic shaping, port monitoring and filtering, and other network management tasks, because each segment can potentially have its own discrete configuration that fits the needs of the users in the group while also suiting the security needs of the network. In effect, network segmentation brings the firewall to the workgroup level where it can be used in conjunction with the perimeter firewall(s) to further secure the network overall.

Network segmentation addresses two potential threats: those that come from outside the network, and those that come from within. The classic case of a clever employee in the Engineering department who finds a way into the Human Resources department's file server is a good example of a situation in which network segmentation would provide benefit on the local area network (LAN). The firewall sitting at the head of the Human Resources department's segment screens traffic to prevent access from computers in unauthorized departments. Likewise, the Engineering department's segment should be protected from other segments.

Network segmentation lets you structure the network into discrete security zones with the capability for unique, rule-based traffic management for each segment. You can segment the network in several ways. For example, you might choose to deploy a hardware-based firewall at each segment and physically segment the network. Or, you can deploy a single, centralized firewall with virtual LAN/segmentation capability that serves to protect individual groups. The solution that you choose ultimately depends on the network topology and the security needs of each group. At the very least you should isolate those segments that pose the most risk, such as wireless networks, from the rest of the network and impose aggressive rules to prevent unauthorized traffic to and from those segments.

To determine the best network segmentation solution for your network, start by reviewing the network structure. Then you can identify the segments that pose the greatest risk and start to build a solution. It is likely that your existing firewall vendor can offer technical information and products to help you deploy a solution.

Trey Research chose to segment the older systems in its headquarters office by putting its Windows NT 4.0 servers and Windows 98 clients on a separate network segment, and then placing a firewall between that segment and the rest of the corporate network. By enabling the use of network address translation (NAT) on that firewall, it becomes easy for Trey’s engineers to block or filter inbound and outbound traffic to that segment, giving them additional defensive capability.

## Personal Firewalls

Although the protections and measures built into the TCP/IP stack of Windows NT are a start, they have significant limitations that make them unsuitable for many deployments. They also do nothing to protect Windows 98 clients. Software firewalls, also called personal firewalls, can often provide additional protection. These specialized applications sit on top of the network stack to intercept network activity, categorize it against their configured database of permitted traffic, and allow or deny the attempt.

The big advantage that personal firewalls provide is that they can be specifically tuned to the traffic patterns of each individual computer. One possible disadvantage, however, is that because they are an application, they can be accessed and interfered with more easily, whether by accident or malicious intent.

Still, personal firewalls add an extra layer of security to older servers and clients and provide important capabilities that older systems lack, such as the ability to restrict traffic on particular ports to specified hosts and reduce the threat surface exposed by required services. Depending on where the firewall hooks into the networking stack, it can block hostile traffic before reaching vulnerabilities in the operating system or listening applications.

Personal firewalls also help limit the damage that results from Trojans, viruses, and worms. Such malware often initiates outbound traffic as well as listening to ports for illicit connections. This traffic has multiple purposes, ranging from relaying spam (both to internal and external messaging systems), to scanning other hosts and networks for vulnerabilities and openings, and it wastes disk space, processor cycles, memory, and network bandwidth. Such applications and their connection initiation often cause Denial-of-Service (DoS) periods as a secondary effect in addition to their infection and clean-up problems.

In addition to expanded ingress and egress filtering capabilities, many personal firewalls can also allow or deny network access based on the executable that requests it. This functionality can be used to block specific applications—or only allow pre-approved applications—from ever accessing the network, regardless of what ports and protocols they use. This prevents users from circumventing security configurations with protocol-agile applications such as peer-to-peer file sharing, instant messaging, or other applications that tunnel their connections through Hypertext Transfer Protocol (HTTP).

Personal firewalls can require a bit of time to set up and configure properly because they scrutinize literally every transaction, at least until a transaction is properly categorized and classified as allowed or denied. They require detailed knowledge of every bit of network traffic generated by authorized applications and services, because a single blocked interaction could subtly cripple the execution of necessary programs. They also require maintenance, patching, and updating. Look for personal firewalls that have features intended for corporate and enterprise use, such as the ability to centrally manage and maintain configuration databases.

The ISA Firewall Client, used in conjunction with ISA Server, provides more sophisticated filtering and policy enforcement capabilities; it allows traffic control based on the user's identity, as well as the origin or destination of the traffic. All traffic can be monitored and controlled through strategically located servers, and the enterprise-wide policies can be updated easily and quickly.

The Routing and Remote Access Service download for Windows NT is another option. Although RRAS mainly provides support for dynamic routing protocols such as the Routing Information Protocol (RIP) and dial-up and virtual private networking (VPN) capabilities, it also gives the ability to define access control on incoming network traffic above and beyond the basic port filtering built in to Windows NT. RRAS is a free download, available from the Routing and Remote Access Service Download page at www.microsoft.com/ntserver/nts/downloads/winfeatures/rras/rrasdown.asp.

Using RRAS or any add-on solution has one significant disadvantage, however. Add-ons are software subsystems, running as Windows services, and as such are started after the network interfaces and protocols are initialized. Thus, there is a window of opportunity during which hostile traffic can slip through during a reboot or service outage. The use of personal firewalls, the ISA Server Firewall Client, or RRAS without other forms and layers of protection will not provide absolute security. As always, multiple defenses, used to offset and counteract weaknesses in each layer, provide the most superior protection.

Trey Research is already using the Internet Connection Firewall (ICF), part of Windows XP, on some of its systems. Trey chose to purchase a license for, and deploy, a third-party personal firewall product for its computers running Windows 98 and Windows NT Workstation 4.0. This provides both ingress and egress filtering for internal traffic on the company's segregated network, and it provides centralized data collection that helps Trey maintain visibility into network traffic type and volume.

## Windows NT Port and Packet Filters

One of the best methods of protecting networked computers is to limit what types of network traffic they receive and process, which usually requires some sort of packet filter. Most administrators think of routers and network chokepoints when they design packet-filtering strategies. However, Windows NT comes with a basic packet filtering capability, known as TCP/IP security. Although this facility does not provide sufficient protection by itself, it does make an excellent secondary layer of defense when used in conjunction with stateful packet-filtering devices.

The main advantage of the Windows NT built-in TCP/IP security features is that they are implemented within the TCP/IP network stack as an integral part of the protocol drivers. The benefit of this depth of integration is that all of the settings are always active as long as the protected interfaces are active; there is never any window of time, such as start-up, during which network traffic is not being filtered. TCP/IP security is transparent to applications, although it can interfere with some personal firewall software.

Despite its name, TCP/IP security allows port-by-port filtering of TCP and User Datagram Protocol (UDP), as well as other IP protocols. Active filters block inbound traffic but permit outbound traffic and responses to TCP connections initiated by the local host. There are some limitations, however:

- Port filters either allow or block traffic from *all* hosts. You cannot establish any finer level of granular control, as is possible with the IP Security (IPsec) Extensions capability built into Windows 2000, Windows XP, and Microsoft Windows Server™ 2003.
- The filters are not truly stateful and cannot be linked dynamically to allow traffic for secondary connections. The IPsec implementation in later versions of Windows provides for allowing secondary connections, as do most hardware firewalls.

In addition to the simple port filtering capability just described, the Windows NT TCP/IP stack provides many tunable parameters that are especially interesting for closing or mitigating threatening network traffic. Over time, a variety of attacks have been developed that exploit flaws in the Windows NT 4.0 TCP/IP networking code; even though all of these flaws have been addressed by service packs and security updates, it may still be valuable to apply these changes to give your network additional protection. These adjustments usually require direct editing of the registry by using the **regedit32** or **regedit** tools. Microsoft Knowledge Base (KB) article 120642, “TCP/IP and NBT Configuration Parameters for Windows 2000 or Windows NT” at http://support.microsoft.com/?kbid=120642 provides a long list of the tunable parameters; the ones that are most pertinent to network and system security are discussed in the following sections.

Trey Research has defined port and packet filtering on its older hosts to disallow traffic from ports not listed in KB article 150543, "Windows NT, Terminal Server, and Microsoft Exchange Services Use TCP/IP Ports" at http://support.microsoft.com/?kbid=150543 or used by other applications on its network from leaving their internal network and traveling to the Internet. This prevents sensitive network traffic from being broadcast on an uncontrolled network.

### SYN Flooding Protection

SYN (synchronization) flooding is a common vector of attacks against TCP services. When a TCP client initiates a new connection, it sends an empty TCP packet to the listening server with the SYN flag set, indicating that it is requesting a new connection. The server sends back a response packet with both the SYN and ACK (Acknowledgment) flags. The client then responds with an ACK packet, completing the three-way handshake, and the connection is then open.

Until the final acknowledgment is received, the TCP/IP driver assigns this connection the SYN_RCVD (SYN Received) state. If for some reason Windows does not receive a response to the SYN+ACK packet, it will wait, by default, for one second and then retransmit the SYN+ACK packet. A second retransmission will occur after another timeout of three seconds, with a final retransmission occurring after another six seconds. Each connection request requires the server to allocate a certain amount of memory and other kernel structures; a flood of incoming requests can rapidly exhaust resources and cause a DoS. The applicable registry keys are the following:

**HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\SynAttackProtect** (REG_DWORD):

- **0** = Disabled (default)
- **1** = Delay creation of the route cache entry until connection is established
- **2** = Delay notifying the Winsock driver until the three-way handshake is complete (recommended)

**HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\TcpMaxHalfOpen** (REG_DWORD):

- This key defines the maximum number of connections that the IP stack will allow to be in the SYN_RCVD state before triggering the SynAttack state. The default value is **100**.

**HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\TCPMaxHalfOpenRetried** (REG_DWORD):

- This key defines the maximum number of connections that are both in the SYN_RCVD state and have been retransmitted more than once before SynAttack triggered. The default value is **80**.

**HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\TCPMaxPortsExhausted** (REG_DWORD):

- This key defines the number of refused connect requests caused by having no backlog before the SynAttack is triggered. The default value is **5**.

**HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\TCPMaxConnectResponseRetransmissions** (REG_DWORD):

- This key defines the number of retransmission attempts during the SYN_RCDV state. The default value is **3**, which causes all three retransmissions. Setting this value to **1** will cause only one retransmission attempt.

The KB article 146241, "Internet Server Unavailable Because of Malicious SYN Attacks" at http://support.microsoft.com/?kbid=142641 discusses these settings in more detail.

On older hosts, Trey Research has tightened the settings of all of the registry keys listed previously, reducing the chance that a SYN attack will negatively impact those systems. Trey Research did not use the most restrictive settings, because there is a legitimate need for regular SYN traffic to traverse the network, and it does not want to interrupt the traversal of normal network traffic. Trey's system administrators understand that the changes they made to the default settings may need to be revisited and are proactively monitoring the amount of SYN traffic on their network.

### Backlog Size Control

Applications that make use of TCP/IP use the Winsock application programming interface (API), which is provided by **afd.sys**, the Winsock kernel mode driver. The Winsock API provides the mechanisms for applications to open client connections and establish listeners on ports for server connections. The **listen()** function is used to tell Winsock to listen for connection attempts to a specific port and forward them to the application.

One parameter that this function must specify is the backlog size, which is a queue that holds pending incoming connections until the stack can handle them. The backlog provides a maximum queue length and by default is set to **200** on Windows NT Server and to **5** on Windows NT Workstation. Windows NT 4.0 Service Pack 2 (SP2) introduces the dynamic backlog feature, which permits the TCP/IP stack to adjust the size of the backlog queue as needed to respond to current network conditions.

This feature that adjusts the backlog queue is turned off by default and must be enabled by using the following registry entries. In addition, the calling application must request a backlog queue larger than the **MinimumDynamicBacklog** parameter in order to benefit from this feature. The applicable registry keys are the following:

**HKLM\SYSTEM\CurrentControlSet\Services\AFD\Parameters\EnableDynamicBacklog** (REG_DWORD):

- **0** = Disabled (default)
- **1** = Enabled (recommended)

**HKLM\SYSTEM\CurrentControlSet\Services\AFD\Parameters\MinimumDynamicBacklog** (REG_DWORD):

- This key defines the minimum number of entries in the backlog queue; if the number of available entries drops below this minimum, you need to create more entries. **20** is the recommended value.

**HKLM\SYSTEM\CurrentControlSet\Services\AFD\Parameters\MaximumDynamicBacklog** (REG_DWORD):

- This key defines the maximum number of entries that can be created in the backlog queue. Setting this number higher than **5,000** per 32 megabytes (MB) of system random access memory (RAM) can lead to memory exhaustion under attack. Remember that a separate backlog queue is created for each network service. Because Trey Research’s target servers have 512 MB of system RAM, the formula used to determine the upper limit of this registry value is:
- (512/32)*5000=144000.
- For workstations (all of which have 256 MB of RAM), the calculated value is (256/32)*5000 = 72000.

**HKLM\SYSTEM\CurrentControlSet\Services\AFD\Parameters\DynamicBacklogGrowth-Delta** (REG_DWORD):

- This key defines the number of new connections to add to backlog at one time when more are required. **10** is the recommended value.

The KB article 146241, "Internet Server Unavailable Because of Malicious SYN Attacks" at http://support.microsoft.com/?kbid=142641 discusses these settings in more detail.

Trey Research elected to set the recommended values for all registry keys listed previously in order to control the number of incoming connections so that the company's systems do not become overwhelmed by incoming requests.

### TCP Keep-Alive Timers

TCP keep-alive timers are an advanced TCP feature that keeps idle connections alive. This function becomes especially important when those connections pass through firewalls and Network Address Translation (NAT) devices, which regularly purge aged entries from their connection and masquerading tables.

Windows NT provides the capability to enable and define TCP keep-alive timeouts. Lowering this value helps prevent dead connections from keeping resources in use for too long. To set this value, use the following registry keys:

**HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\KeepAliveTime** (REG_DWORD):

- This key defines the number of milliseconds between keep-alive checks; by default this value is **7,200,000** (two hours).

**HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\KeepAliveInterval** (REG_DWORD):

- This key defines the number of milliseconds between keep-alive packet retransmissions if a response is not received; by default this value is **1000** (one second).

Trey Research did not alter its TCP keep-alive timer settings, because they have not, to date, experienced any problems with the amount of time that connections are kept in use.

### Path Maximum Transmission Unit Discovery

Path Maximum Transmission Unit (MTU) discovery is a feature that allows Windows to automatically discover the largest supported packet size on all network segments between hosts. It works by sending out large packets that have the "do not fragment" bit set. When an intervening router cannot forward this packet because it is larger than the MTU for the segment, it returns an Internet Control Message Protocol (ICMP) message. Windows then reduces the packet size and tries again until the packet is sent through to the destination.

By setting the proper MTU for remote hosts, Windows avoids generating fragmented packets, which reduce performance and increase the chance of lost data and retransmissions. Fragmented packets can be a security risk as well; fragmented packet handlers are ripe source of buffer overflows, and the ability to filter out fragments at the border can reduce a lot of attacks. Path MTU Discovery should be left enabled but will require ICMP type 3, code 4 to be routed through the firewalls. Use the following registry keys to set this value:

**HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\EnablePathMTUDiscovery** (REG_DWORD):

- **1** = Enabled (default, recommended)
- **0** = Disabled

A value of 0 sets the MTU size to 576 bytes for all traffic outside of configured local subnets. Additionally, with this setting, Windows will not honor requests to change the MTU.

**HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\EnablePathMTUBHDetect** (REG_DWORD):

- **0** = Disabled (default)
- **1** = Enabled

A value of 1 allows Windows to attempt detection of black hole routers during Path MTU Discovery; these routers silently discard packets with the “do not fragment” flag set if they are too large, instead of sending the correct ICMP reply. Enabling this value will cause more retransmission attempts.

Trey Research did not alter the default values for path MTU discovery because disabling this feature might result in some remote systems becoming unreachable. Trey Research does not want to interrupt business transactions in cases where systems within the communication path cannot support reducing the MTU size.

### Source Routing

Source routing allows applications to override the routing tables and specify one or more intermediate destinations for outgoing datagrams. Although this capability is marginally useful for troubleshooting, it is extremely unwise to use it on modern production networks. Successful attackers can use this feature to transparently direct all network traffic to a centralized collection point for packet capture. Disable source routing by using the following registry key; additionally, ensure that your border routers are configured to drop all IP datagrams with the source routing option set:

**HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\DisableIPSourceRouting** (REG_DWORD):

- **0** = Enabled (default)
- **1** = Disabled when IP forwarding is enabled
- **2** = Disabled completely (recommended)

Trey Research disabled source routing on all hosts within its control to help prevent the capture of network data by attackers.

### Dead Gateway Detection

Dead Gateway Detection allows Windows to detect when a default gateway appears to have stopped responding and fail over to additional configured default gateways. In practice, this capability is rarely used and provides an opportunity for DoS attacks. Use the following registry key to control this function:

**HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\EnableDeadGWDetect** (REG_DWORD):

- **0** = Disabled (recommended)
- **1** = Enabled (default)

Trey Research enabled dead gateway detection on all hosts to help alleviate the possibility of DoS attacks on its network.

### Router Discovery

Router discovery uses the ICMP Router Discovery messages to locate and configure default routes and gateways. Again, attackers can use this ability to redirect network traffic for various purposes, including sniffing and man-in-the-middle attacks. Routers that support this feature must send the ICMP Internet Router Discovery Protocol (IRDP) message, which should be disabled. In additional, Dynamic Host Configuration Protocol (DHCP) must be configured with the appropriate option so that the interface will accept the IRDP messages.

The following registry key should be edited to prevent malicious configuration changes; if the capability to use multiple routes is necessary, consider deploying the RRAS add-on and use a securable routing protocol.

**HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\PerformRouterDiscovery** (REG_DWORD):

- **0** = Disabled (recommended)
- **1** = Enabled

Trey Research disabled router discovery on all hosts to prevent unapproved configuration changes within its network.

### ICMP Redirects

ICMP redirects are another potential vulnerability because they allow an arbitrary sender to forge packets and alter the victim's routing table. This feature is enabled by default and should be disabled on the following key to avoid malicious effects:

**HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\EnableICMPRedirect** (REG_DWORD):

- **0** = Disabled (recommended)
- **1** = Enabled (default)

Trey Research disabled ICMP redirects on all hosts to protect the routing table on each host.

### RPC Port Closures

Many services by default are listening on network interfaces, including the local loopback interface. For example, the Microsoft remote procedure call (RPC) port mapper listens on TCP/135, UDP/135, TCP/1027, and TCP/1028. Three of these services—the RPC client, RPC server, and RPC end-point mapper—can be configured to close all open ports. However, these changes must be carefully tested because they can break functionality, not only with remote hosts, but also between local services.

Microsoft Exchange and Microsoft SQL Server™ are the most commonly deployed applications that require RPC. Additionally, RPC calls are used during remote management of servers. The common built-in utilities that are dependent on RPC services are:

- DHCP Manager
- DNS Administrator
- WINS Manager
- Performance Monitor
- Event Viewer
- Registry Editor

- Server Manager
- User Manager

To determine whether RPC is being used by an application, install Network Monitor Tools and Agent from the Windows NT 4.0 CD and use it to assess whether the application is using RPCs on the ports specified previously.

Remove the following registry keys to disable the RPC client:

**HKLM\SOFTWARE\Microsoft\RPC\ClientProtocols\ncacn_ip_tcp**

**HKLM\SOFTWARE\Microsoft\RPC\ClientProtocols\ncacn_ip_udp**

Remove the following registry keys to disable the RPC server:

**HKLM\SOFTWARE\Microsoft\RPC\ServerProtocols\ncacn_ip_tcp**

**HKLM\SOFTWARE\Microsoft\RPC\ServerProtocols\ncacn_ip_udp**

The RPC end-point mapper (rpcss.exe) opens multiple ports if the RPC server is enabled, but it can alternatively be configured to reject non-local Distributed Component Object Model (DCOM) connection attempts. Be aware that this will have a significant impact on your ability to remotely manage your systems; the Directory Services client extension described in Chapter 4, "Hardening Microsoft Windows NT 4.0," will be particularly affected, because it relies heavily on the Windows Management Instrumentation (WMI) providers, which are DCOM objects. You can disable non-local DCOM connections by using the following registry key:

**HKLM\SOFTWARE\Microsoft\OLE\EnableDCOM** (REG_SZ):

- Set the value to “**N**” to disable remote DCOM connections.

Trey Research did not alter the settings for RPC ports on its internal hosts, because they need these ports active in order to use numerous network services, and closing RPC hosts would have left several applications inaccessible to users.

# Implementation

## Prerequisites

For these implementation details to work correctly, you must have the basic Trey Research infrastructure implemented as introduced in Chapter 2, "Applying the Security Risk Management Discipline to the Trey Research Scenario."

## Overview

Implementing this solution scenario will involve performing the following activities:

- Configuring native Windows NT 4.0 port filtering.
- Configuring Windows NT 4.0 IP tuning parameters.

Registry files for most of these settings can be found in the Tools and Templates that are included as part of this guidance.

## Configuring Native Windows NT 4.0 Port Filtering

Prior to enabling TCP/IP port filtering on your Windows NT servers, you must know exactly *which* traffic is necessary for the proper operation of that server in its given role. Microsoft provides a base list of well-known ports that Windows NT 4.0 uses; it is available in "Port Reference for MS TCP/IP" at www.microsoft.com/technet/prodtechnol/winntas/support/port_nts.mspx.

Other applications should have their own documentation available. For Windows Server systems, you can consult KB article 832017, "Port Requirements for the Microsoft Windows Server System" at http://support.microsoft.com/?kbid=832017. For other applications not included in these lists, either consult the application vendor's documentation or use network traffic monitoring tools such as Netmon.exe (which comes as a component of Windows NT 4.0), Netstat (a built-in Windows NT utility that shows which ports are currently in use), or TCPView (a free tool available from Sysinternals at www.sysinternals.com) to verify the ports used by the application.

Configuring Windows NT port filtering is a simple process using the following procedure.

▶ **To configure native Windows NT port filtering**

1. To open the Network Control Panel click **Start**, select **Settings**, click **Control Panel**, and then double-click **Network**.
2. On the **Protocols** tab, select **TCP/IP Protocol**, and then click **Properties**.
3. On the **IP Address** tab, click .**Advanced**.
4. In the **Advanced IP Addressing** dialog box, select the **Enable Security** check box.
5. Click Configure.
6. In the **TCP/IP Security** dialog box, pick the relevant adapter (if the server is multi-homed).
7. By default, no filters are defined, and all TCP, UDP, and IP traffic is permitted (see Figure 3.1). To enable filters, select the **Permit Only** check box on TCP, UDP, or IP Protocols, and add the port and protocol numbers on which you want to allow traffic on. Be sure to allow for basic infrastructure services.

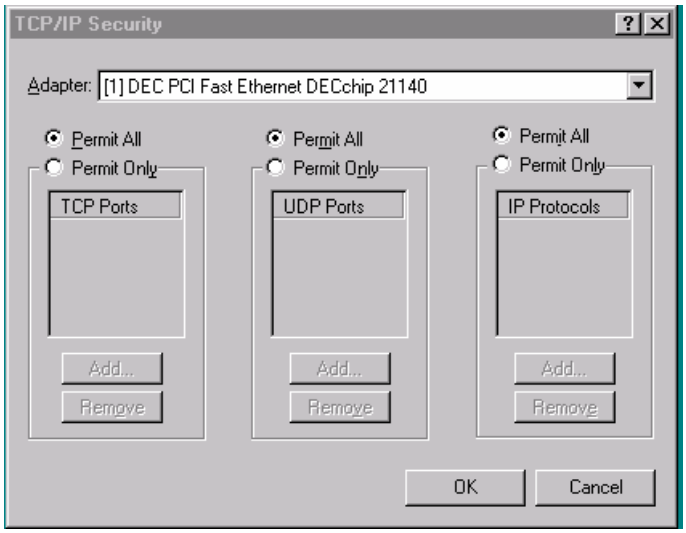


**Figure 3.1**
*Configuring native Windows NT port filtering*

## Configuring Windows NT 4.0 IP Tuning Parameters

Configuring the Windows NT IP tuning parameters is a matter of editing the following registry parameters.

---

**Caution**: Changing the parameters used to tune the TCP/IP stack changes the way that the network subsystem communicates with routers, switches, and other computers on the network. These changes may impact the performance or stability of production applications. Before making any changes to the production environment, you should thoroughly test your proposed changes in a lab environment that replicates the behavior and configuration of your production servers and clients.

---

▶ **To configure the Windows NT IP tuning parameters for servers**

1. Run Registry Editor (Regedt32.exe).
2. Ensure that the following registry entries (TCP_Params.reg) are applied to the **HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters** key:
   - SynAttackProtect (REG_DWORD) = 2
   - TcpMaxHalfOpen (REG_DWORD) = 100
   - TCPMaxHalfOpenRetried (REG_DWORD) = 80
   - TCPMaxPortsExhausted (REG_DWORD) = 5
   - TCPMaxConnectResponseRetransmissions (REG_DWORD) = 1
   - KeepAliveTime (REG_DWORD) = 7200000
   - KeepAliveInterval (REG_DWORD) = 1000
   - EnablePathMTUDiscovery (REG_DWORD) = 1
   - EnablePathMTUBHDetect (REG_DWORD) = 0
   - DisableIPSourceRouting (REG_DWORD) = 2
   - EnableDeadGWDetect (REG_DWORD) = 0
   - PerformRouterDiscovery (REG_DWORD) = 0
   - EnableICMPRedirect (REG_DWORD) = 0

3. Ensure that the following registry entries (AFD_Params.reg) are applied to the **HKLM\SYSTEM\CurrentControlSet\Services\AFD\Parameters** key:
   - EnableDynamicBacklog (REG_DWORD) = 1
   - MinimumDynamicBacklog (REG_DWORD) = 20
   - MaximumDynamicBacklog (REG_DWORD) = 144000
   - DynamicBacklogGrowth-Delta (REG_DWORD) = 10
4. Ensure that the following registry entries (DisableDCOM.reg) are applied to the **HKLM\SOFTWARE\Microsoft\OLE** key:
   - EnableDCOM (REG_SZ) = N

---

**Caution**: Disabling DCOM will disable several basic remote management capabilities and other applications, such as the Directory Services client extensions, WMI, the ability to remotely manage print and file shares, and more. Carefully test this change in your lab before you apply it to production servers.

---

5. If you are certain that no RPC-based services, applications, or utilities are in use, delete the following subkeys:
   - HKLM\SOFTWARE\Microsoft\RPC\ClientProtocols\ncacn_ip_tcp
   - HKLM\SOFTWARE\Microsoft\RPC\ClientProtocols\ncacn_ip_udp
   - HKLM\SOFTWARE\Microsoft\RPC\ServerProtocols\ncacn_ip_tcp
   - HKLM\SOFTWARE\Microsoft\RPC\ServerProtocols\ncacn_ip_udp
6. Exit Registry Editor.

▶ **To configure the Windows NT IP tuning parameters for workstations**

1. Run Registry Editor (Regedt32.exe).
2. Ensure that the following registry entries (NT4WS_TCP_Params.reg) are applied to the **HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters** key:
   - SynAttackProtect (REG_DWORD) = 2
   - TcpMaxHalfOpen (REG_DWORD) = 100
   - TCPMaxHalfOpenRetried (REG_DWORD) = 80
   - TCPMaxPortsExhausted (REG_DWORD) = 5
   - TCPMaxConnectResponseRetransmissions (REG_DWORD) = 1
   - KeepAliveTime (REG_DWORD) = 3600000
   - KeepAliveInterval (REG_DWORD) = 1000
   - EnablePathMTUDiscovery (REG_DWORD) = 1
   - EnablePathMTUBHDetect (REG_DWORD) = 0
   - DisableIPSourceRouting (REG_DWORD) = 2
   - EnableDeadGWDetect (REG_DWORD) = 0
   - PerformRouterDiscovery (REG_DWORD) = 0
   - EnableICMPRedirect (REG_DWORD) = 0

3. Ensure that the following registry entries (NT4WS_AFD_Params.reg) are applied to the **HKLM\SYSTEM\CurrentControlSet\Services\AFD\Parameters** key:
   - EnableDynamicBacklog (REG_DWORD) = 1
   - MinimumDynamicBacklog (REG_DWORD) = 20
   - MaximumDynamicBacklog (REG_DWORD) = 72000
   - DynamicBacklogGrowth-Delta (REG_DWORD) = 10
4. Exit Registry Editor.

# Summary

Extending the concept of the perimeter network is a useful measure for securing older systems, one that offers additional protection within the overall design and configuration of the entire network. These natural chokepoints, combined with the traffic control provided by personal firewall software and the native port filtering and IP tuning capabilities of Windows NT 4.0, give you multiple layers of protection that enhance the specific hardening instructions that subsequent chapters in this guidance offer.

## More Information

- The NTBUGTRAQ mailing list is an excellent resource devoted to the discussion of current Windows security issues. The list moderator, Russ Cooper, is not affiliated with Microsoft but is extremely knowledgeable. To subscribe to the list or read the archives, see www.ntbugtraq.com/.
- The CERT Coordination Center (CERT/CC) is a central clearing house for security advisories and bulletins. More information is available at www.cert.org/.
- The SANS Institute offers a variety of training, certification, and research focused on computer security issues. More information is available at www.sans.org/.
- The US Computer Emergency Readiness Team (US-CERT) focuses on computer security advisories and threats for the United States, but is a good source of information on current threats. More information is available at www.us-cert.gov/.
- The National Security Agency offers a variety of secure configuration guides covering multiple operating systems and applications at www.nsa.gov/snac/.
- RRAS for Windows NT 4.0 is available on the Routing and Remote Access Service Download page at www.microsoft.com/ntserver/nts/downloads/winfeatures/rras/rrasdown.asp.

# 4

# Hardening Microsoft Windows NT 4.0

Microsoft® Windows NT® version 4.0 lacks some of the enhanced security functionality that later versions of the Microsoft Windows® operating system have. Still, it has a number of extremely useful features and techniques that significantly increase the security of your systems. This chapter covers these topics in detail and explains how to use these features to harden your Windows NT 4.0 systems, whether they are workstations or servers.

These topics include information on how to:

- Install the initial operating system and patch baseline.
- Harden the boot sequence.
- Install the Directory Services Client add-in.
- Use system policies and the Security Configuration Manager.
- Choose the Windows NT LAN Manager (NTLM) authentication level.
- Define effective password policies.
- Use account and password lockouts.
- Harden the file system.
- Harden services.
- Perform other hardening measures.

# Windows NT Host Security Design

Hardening the Windows NT 4.0 operating system requires assessment and understanding of multiple topics.

## Installing the Initial Operating System and Patch Baseline

The first critical step in hardening Windows NT-based systems is to install the latest set of security patches for the base operating system. Ideally you should have performed this step as part of your normal patch management process, as described in Chapter 6, “Patch Management.” This process should begin with a complete inventory to determine which computers are at which revision levels. You can use the Microsoft Baseline Security Analyzer (MBSA) or Microsoft Systems Management Server (SMS) to scan Windows NT 4.0 systems, either locally or remotely. After you complete the inventory, you should inspect inventory results to ensure that each Windows NT system has the necessary set of baseline updates installed. These updates should include:

- Service Pack 6a (SP6a) for Windows NT 4.0, which is the most recent service pack for Windows NT 4.0; it contains a complete and regression-tested set of fixes that are designed to be installed as a unit.
- The Post-Windows NT 4.0 Service Pack 6a Security Rollup Package (SRP). The SRP incorporates several patches that were released after the original SP6a release and is a prerequisite for follow-up patches from Windows Update.
- An up-to-date set of security patches. Microsoft has not released an integrated SRP or service pack for Windows NT 4.0 since the release of the SRP, but individual components such as portions of the operating system, Microsoft Internet Explorer, and add-on utilities such as the Routing and Remote Access Service (RRAS) and Internet Information Services (IIS) have been updated. You can obtain a current list of patches from Microsoft TechNet, or manually review the list of patches available through Windows Update.
- An updated version of Internet Explorer. Internet Explorer 6.0 Service Pack 1 (SP1) is the current version; it incorporates hundreds of security fixes and improvements released since the versions of Internet Explorer available with released versions of Windows NT. Depending on your environment, you may choose to download Internet Explorer directly from the Microsoft Web site, order it on a CD–ROM, or install it on multiple computers by using the Internet Explorer Administration Kit.

## Hardening the Boot Sequence

After the system has a correct and complete minimum patch set, you must next increase the degree of protection available to the system at boot time. This protection takes two forms: zeroing the boot menu timeout to make it more difficult for an attacker to boot into an alternate operating system, and using the built-in Syskey utility to encrypt information stored in the Security Account Manager (SAM) database.

Windows NT Server stores user account information, including a derivative of the user account password, in a secure portion of the registry protected by access control and an obfuscation function. The account information in the registry is accessible only by members of the Administrators group. Windows NT Server, like other operating systems, allows administrators access to all resources in the system. For installations that want enhanced security, strong encryption of account password derivative information provides an additional level of security to prevent administrators from intentionally or unintentionally accessing password derivatives by using registry programming interfaces. Syskey also protects SAM data from various kinds of offline attacks mounted by booting into an alternate operating system and accessing the SAM files.

Syskey generates a random encryption key that is used to encrypt the SAM data. After the SAM data has been encrypted, the key must be loaded at boot time and used to decrypt the in-memory copy of the SAM data. Syskey operates in three modes, as shown in Figure 4.1:

- In mode 1 (**Store Startup Key Locally**), the encryption key is super encrypted and stored on the local computer. At boot time, the key is decrypted and loaded, allowing the computer to be restarted without administrator intervention.
- In mode 2 (the **Password Startup** button and associated text fields), the encryption key is super encrypted with an administrator-specified pass phrase. At boot time, the administrator must type the pass phrase into the console to complete the boot. The system will not boot until the pass phrase is entered.
- In mode 3 (the **Store Startup Key on Floppy Disk** button), the random Syskey key is stored on a floppy disk, which must be inserted at a prompt during the boot cycle. This floppy disk must be stored and maintained securely.

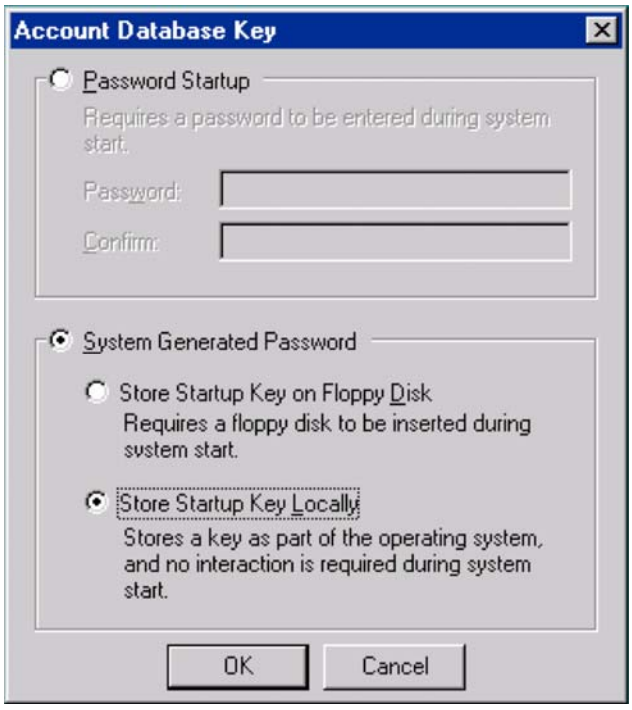


**Figure 4.1**
*Syskey's mode selection dialog box*

Trey chose to implement Syskey protection on all of its computers running Windows NT. (Syskey is already enabled by default in Windows 2000, Microsoft Windows Server™ 2003, and Windows XP.). Most of Trey's servers, and all of its workstations, are set to use Syskey mode 1; Trey chose this mode for widespread use because it offers a reasonable balance of security and convenience. For high-value servers, Trey is using mode 2 even though it requires an administrator to visit the computer and enter the password each time it is rebooted. Because of the requirement to securely store and maintain a boot floppy disk for each protected computer, Trey elected not to use mode 3 on any of its systems.

## Installing the Directory Services Client Add-In

Windows NT 4.0 systems can only fully participate in Windows NT domains; if there are Microsoft Active Directory® directory service domains, Windows 98 and Windows NT can natively participate in them only through the network basic input/output system (NetBIOS) interface. Windows 2000 and Windows Server 2003 provide NetBIOS compatibility by emulating a Windows NT primary domain controller (PDC). Even though Active Directory provides transitive two-ways trusts, Windows NT 4.0 systems can take advantage only of explicit one-way trusts, whether the system is using an Active Directory domain controller or a Windows NT BDC.

Instead of maintaining Windows NT-style domains, it is possible for computers running Windows NT and Windows 98 to participate in Active Directory domains by adding the Directory Service Client (DSClient) add-in. DSClient enables these systems to participate in Active Directory domains. With this software, these systems have use of many of the Active Directory features, such as the ability to use transitive trust relationships within the forest. Transitive trust relationships allow authorized users to access appropriate resources in any domain in the forest. The DSClient 2003 update runs on Windows NT 4.0 with SP6a and Internet Explorer 6 SP1; the DSClient version supplied with Windows 2000 also requires SP6, but can use Internet Explorer 4.01 or later. (Requirements for using the DSClient with Windows 98 are described in Chapter 5, “Hardening Microsoft Windows 98.”)

Once installed, DSClient does not provide support for all Active Directory features; in particular, it does not support Kerberos authentication, Group Policy application, or the use of service or user principal names (UPNs) for authentication. However, it does enable support for the following Active Directory features:

- Active Directory Service Interfaces (ADSI). ADSI provides a common programming API for Active Directory-aware scripts and programs. With ADSI, it is possible to script a large number of directory operations that would otherwise be unscriptable under Windows NT.
- Distributed file system (DFS) fault tolerance client. DFS fault-tolerant and fail-over file shares provide Active Directory-integrated, distributed file share resources.
- Active Directory Windows Address Book (WAB) property pages. User object pages, accessed through the **Search** menu, provide the capability for authorized users to update properties (such as addresses and phone numbers) on user objects.
- NTLM version 2 authentication. NTLM provides improved authentication features and gives the best level of authentication security other than Kerberos.

- Site awareness. Systems using the DSClient are aware of Active Directory sites and will use a domain controller in their local site, even for password change operations, if Domain Name System (DNS) is properly configured and sites are properly registered in Active Directory. See Microsoft Knowledge Base (KB) article 249841, "How Windows 98 Active Directory Client Extension uses Active Directory site information" at http://support.microsoft.com/?kbid=249841 for more details on how DSClient affects logon and password change behavior.
- Search for objects in Active Directory. Users can find printers and users in Active Directory from the **Search** menu.
- Reduced dependency on the PDC. Clients may connect to any domain controller in the domain for password changes.

## Using System Policies and the Security Configuration Manager

System policies are a specific configuration management measure that Microsoft first made available with Windows 95 and Windows NT 4.0. By themselves, they do not increase the security of the system; they will, however, allow you to restrict access to specified resources in a way that enhances your other security measures, making it more likely that users will be unable to circumvent your policies and restrictions either through intent or ignorance. When these policies are properly designed and applied, they are a critical part of ensuring the integrity of your older systems. In fact, they are so useful that Microsoft expanded upon them for the Group Policy Object (GPO) functionality in Windows 2000 and later.

A system policy is a set of one or more restrictions that are applied to the **HKEY_CURRENT_USER** and **HKEY_LOCAL_MACHINE** registry hives. When the user logs on to a computer, a successive series of policies are checked for and (if found) applied, based on the user name, global group memberships in the domain, and computer-specific policies. System policies are intended to be used in a Windows NT domain implementation and complement other domain-aware features such as user profiles.

It is important to understand the difference between system policies and user profiles. User profiles are the collection of user-specific configuration directories, files, shortcuts, and registry hives loaded under **HKEY_CURRENT_USER** (Ntuser.dat for Windows NT, User.dat for Windows 98); they control the appearance of the desktop, browser favorites and bookmarks, and other user-configurable options in both the operating system and any applications. Conversely, system policies are an extra set of registry entries that are applied to the computer after the active user profile has been located and loaded; they specify particular features of the operating system to which the current user should not have access.

### Working with System Policies

Microsoft provides a detailed white paper, "Guide to Microsoft Windows NT 4.0 Profiles and Policies," which explains user profiles and system policies, how they are used, and how they are used together. (See the "More Information" section for the white paper's location).

System policies provide you with many controls over earlier versions of the desktop and user interface. If you are familiar with Group Policies in Windows 2000, you will be familiar with the basic capabilities of system policies. In general, though, they allow you to:

- Specify a banner or disclaimer that will be displayed before users are permitted to log on, such as one advising them of any applicable legal usage policies.
- Prevent the logon dialog box from displaying security-sensitive information such as the user name of the previous user or the shutdown button.
- Specify the behavior of logon scripts, cached roaming profiles, slow network detection, the Start banner, and Welcome Tips.
- Specify the location of various shared folders on the **Start** menu, including the Startup folder, allowing you to ensure that particular applications are always launched at logon.
- Restrict the appearance of the desktop: backgrounds, icons, colors, and available commands on the **Start** menu.
- Specify the behavior of various file system-related features, such as shell extensions, read-only file access time updates, and long file names.
- Restrict access to network resources, available drive letters, and the ability to map drives in the Windows Explorer. See KB article 156698, "Disabling Access to Network Resources Using System Policies" at http://support.microsoft.com/?kbid=156698 and article 220955, "Using System Policies to Hide Specific Drive Letters" at http://support.microsoft.com/?kbid=220955 for more information.
- Restrict the creation of hidden file shares.
- Specify and restrict printer, Simple Network Management Protocol (SNMP), and RRAS settings.
- Restrict the ability of the shell to launch particular executables.
- Restrict the ability to launch registry editors and the command prompt.

System policies have many similarities to Group Policies, but there are a few limitations and differences that you need to keep in mind:

- System policies are not hierarchical. Given the flat nature of the Windows NT domain model, you do not have the same flexibility to define overlapping, complimentary system policies that you do with Group Policies under Active Directory. You can define policies for individual users, default users, groups, individual computers, and default computers. A subsequent section in this chapter describes the method used to apply them.
- System policies make their changes directly to the appropriate registry hive. They stay in force until something causes them to be changed, such as an alternate policy. Group Policies in Active Directory make their changes to a special part of the registry, which Windows then uses to override the normal registry entries. This system policy practice of making direct writes to the registry is known as "tattooing the registry" and means that system policies cannot make any assumptions about the state of any setting that is being managed.

System policies are created by using the System Policy Editor (SPE) utility and administrative template files (.adm). These template files provide the SPE categories and subcategories, registry keys and values that control the specific settings, as well as any options, restrictions, and default values that may apply. You may create your own custom template files and add in specific registry settings that are not covered by the default templates supplied with the SPE.

After SPE is run and you create a group of settings that are then associated with a specific user, group, computer, or default user or computer, you will have a policy file named Ntconfig.pol. By default, computers running Windows NT are set to download appropriate policy files from the NETLOGON share on the domain controllers. Policy files should be placed in this share on the PDC; the contents of this share are then replicated to the BDCs via the Directory Replicator service. Computers running Windows NT can download all applicable policies from the domain controller they use for logon. These behaviors are the same whether the domain controllers are Windows NT or Windows 2000 or Windows Server 2003, because Windows 2000, Windows XP, and Windows Server 2003 clients will ignore policy files in the NETLOGON share in preference for the Active Directory-integrated GPOs.

System policies are loaded and applied on Windows NT in the following manner:

1. If a user-specific policy exists, it is loaded, the registry is modified, and the policy processing skips to step 4.
2. If the Default User policy exists, it is loaded and applied.
3. If Group Policies exist, they are loaded and applied in ascending order of priority as applicable. If a Group Policy conflicts with the Default User policy, it will take precedence unless the Group Policy setting is to "ignore."
4. If a computer-specific policy exists, it is loaded, the registry is modified, and the policy processing skips to step 6.
5. If a Default Computer policy exists, it is loaded, and the registry is modified.
6. The policy application is complete.

By default, Windows NT is set to download and apply system policies. This behavior is described in KB article 168231, "System Policies Are Not Applied in Windows NT" at http://support.microsoft.com/?kbid=168231, and is controlled by the following registry value:

**HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\Update\UpdateMode** (REG_DWORD)

- A value of **0** disables the application of system policies.
- The default value of **1** enables Automatic mode. Windows will look for the Ntconfig.pol system policy file on the authenticating domain controller as described previously.
- A value of **2** enables Manual mode. Windows will look at the **NetworkPath** (REG_SZ) value (in the same key) and attempt to find the policy file located there.

It is important to verify the proper functioning of system policies. KB article 154120, "Debugging User Profiles and System Policies in Windows NT 4.0" at http://support.microsoft.com/?kbid=154120 describes the process of replacing userenv.dll with the check build version, allowing you to create a log file that you can use to debug policy application. This procedure is good on all versions and service pack levels of Windows NT 4.0, including Terminal Server edition.

Even though the default system policy behavior depends on a functional Windows NT domain infrastructure, computers running Windows 2000, stand-alone computers running Windows NT, and users in the local account database can still benefit from system policies. KB article 168579, "How to Set Up Locally-Based System Policies" at http://support.microsoft.com/?kbid=168579 gives specific directions for two methods of configuring a Windows NT system so that it can provide system policies for users in the local account database. If you are using computers running Windows 2000, Windows XP, or Windows Server 2003 on a Windows NT domain, KB article 274478, "Group Policies

for Windows 2000 Professional Clients in Windows NT 4.0 Domain or Workgroups" at http://support.microsoft.com/?kbid=274478 gives the process necessary to allow Windows 2000-compatible policies to be distributed by Windows NT domain controllers.

## Planning Considerations for System Policy Deployment

Keep in mind these potential system policy complications as you develop and deploy your policies, in order to ensure the best level of security:

- Because of various bugs in the system policy implementations in Windows NT 4.0, you must have at least SP3 to ensure that all of the important patches related to system policies have been applied to your systems. Generally, this is not a problem unless you have specific requirements for supported applications, because other security requirements virtually demand that you have all your Windows NT systems on SP6a.
- System policies require up to two logons and logoffs in order to ensure that system policies are being located, downloaded, and applied.
- Because policies do not automatically remove themselves from a computer, you should create a group-specific policy for your administrator accounts. This policy should be set to remove, at a minimum, the restrictions necessary for your administrators to re-grant themselves access to those features they may need; a common Administrator Group Policy is to simply restore all restricted features.
- Carefully read the individual settings in your policies and be sure that you are correctly parsing any potential double negatives. Some settings will require you to enable them in order to turn off the specified behavior.
- You should double-check the location of your system policy files. They should be in your NETLOGON share, the location of which varies according to whether the primary domain controller is running Windows NT 4.0 or Windows 2000 or Windows Server 2003.
- It is important to ensure that your Directory Replication service is working correctly and that the contents of the NETLOGON share on the PDC are getting properly copied to all the BDCs.
- You should ensure that the modification time of your policy files is updated every time you modify and deploy a new policy. Some clients with older service packs will cache the policy files, and they use the modification time as their indicator to refresh the file. Better yet, ensure that all of your computers running Windows NT are on at least SP3, and preferably on SP6a.
- If you have a mix of Windows NT 4.0 and Windows 2000 or Windows Server 2003 domain controllers, it is possible to establish a replication process to bridge between Windows NT's Directory Replication service and Windows 2000's File Replication Service. See KB article 317368, "HOW TO: Use Lbridge.cmd to Replicate System Policies Between Windows 2000 and Windows NT 4.0 Domain Controllers" at http://support.microsoft.com/?kbid=317368 for details on this process.
- Windows NT 4.0 Terminal Server Edition poses an interesting set of complications for the use of system policies. See KB article 192794, "How to apply System Policy settings to Terminal Server" at http://support.microsoft.com/?kbid=192794 for more information.

It is worth repeating that system policies are not an adequate substitute for proper application of registry and file system access control lists (ACLs), nor will they allow you to skimp on other system hardening measures. As an example, consider the system with a system policy that restricts the user to only running the Microsoft Office application binaries. The user could make use of the standard Office **File** menu to create new folders, copy executables such as cmd.exe or either of the registry editors to writeable locations, and rename them to executable names permitted by the system policy – thus circumventing the policy.

There are specific areas of weakness that you need to consider when deploying system policies in order to supplement them with other security measures:

- Any restricted executables in system policies should be specified by full pathnames and should not rely on the default search path.
- You should consider restricting the use of the **Tools** and **Views** menus in Windows Explorer. These menus contain many options that can be used to overcome policy restrictions.
- Registry files (.reg) can normally be executed even if access to **RegEdt32** and **RegEdit** has been removed, because the default file name extension association is still in place.
- No user-writeable directories, such as their temporary directory and profile directory, should ever be part of the default search path.
- The World security ID (SID) in the registry has more access than it should. You should strongly consider restricting it to only the Query/Enumerate/Read permissions; however, this may break older applications and require extensive testing to determine the specific permissions that you must apply in order to retain functionality.
- All executables in system directories should be carefully scrutinized. Even if you remove the **Execute** permission from them, if a user has **Read**, he or she can copy them into a writeable location and attempt to run them from there.

## Security Configuration Manager

Security Configuration Manager (SCM) was originally designed for Windows 2000, but Microsoft made it accessible to Windows NT 4.0 SP4 and later. SCM is available on the SP6a CD-ROM or by download from the Microsoft FTP site. The chief tool in SCM is the Security Configuration Editor (SCE), which you use to create and manage security templates and system policies.

You can use the SCE to perform several useful tasks. For example, with it you can:

- Build security templates that specify auditing, user rights, and security settings for computers running Windows NT.
- Apply templates automatically to one or more computers in a domain.
- Scan one or many computers to assess their levels of compliance with a particular template.

By default, installing SCE adds a set of templates in the %winnt%\security\templates folder. Each template is intended for a particular security environment and computer type:

The basic templates (Basic*4.inf) apply a standard level of security to target computers, including a six-week password age limit and a basic set of registry key, file system, and user rights permissions.

The compatible templates (Comp*4.inf) apply stronger security than the basic templates. However, security settings that can interfere with older clients or applications remain turned off; the goal of this template set is to raise security without introducing application compatibility problems.

The high security templates (HiSec*.inf) apply some additional security policies, including increasing the minimum password length to eight characters from seven and using a more restrictive set of file system and registry permissions. These templates can cause problems with older applications that depend on file system or registry permissions or user rights assignments.

The secure templates (Secur*.inf) are the most protected standard templates, but they enable features—including Server Message Block (SMB) signing—that will limit interoperability between secured computers and Windows 95/98 clients that are not similarly configured.

**Table 4.1: Predefined Templates That Ship with the Windows NT SCE Toolset**

| If you want | To secure | Use |
|---|---|---|
| Basic security | Primary / backup domain controllers | BasicDC4.inf |
| | Member servers | BasicSv4.inf |
| | Workstations | BasicWk4.inf |
| Improved security with good application compatibility | Primary / backup domain controllers | CompDC4.inf |
| | Member servers | CompDC4.inf |
| | Workstations | CompWS4.inf |
| High security with reduced application compatibility | Member server or domain controller | HiSecDC4.inf |
| | Workstation | HiSecWS4.inf |
| High security with reduced Windows 98 connectivity | Member server or domain controller | SecurDC4.inf |

Because the SCM was originally designed for Windows 2000, it provides a familiar interface for those administrators who use Group Policies. This familiar interface includes updating the integrated Windows NT ACL editor to the same ACL editor used on Windows 2000. This action has important consequences on ACL behavior throughout the entire system, because SCM requires the Windows 2000 ACL inheritance model, which applies inherited ACLs dynamically. Rather than copying access control entries (ACEs) from the parent object onto the child object, the new subsystem references the **Allow Inheritable Permissions** policy to see if it is enabled. The updated inheritance system will be used by the normal ACL editors in the Windows Explorer as well as through the policies applied by the SCM. However, it will not be directly accessed by the Registry Editor (regedt32.exe), so policies applied by SCM that change ACLs in the registry can create ACLs that cannot be manipulated by the Registry Editor. Further information about the new ACL editor system and its consequences, including directions for disabling it (which in turn disables SCM), can be found in KB article 195509, "Installing Security Configuration Manager from SP4 Changes Windows NT 4.0 ACL Editor" at http://support.microsoft.com/?kbid=195509.

The SCM lets you configure the following areas:

- Account Policies, which includes password and account lockout policy options.
- Local Policies, which includes audit, user rights assignment, and security policy options.
- Event Log policies.
- Restricted Groups, which provides the ability to lock down the membership of designated groups.
- System Services, which provides options to configure various services and transports.
- Registry permissions.
- File System permissions.

The SCM GUI provides a listing of the various policies that have been defined and that are stored in the %winnt%\security\templates folder. The graphical user interface (GUI) lets you define, edit, and maintain multiple policies suitable for the different older systems you maintain, as well as analyze the system and compare its compliance with a selected policy template.

The primary difference between SCM and Active Directory policies is that Windows NT does not provide an automated way of deploying policy templates throughout a related group of computers, so you still must provide some way of distributing configuration templates to the appropriate computers. Distribution will require either manual intervention or some sort of scripting, combined with some mechanism to ensure that policies are automatically applied on a regular basis.

Use SCM to assist in making and deploying the following basic account policy changes on all systems:

- Rename the Guest account (in addition to disabling it) by using the SCM (via the Local Policies | Security Options | Change Guest account name to setting). Even though it is disabled, ensure that it has an extremely strong password. Regularly audit to ensure that it is not included in any groups other than Domain Guests, which should have no other members.
- Rename the Administrator account to something that is not obvious (using the **Local Policies | Security Options | Change Administrator account name to** setting), then create a decoy account named Administrator with no privileges or permissions. Treat this decoy account as you did the Guest account by setting a strong password for it and disabling it. Do not make the common mistake of labeling this account as a decoy in the comments; this will defeat the purpose of the account. You also need to regularly audit to ensure that the decoy account is not added to any groups; be sure to watch out for authentication attempts against this account. Although this measure does not prevent attackers inside the network from using the well-known SID to find the real name of the administrator account, it does slow down sloppy attackers and attempts from external networks, giving you advanced warning of hostile intentions.
- Prohibit the enumeration of account names and file shares for anonymous users (the null session). By default, Windows permits a remote system to connect with no credentials and list this information. This technique is a common route of attack. Use the **Disallow enumeration of account names and shares by anonymous** setting in the Security Policies object of the Local Policies node.

## Choosing the NTLM Authentication Level

Windows NT 4.0 clients earlier than SP4 use NTLMv1 authentication, which has been proven to have design flaws that make it vulnerable to interception and decryption by attackers. NTLM authentication was designed as a successor to the even weaker LAN Manager (LM) authentication protocol, which requires the storage of all password information in a weaker hash known as the LM hash (which the subsequent section describes in more detail). The LM hash is easily cracked and broken, leading to compromise of the password.

LM, NTLM, and NTLMv2 are used to authenticate such operations as:

- Joining a domain.
- Authenticating between Active Directory forests.
- Authenticating to Windows NT 4.0 domains.
- Authenticating to computers running Windows NT 4.0 or Windows 98 acting as file or print servers.
- Authenticating to servers that are not domain members.

Because Trey has a mix of computers running Windows Server 2003 (for domain controllers), Windows 98, and Windows NT 4.0, it needs to choose an NTLM authentication level that provides the best possible security while still allowing existing clients to operate properly. Windows NT, Windows 2000, and Windows Server 2003 support six levels of LM compatibility, controlled through the **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\LSA\LMCompatibility Level** registry key:

- **0 (Send LM & NTLM responses)**. This setting offers the most interoperability, because clients may use LM or either version of NTLM to authenticate. However, it allows insecure LM traffic on the network.
- **1 (Send LM & NTLM–use NTLMv2 session security if negotiated)**. This setting offers more security, but it requires the DSClient software for Windows 98. When this value is chosen,NTLMv2 can be negotiated if both computers support it. This value is the best setting to use during deployment; all computers with DSClient will use NTLMv2, while all other clients with earlier version of Windows can still authenticate.
- **2 (Send NTLM response only)**. This value will restrict the use of Windows 98 clients without the DSClient installed because servers will not answer LM requests from unmodified computers running Windows 98.
- **3 (Send NTLMv2 response only)**. Use this value only if all clients running versions of Windows prior to Windows NT 4.0 SP4 have DSClient installed.
- **4 (Send NTLMv2 response only\refuse LM)**. Like the previous setting, with the addition that if LM authentication is requested, no NTLMv2 response is sent.
- **5 (Send NTLMv2 response only\refuse LM & NTLM)**. Like the previous setting, except that NTLM requests are ignored, as well. This setting is only appropriate for networks where all computers are either running Windows 2000 or later or running Windows 98 with the DSClient installed.

## Upgrading NTLM Authentication for Windows 98

The DSClient extension allows Windows 98 systems to use NTLMv2 for authentication. By default, NTLMv2 session security encryption uses a 56-bit maximum key length. To enable 128-bit NTLMv2 support, ensure that Internet Explorer 4 or higher is installed and that 128-bit support is installed and configured. This preparation must be performed before you install the DSClient.

After you have distributed the DSClient software to your client computers, you can restrict which versions of the LM protocol family can be used for authentication. You must configure these restrictions on all domain controllers and other servers that provide accounts used for remote authentication.

## Upgrading NTLM Authentication on Windows NT 4.0

Windows NT 4.0 SP4 and later include the ability to selectively use and respond to LM, NTLM, and NTLMv2 authentication requests. This change requires the addition of a registry key (described in the following implementation section). Trey chose to set the compatibility level to 5 (allow NTLMv2 only) on its Windows NT systems, but it waited to make this change until all Windows 98 systems had the DSClient extension installed.

## Upgrading NTLM Authentication on Windows 2000 and Windows Server 2003

Trey's IT administrators created a new GPO on the company's domain controllers to set the NTLM authentication level for computers and domain controllers and attached it to the Domain Controllers organizational unit (OU). They accomplished this by modifying the **Computer Configuration\Windows Settings\Security Settings\Local Policies\Security Options\LAN Manager Authentication Level** setting. Be sure to use the **Domain Controllers** policy for domain controllers, the **Domain** policy for Active Directory members, and local policy for stand-alone computers or computers that are members of Windows NT domains.

After DSClient is fully deployed on all Windows 98 systems, it is appropriate to change the LM authentication level setting to **3**, **4**, or **5**, which disallows the use of LM authentication. Until then, systems running earlier versions of Windows without DSClient will lose access to all network data.

Because Trey's environment includes Windows 98 clients, all of which have the DSClient extension, Trey set this value to **5** on Windows NT-based computers (Windows NT, Windows 2000, and Windows XP) and to **3** on computers running Windows 98. Otherwise, settings higher than **3** on computers *not* running Windows 98 will prevent those computers without the DSClient extension from authenticating.

Refer to KB article 305379, "Authentication Problems in Windows 2000 with NTLM 2 Levels Above 2 in a Windows NT 4.0 Domain" at http://support.microsoft.com/?kbid=305379 for information about a hotfix that ensures that this setting works in networks with a mix of Windows 2000 and Windows NT 4.0 systems.

Windows systems create communication channels, known as secure channels, which are used to authenticate computer accounts and user accounts in trusted domains. Secure channels prevent man-in-the-middle attacks and allow clients to access account databases in trusted domains. Secure channels can be digitally encrypted or signed, but all domain controllers must be running Windows NT 4.0 SP6a or higher. Because secure channel signing affects all clients in the domain, as well as clients in trusting domains, think carefully before enabling this feature.

## Defining Effective Password and Lockout Policies

One of the fundamental principles of security is to avoid transmitting or storing plaintext passwords. A password stored in plaintext is easily revealed, both from examination of the SAM or from sniffing network sessions. Modern Windows authentication protocols use encryption algorithms known as hashes to protect all transmission of authentication credentials as well as to store that password from the moment the user enters it. However, it is necessary to define policies that control password length and strength, as well as how many failed logon attempts are allowed before the account is locked out.

The Microsoft Security Guidance Kit (SGK) (available at www.microsoft.com/security/guidance/order/default.mspx) contains a number of reference documents that can help in choosing and implementing a strong password and account lockout policy. Additional guidance is available in the "Account Passwords and Policies" white paper at www.microsoft.com/technet/prodtechnol/windowsserver2003/technologies/security/bpactlck.mspx.

## Hardening the File System

In conjunction with boot hardening, you need to increase security of the file systems and storage devices on each computer running Windows NT 4.0 in the organization.

The first critical step is to make sure that all of your computers running Windows NT 4.0 are using NTFS file system (NTFS). NTFS provides file and folder permissions, in addition to the share-level permissions offered in Windows NT, Windows 98, and their descendants. Using NTFS permissions restricts the ability of users on a computer to read files belonging to other users, and it makes it more difficult for attackers who gain physical access to the console to steal or modify data without booting into an alternate operating system. Windows NT includes a command that nondestructively converts file allocation table (FAT) or FAT32 volumes to NTFS, although it requires exclusive access to the volume.

When a volume is converted to NTFS with the convert utility, the default permission applied is **Everyone:Full Control**. You need to change this permission, because it is unnecessarily permissive; existing NTFS systems also might have unnecessarily lax permissions.

The National Security Agency (NSA) has developed a set of recommended permissions for Windows NT servers and workstations and published them in the *Guide to Securing Microsoft Windows NT Networks*. You can apply these recommendations manually with scripts that use the xcacls tool (described in KB article 318754, "HOW TO: Use Xcacls.exe to Modify NTFS Permissions" at http://support.microsoft.com/?kbid=318754) or by using predefined SCM templates. The recommendations contained in chapters 11 and 13 of the guide are quite extensive, so you must refer to that publication for detailed instructions for implementing them.

## Hardening Services

There are several configuration changes centered on securing services and accounts that are usually appropriate for all Windows NT environments.

One of the many improvements that Windows NT 4.0 SP3 delivered was the creation of the Authenticated Users security group. This group is intended to replace the Everyone security principal in all areas where you want to prohibit any user or process, regardless of authentication status, to be able to access files and services. This group should be replaced in every file ACL that you create and manage, unless you truly need to allow anonymous connections. For most services, this should not be necessary; IIS, as an example, maps all anonymous requests to a specified account. Third-party applications should be carefully tested, however, because they may depend on inclusion of the Everyone group.

In addition to the Authenticated Users group, you also need to disable anonymous connections to the registry and restrict the ability of anonymous users to enumerate the domain user account list and share lists on servers via the following registry entries:

**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\LSA**

**RestrictAnonymous** (REG_DWORD)

- A value of **0** disables this restriction and allows anonymous users (null sessions) to enumerate the user list and share list.
- A value of **1** enables this restriction; null sessions will not be allowed to enumerate the user list and share list.

**RestrictNullSessAccess** (REG_DWORD)

- A value of **0** disables this restriction and allows anonymous users (null sessions) to connect to the registry.
- A value of **1** enables this restriction; null sessions will not be allowed to connect to the registry.

KB article 143474, "Restricting information available to anonymous users" at http://support.microsoft.com/?kbid=143474 demonstrates how you can use the Authenticated Users group in conjunction with manual registry edits to completely restrict anonymous connections to the registry and enumerations of users and shares. KB article 153183, "How to Restrict Access to the Registry from a Remote Computer" at http://support.microsoft.com/?kbid=153183 gives additional information on restricting remote access to the registry for authenticated users.

Another important change that you must make, especially on domain controllers, is to properly secure the Directory Replicator Service (DRS). This service allows the contents of specified directories to be replicated to multiple servers and is important in ensuring that system policies and other critical files are propagated to all domain controllers. It can be extremely useful in maintaining load-balanced file services for critical applications.

Replication connections are configured in the Server Manager for Domains, by clicking the **Replication** button. For each replication partnership, there is an Exporter (the source) and an Importer (the destination). The details of each partnership must be separately configured on each side of the connection (similar to the way that trust relationships are created between domains).

By default, the DRS uses the Local System account, something you should change right away. Create a separate domain user account specifically for the DRS to use so that the same account is used on all your computers; do not use local accounts. After you have created the accounts, you can use the Service Manager to ensure that the DRS is running as this account and not the Local System account.

Many services are enabled by default and can be disabled to reduce the threat surface for your computers running Windows NT; still others you can configure to run as non-privileged accounts, rather than the Local System account. For more information about securing the DRS and disabling unnecessary services, including a thorough set of recommendations for each service based on the role of the computer, see Chapter 10 of the *Microsoft NT 4.0 Security, Audit, and Control Technical Reference*, which is available for preview on Microsoft TechNet at www.microsoft.com/technet/prodtechnol/winntas/maintain/nt4sac/sacch10.mspx.

## Other Hardening Measures

A number of other hardening measures may be appropriate for various environments.

### Disabling Auto Generation of 8.3 Filenames

Windows NT supports 8.3 file name formats for backward compatibility with16-bit applications. The 8.3 file name convention is a naming format that allows file names up to eight characters long, which means that an attacker only needs eight characters to refer to a file that may be 20 characters long. For example, a file named Thisisalongfilename.doc, could be referenced by its 8.3 file name, Thisis~1.doc. If you avoid using 16-bit applications, you can turn this feature off. Disabling short name generation on an NTFS partition also increases directory enumeration performance.

Attackers could use short file names to access data files and applications with long file names that would normally be difficult to locate. An attacker who has gained access to the file system could access data or execute applications.

You can control this through the following registry entry:
**HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\FileSystem\NtfsDisable8dot3NameCreation**

The 16–bit applications in your organization will not be able to access files with names longer than the 8.3 format allows if those files are stored on computers where this change has been made. Applying this setting to an existing server that already has files with auto generated 8.3 file names does not remove them. To remove existing 8.3 file names, you will need to copy those files off the server, delete the files from the original location, and then copy the files back to their original locations.

### Disabling Autorun

The Autorun feature begins reading from a drive on your computer as soon as media is inserted into it. As a result, the setup file of programs and the sound on media starts immediately. To prevent a possible malicious program from starting when a CD or DVD is inserted, Group Policy disables Autorun on all drives.

An attacker with physical access to the system could insert Autorun enabled media into the computer and automatically launch malicious code that will run in the context of the currently-logged-in user.

You can control this through the following registry entry:
**HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\CDROM\Autorun**

After you make this change, Autorun will no longer work when autorun-enabled discs are inserted into drives on the computer.

## Removing OS/2 and POSIX Subsystems

OS/2 is a family of Microsoft protected-mode, virtual-memory, multitasking operating systems for personal computers based on the Intel 80286 and 80386 processors. The Portable Operating System Interface for UNIX (POSIX) is an Institute of Electrical and Electronic Engineers (IEEE) standard that defines a set of operating system services. The OS/2 subsystem is required if the server needs to significantly interact with OS/2 clients; the POSIX subsystem is required to run applications that use POSIX services.

These subsystems introduce a small degree of security risk relating to processes that can potentially persist across logons. That is, if a user starts a process and then logs off, a potential exists that the process will be accessed by the next user who logs on to the system. The process started by the first user may retain the user's system privileges.

In order to disable these subsystems, the binary files and following registry entries required for these subsystems can be deleted:

- The **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\OS/2 Subsystem for NT** registry key and subkeys.
- The **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Session Manager\Environment\OS2LibPath** registry key.
- The POSIX and OS/2 subkeys of the **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Session Manager\Environment\Subsystems** registry key.
- The **\winnt\system32\os2** directory and subdirectories.

Applications that rely on the OS/2 or POSIX subsystems will no longer operate.

## Increasing Object Protection Levels

The Windows NT kernel may allow callers to change attributes of kernel objects under various conditions. In some circumstances, this may allow malicious callers to escalate their privileges.

You can control this through the following registry entry:

- **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Session Manager\EnhancedSecurityLevel**

This change only allows kernel-mode code to change kernel object attributes for the current process.

There is no potential impact for user-mode applications.

## Preventing Users from Adding Printer Drivers

Windows NT default permissions allow users to install additional printer drivers. In some circumstances, this may allow malicious or poorly written drivers to escalate privileges.

You can control this through the following registry entry:

- **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Print\Providers\LanMan Print Services\Servers\AddPrinterDrivers**

This change only allows members of the Printer Operators and Administrator groups to add new printer drivers.

There is no potential impact for user-mode applications.

## Confirming That Automatic Administrator Logon Is Disabled

Windows NT may be configured to allow the automatic logon of the administrator account when the computer is restarted. This exposes the console and creates a new registry key that contains the administrator password in plain-text format. KB article 97597, "How to Enable Automatic Logon in Windows NT 3.x and 4.0" at http://support.microsoft.com/?kbid=97597 discusses this setting in more detail.

You can control this through the following registry entries:

- **HKEY_LOCAL_MACHINE\Software\Microsoft\WindowsNT\CurrentVersion\Winlogon\AutoAdminLogon**
- **HKEY_LOCAL_MACHINE\Software\Microsoft\WindowsNT\CurrentVersion\Winlogon\DefaultPassword**

There is no potential impact for user-mode applications.

## Disabling Notifications for the Novell Client

Windows NT is configured by default to notify the Novell networking client (FPNWCLNT.DLL) when passwords are changed; however, unless the client is installed, a copy of this dynamic-link library (DLL) is not installed. This allows an attacker to craft a replacement DLL and hijack all user password change events.

To disable this behavior, ensure that the value "FPNWCLNT" is not contained in the **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\LSA\Notification Packages** registry entry.

Password synchronization between Windows and Netware is disabled in environments that use Netware.

# Implementation

## Implementation Prerequisites

For these implementation details to work correctly, you must have a basic Trey Research infrastructure implemented as introduced in Chapter 2, "Applying the Security Risk Management Discipline to the Trey Research Scenario."

## Implementation Overview

To implement this solution scenario, you will need to perform the following activities:

- Establish a baseline and patch the operating system
- Enable the Syskey
- Reduce the boot timeout
- Install the SCM
- Load the SCM snap-in
- Apply the Trey policy template
- Prevent the storage of LM password hashes
- Set the NTLM authentication level
- Convert the server volumes to NTFS
- Perform other hardening measures

## Establishing a Baseline for the Operating System

Whenever you build a new computer or reinstall the operating system on an existing system, you must install the correct set of security patches to ensure that the computer is adequately protected. Patch management is a complicated topic covered more fully in Chapter 6, “Patch Management”; however, the process of baselining new systems is a necessary part of securing computers running Windows NT. Every computer running Windows NT in your organization should have the following patches installed:

- Service Pack 6a (SP6a) for Windows NT 4.0, which is the most recent service pack for Windows NT 4.0. It contains a complete and regression-tested set of fixes that are designed to be installed as a unit. Since the release of Windows NT 4.0 SP6a, U.S. law on exporting cryptographic software has changed. You might be able to increase the encryption strength available to your computers running Windows NT 4.0 by upgrading to the high encryption version of SP6a. Installing the high encryption version of SP6a on a standard encryption version will automatically upgrade the operating system. Windows NT 4.0 Service Pack 6a is available at www.microsoft.com/ntserver/nts/downloads/recommended/sp6/128bitx86/.

▶ **To check the encryption version of Windows NT 4.0**

1. As an Administrator, log on to the target computer running Windows NT 4.0.
2. Use Windows Explorer to open the \winnt\system32 folder.
3. Locate the schannel.dll file, select it, and then click **File** and **Properties**.
4. Click the **Versions** tab.
5. Review the contents of the **Description** field.
   If it says “Export Version”, you have the standard encryption version; if it says “U.S. domestic version,” you have the high-encryption version.
   - The Post-Windows NT 4.0 Service Pack 6a Security Rollup Package (SRP) incorporates several patches that were released after the original SP6a release and is a prerequisite for follow-up patches from Windows Update. It is documented in KB article 299444, "Post-Windows NT 4.0 Service Pack 6a Security Rollup Package (SRP)” at http://support.microsoft.com/?kbid=299444; the Security Rollup Package itself can be downloaded from http://download.microsoft.com/download/winntsp/Patch/q299444/NT4/EN-US/Q299444i.exe.
   - Microsoft has not released an integrated SRP or service pack for Windows NT 4.0 since the release of the SRP, but individual components such as portions of the operating system, Internet Explorer, and add-on utilities such as the Routing and Remote Access Service (RRAS) and Internet Information Services (IIS) have been updated. You can obtain a current list of patches from Microsoft TechNet, or you can manually review the list of patches available through Windows Update.
   - Internet Explorer 6.0 Service Pack 1 (SP1) is the current version, available at www.microsoft.com/windows/ie/. It incorporates hundreds of security fixes and improvements. Depending on your environment, you may choose to download Internet Explorer directly from the Microsoft Web site, order it on a CD, or install it on multiple computers by using the Internet Explorer Administration Kit, available at www.microsoft.com/windows/ieak/.

Because of the length of time since the release of Windows NT, Trey’s IT staff had already applied SP6a and the SRP to all of the computers running Windows NT. However, it lacked a consistent policy for ensuring that other necessary patches released by Microsoft had been applied, which led directly to the company's adoption of the patch management methodology described in Chapter 6, "Patch Management."

## Enabling Syskey

You must enable Syskey to allow encrypted protection of the computer SAM. Trey chose to apply this protection to all of the company's computers, servers, and workstations alike running Windows NT.

▶ **To enable Syskey**

1. As an Administrator, log on to the target computer running Windows NT 4.0.
2. Start the Rdisk.exe utility from a command prompt or by using the **Run** command.
3. When Rdisk starts, click the **Update Repair Info** button, and then when prompted, click **Yes**.
4. After the repair information update completes, insert a blank floppy disk into the floppy disk drive and then, in the **Repair Disk Utility** dialog box, click **Yes**.

5. Click **OK** to confirm that you want the repair disk to be created.
6. When the disk creation is finished, click **OK** to acknowledge the security warning, and then click **Exit** to close the Rdisk utility.
7. Start the Syskey.exe utility.
8. When the Securing the Windows NT Account Database window appears, click the **Encryption Enabled** button, and then click **OK**.
9. Click **OK** in the confirmation dialog box.
10. From the **Account Database Key** dialog box, click **Store Startup Key Locally** (see Figure 4.1), and then click **OK**.
11. When the **Success** dialog appears, click **OK**.
12. Restart the computer.

---

**Note:** On domain controllers, the recovery data gathered by Rdisk will probably be too large for a single floppy disk, but it is also saved to the **%systemroot%\repair** directory. Having an up-to-date ERD allows quick recovery in case of a problem with Syskey. The Rdisk /s command is more fully described in KB article 122857, "RDISK /S and RDISK /S- Options in Windows NT" at http://support.microsoft.com/?kbid=122857.

---

**Note:** For more information on how Syskey protects SAM data against attack, see KB article 143475, "Windows NT System Key Permits Strong Encryption of the SAM" at http://support.microsoft.com/?kbid=143475.

---

## Reducing the Boot Timeout

This procedure explains how to modify the boot timeout value.

▶ **To reduce the boot timeout**

1. As an Administrator, open a command prompt.
2. Change to the root directory of the system volume (usually C:\).
3. Reset the read-only, hidden, and system attributes on the **Boot.ini** file by using the **attrib** command:
   ```
   attrib –s –h –r boot.ini
   ```
4. Open the **Boot.ini** file with a text editor such as Notepad.
5. Find the line that says **timeout=30** and change the **timeout** value to **0**.
6. Save the file, and then close the text editor.
7. Restore the read-only, hidden, and system attributes on the **Boot.ini** file by using the **attrib** command:
   ```
   attrib +s +h +r boot.ini
   ```

## Installing the SCM

This procedure explains how to install the SCM tools. The SCM can run on any computer running Windows NT 4.0 SP4 or later, on both Server and Workstation editions. For ease of use, the Trey IT staff installed the SCM on their personal workstations, where it can be run against any of the computers running Windows NT 4.0 in their domain.

▶ **To install the SCM**

1. As an Administrator, open Windows Explorer.
2. Create the directory c:\temp if it does not already exist.
3. Download the SCM installer from www.microsoft.com/ntserver/nts/downloads/recommended/scm/default.asp and save the downloaded file to C:\Temp.
4. Open C:\Temp and double-click the **Scesp4i.exe** application.
5. When prompted by the installation utility, specify a temporary file path of C:\Temp\scminstall in which to uncompress the support files, and then click **OK**.
6. Use Explorer to open C:\Temp\scminstall.
7. Double-click the **Mssce** utility to install the SCE and the Microsoft Management Console (MMC) tool required for the GUI version of the SCE. The installer will automatically place the SCM components in the correct location on the computer where you are installing.

## Loading the SCM Snap-in

After the SCM has been installed on a workstation or server, you can execute it. This procedure explains how to load the SCM snap-in into the MMC.

▶ **To load the SCM snap-in**

1. As an Administrator, log on to a computer where you have previously installed the SCM.
2. Launch the MMC.
3. Select **Console**, and then select **Add/Remove Snap-In**.
4. Click **Add**.
5. Select **Security Configuration Manager**.
6. Click **OK**, and then click **OK** again.

## Applying the Trey Policy Template

This procedure explains how to use the SCE GUI and command-line interface (CLI) tools to apply the Trey policy template. Trey decided to use the Comp* templates because they provide the best balance between backward compatibility and security.

▶ **To customize the policy template from the SCE GUI**

1. Log on to a test Windows NT 4.0 workstation.
2. Make backup copies of %systemroot%\security\templates\compws4.inf and %systemroot%\security\templates\compdc4.inf.
3. Use Notepad to open %systemroot%\security\templates\compws4.inf.
4. Search for the line that contains "**MACHINE\System\CurrentControlSet\Services\Rdr\Parameters\EnableSecuritySignature**".

5. Edit that line so that the value for **EnableSecuritySignature** is **1**.
6. Save the file, and then close Notepad.
7. Use Notepad to open %systemroot%\security\templates\compdc4.inf.
8. Repeat steps 4-6 with the new file.

▶ **To apply the policy template from the SCE GUI**

1. As an Administrator, launch the Security Configuration Manager.
2. Select the **Security Configuration Manager** node.
3. Select the **Configurations** node.
4. Select the default configuration file directory (%systemroot%\security\templates) to show the configuration templates.
5. Select the appropriate configuration file (compws4.inf or compdc4.inf).
6. Familiarize yourself with the various objects and settings in the policy template.
7. Select the **Security Configuration Manager** node.
8. Right-click the **Database** node, and then select **Import Configuration**.
9. Choose the corresponding policy template, and then click **OK**.
10. Right-click the **Database** node, and then select **Configure System Now**.
11. Click **OK** to accept the default log file path.
12. Wait for the policy to apply, and then review the log file results.
13. Close the log file.
14. Close the MMC.

▶ **To apply the policy template from the command line**

1. Log on to the target computer.
2. Open a command prompt.
3. Run the following command:

   ```
   secedit /configure /cfg c:\winnt\security\templates\treysec.inf
   /overwrite
   ```

4. Examine the results.
5. Close the command prompt.

## Preventing the Storage of LM Password Hashes

This procedure explains how to configure your domain controllers to prevent storage of the LM password hashes using either policies or direct registry editing.

▶ **To prevent the storage of LM hashes by using Group Policies or local policies on Windows 2000 SP2 or later, Windows XP, or Windows Server 2003**

1. Open the Group Policy editor.
2. Expand **Computer Configuration**, **Windows Settings**, **Security Settings**, and **Local Policies**, and then click **Security Options**.
3. In the list of available policies, double-click **Network security: Do not store LAN Manager hash value on next password change**.
4. Click **Enabled**, and then click **OK**.
5. Restart your computer, and then change your password.

## Setting the NTLM Authentication Level

This procedure explains how to configure the level of NTLM authentication used on your domain controllers.

▶ **To set the NTLM level by using Group Policies on Windows 2000, Windows XP, or Windows Server 2003**

1. Log on to a domain controller as an Administrator.
2. Launch Active Directory Users and Computers.
3. Open the Group Policy editor.
4. Right-click the domain to which you want to apply the NTLM authentication level setting, and then select **Properties**.
5. In the **Properties** dialog box, click the **Group Policy** tab.
6. Click the **New** button to create a new Group Policy object (GPO).
7. Type **NTLM Authentication Level** and then press **ENTER**.
8. Click the **Edit** button.
9. Expand **Computer Configuration**, **Windows Settings**, **Security Settings**, and **Local Policies**, and then click **Security Options**.
10. In the list of available policies, double-click **Network Security: LAN Manager Authentication Level**.
11. Select **Send NTLMv2 response only\refuse LM**, which provides the best balance between security and compatibility for mixed networks that include computers running Windows 98 and Windows NT.
12. Click **OK**.
13. Close the Group Policy Editor.
14. In the domain **Properties** dialog box, right-click the NTLM Authentication Level GPO ,and then select the **No Override** command.
15. Click **Close**.

▶ **To set the NTLM level through registry entries on Windows NT 4.0 SP3 or later**

1. Start Registry Editor (regedt32.exe).
2. Locate and then click the following key:
   **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa**
3. On the **Edit** menu, click **Add Value**.
4. For **Data Type**, select **DWORD Value**.
5. For **Value Name**, type **LMCompatibilityLevel** and then click **OK**.
6. For **Data**, type **3** (or your desired level) as a **Decimal** value.
7. Quit Registry Editor.
8. Restart the computer.

▶ **To set the NTLM level through registry entries on Windows 98**

1. Start Registry Editor (regedit.exe).
2. Locate and then select the following subkey in the registry:
   **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control.**
3. On the **Edit** menu, point to **New**, and then click **Key**.
4. Type **Lsa** and then press **Enter**.

5. On the **Edit** menu, point to **New**, and then click **DWORD Value**.
6. Type **LMCompatibilityLevel** and then press **Enter**.
7. On the **Edit** menu, click **Add Value**, and then add the following registry value:
   - **Value Name**: LMCompatibility
   - **Data Type**: REG_DWORD
   - **Value**: 3
8. Quit Registry Editor.
9. Restart your computer.

## Converting a Volume to NTFS

This procedure explains how to convert a FAT volume to NTFS.

▶ **To convert a FAT volume to NTFS**

1. As an Administrator, open a command prompt.
2. Use the **convert** command to specify which drive is to be converted:
   ```
   convert d: /fs:ntfs
   ```
3. If prompted, reboot to allow the conversion utility to gain exclusive access to the selected drive.
   Rebooting will be necessary if you convert the system partition, the volume which contains the current working directory for the command interpreter, or a volume on which applications have files open.

---

**Note:** KB article 214579, "How to Use Convert.exe to Convert a Partition to the NTFS File System" at http://support.microsoft.com/?kbid=214579 describes the **convert** command and its operation in more detail.

---

## Performing Other Hardening Measures

These procedures explain the steps necessary to implement additional hardening measures.

### Disabling Auto Generation of 8.3 Filenames

This procedure disables the automatic generation of 8.3 format filenames for 16-bit applications.

▶ **To disable 8.3 filename generation**

1. Start Registry Editor (regedt32.exe).
2. Locate and then click the following key:
   **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\FileSystem**
3. If the **NtfsDisable8dot3NameCreation** value does not exist, on the **Edit** menu, click **Add Value**; tType **NtfsDisable8dot3NameCreation** and ensure that it is a REG_DWORD type. Then press **Enter**.
4. Select the **NtfsDisable8dot3NameCreation** value.
5. On the **Edit** menu, click **Modify**.
6. Type **1** (or your desired level), and then click **OK**.
7. Quit Registry Editor.

**Note:** Trey has a number of 16-bit applications that must continue to run normally. However, these applications run on a subset of computers in the domain. All file and print servers have had this change applied; it has been selectively applied to member servers and workstations that are not running 16-bit applications.

## Disabling Autorun

This procedure disables the autorun feature for CD-ROM drives.

▶ **To disable autorun**

1. Start Registry Editor (regedt32.exe).
2. Locate and then click the following key:
   **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\CDROM**
3. On the **Edit** menu, click **Add Value**.
4. Type **Autorun** and ensure that it is a REG_DWORD type. Then press **Enter**.
5. On the **Edit** menu, click **Modify**.
6. Type **0** (or your desired level), and then click **OK**.
7. Quit Registry Editor.

## Removing OS/2 and POSIX Subsystems

This procedure disables the OS/2 and POSIX compatibility subsystems.

▶ **To disable OS/2 and POSIX**

1. Start Registry Editor (regedt32.exe).
2. Locate and delete the following key:
   **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Session Manager\Environment\OS2LibPath**
3. If the **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Session Manager\Subsystems** key exists, locate and delete its OS/2 and POSIX subkeys.
4. If the **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Session Manager\Environment\Subsystems** key exists, locate and delete its POSIX and OS/2 subkeys, if present.
5. Locate and delete the following key:
   **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\OS/2 Subsystem for NT**
6. Quit Registry Editor.
7. Open Windows Explorer.
8. Delete the \winnt\system32\os2 directory and all its subdirectories.
9. Restart the computer.

## Increasing Object Protection Levels

This procedure increases the protection levels for kernel objects.

▶ **To increase object protection levels**

1. Start Registry Editor (regedt32.exe).
2. Locate and then click the following key:
   **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Session Manager**
3. On the **Edit** menu, click **Add Value**.

4. Type **EnhancedSecurityLevel** and ensure that it is a REG_DWORD type. Then press **Enter**.
5. On the **Edit** menu, click **Modify**.
6. Type **1** (or your desired level), and then click **OK**.
7. Quit Registry Editor.

## Preventing Users from Adding Printer Drivers

This procedure disables the ability of normal users to add printer drivers.

▶ **To disable addition of printer drivers by normal users**

1. Start Registry Editor (regedt32.exe).
2. Locate and then click the following key: **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Print\Providers\LanMan Print Services\Servers**
3. On the **Edit** menu, click **Add Value**.
4. Type **AddPrinterDrivers** and ensure that it is a REG_DWORD type. Then press **Enter**.
5. On the **Edit** menu, click **Modify**.
6. Type **1** (or your desired level), and then click **OK**.
7. Quit Registry Editor.

## Confirming That Automatic Administrator Logon Is Disabled

This procedure disables the automatic logon of the Administrator account.

▶ **To disable automatic Administrator logon (which is not on by default)**

1. Start Registry Editor (regedt32.exe).
2. Locate and then click the following key: **HKEY_LOCAL_MACHINE\Software\Microsoft\WindowsNT\CurrentVersion\Winlogon**
3. Locate and then click the following value if it exists: **AutoAdminLogon**.
4. On the **Edit** menu, click **Modify**.
5. Type **0** (or your desired level), and then click **OK**.
6. Delete the **DefaultPassword** value if it exists.
7. Quit Registry Editor.

## Disabling Notifications for the Novell Client

This procedure disables the automatic notification of password changes to the Novell networking client.

▶ **To disable the default password change notifications to the Novell client**

1. Start Registry Editor (regedt32.exe).
2. Locate and then click the following key: **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa**
3. Locate and then click the following value: **Notification Packages**.
4. On the **Edit** menu, click **Modify**.
5. Ensure that the value **FPNWCLNT** is not listed, and then click **OK**.
6. Quit Registry Editor.

# Testing the Solution

After the Trey scenario implementation is complete, you are ready to validate your implementation to ensure that it meets the requirements.

## Validation

You can use the information in the following table to test the Trey scenario and validate your implementation of this guidance.

**Table 4.2: Validation Tests**

| Description | Test steps | Expected result |
|---|---|---|
| Validate installation of Internet Explorer 6 SP1 | Start Internet Explorer. On the **Help** menu, click **About Internet Explorer**. | Version information should show 6.0.2800.xxxx. |
| Validate installation of Windows NT 4.0 (SP6a) | Run the Windows NT Diagnostics applet in the Administrative Tools group and then select the **Version** tab. | Version information should show SP6a. |
| Validate the boot sequence | *In mode1* (store Startup Key locally for most of Trey's servers and all of its workstations).<br>At boot time, the key is decrypted and loaded, allowing the computer to be restarted without administrator intervention.<br>*In mode2* (the **Password Startup** button and associated text fields - for high-value servers).<br>At boot time, the administrator must type the pass phrase into the console to complete the boot. | Client can successfully log on. |
| Validate that DSClient successfully installed | Click **Start**, and then **Search** and **For People**. | Ability to search Active Directory indicates successful installation of DSClient. |
| Validate system policies | Attempt to access resources restricted by system policies.<br>Double check the location of the system policies, which is usually in the C:\Winnt\system32\repl\import\scripts folder. | You cannot access forbidden resources.<br>System policies should be available at given location. |
| Validate NTLMv2 authentication | Set the domain controller to require NTLMv2 authentication. | Client can successfully log on. |
| Validate the security configuration manager | Build and apply the security templates automatically to one or more computers in a domain.<br>Check the set of templates in the %winnt%\security\templates folder. | You are able to secure the member servers and workstations.<br>Templates should be available at given location. |
| Check the long file name creation | Create a new file with a name longer than the 8.3 format. | You are not able to access the file with 8.3 format. |
| Validate the autorun functionality | Insert the autorun-enabled discs into drives on the computer. | Autorun will no longer work when autorun- |

| Description | Test steps | Expected result |
|---|---|---|
| | | enabled discs are inserted into drives on the computer. |
| Validate that the OS/2 or POSIX subsystems will no longer operate | A user starts a process and then logs off, a potential exists that the process will be accessed by the next user who logs on to the system. | Applications that rely on the OS/2 or POSIX subsystems will no longer operate. |
| Validate the object protection levels | Malicious users try to change the attributes of kernel objects under various conditions. | System should not allow malicious callers to escalate their privileges. |
| Prevent users from adding printer drivers | Malicious users try to escalate privileges through adding some printer drivers. | Members of the Printer Operators and Administrator groups are able to add new printer drivers. |
| Validate that the automatic administrator logon is disabled | Try to automatically log on as administrator when the computer is restarted. | Automatic Administrator Logon should be disabled. |
| Validate the boot timeout | Restart and check the boot timeout. | The boot timeout should have changed from 30 to 0. |
| Validate the SCM installation | Client **Start**, **Run**, type **mmc** and press **ENTER.** | SCM Microsoft management console should be available. |

# Summary

Even though Windows NT does not have the full complement of modern security features, there are still many features and techniques to mitigate threats and security risks for your Windows NT systems. From initial operating system and patch baseline installation, through boot hardening, deployment of the Directory Services Client add-in, to the effective use of system policies to control security-critical aspects such as NTLM authentication and password generation, to file system and service hardening, and finally to a number of other hardening measures, there is still a lot that you can do at the operating system level to keep Windows NT as secure as it can be.

## More Information

- Windows NT 4.0 Service Pack 6a is available at www.microsoft.com/ntserver/nts/downloads/recommended/SP6/allSP6.asp.
- The Windows NT Server Service Pack 6a FAQ includes security-specific information, and is available at www.microsoft.com/ntserver/support/faqs/sp6faq.asp.
- The Post-Windows NT 4.0 Service Pack 6a Security Rollup Package is available at www.microsoft.com/ntserver/nts/downloads/critical/q299444/default.asp.
- The Internet Explorer Administration Kit is available at www.microsoft.com/windows/ieak/.
- Active Directory Client Extensions for Windows 95/98 and Windows NT 4.0 are available at www.microsoft.com/windows2000/server/evaluation/news/bulletins/adextension.asp.
- The “Guide to Microsoft Windows NT 4.0 Profiles and Policies” white paper is available for downloading at www.microsoft.com/ntserver/docs/prof_policies.doc.
- Information about the System Policy Editor is available on the Using the System Policy Editor page on MSDN at http://msdn.microsoft.com/library/en-us/policy/policy/using_the_system_policy_editor.asp.
- The Security Configuration Manager is available from the Microsoft FTP server at http://www.microsoft.com/ntserver/techresources/security/securconfig.asp.
- A complete guide to the SCM for Windows NT 4.0 is available for download at www.microsoft.com/ntserver/docs/scm-nt4.doc.
- The National Security Agency's "Guide to Securing Microsoft Windows NT Networks" is available at www.nsa.gov/snac/downloads_winnt.cfm.
- More information about service hardening for Windows NT 4.0 servers is available in Managing Server Security at www.microsoft.com/technet/prodtechnol/winntas/maintain/nt4sac/sacch10.mspx.

# 5

# Hardening Microsoft Windows 98

Most organizations with earlier versions of the Windows operating system have substantial populations of desktop and mobile computer clients running Microsoft® Windows® 98, Windows 98 Second Edition (SE), or Windows Millennium Edition (Me). (This chapter refers to these versions collectively as “Windows 98.”) This chapter focuses on what you can do to improve the security (or *harden*) these clients to improve the overall security of your network.

Windows 98 clients are deployed in many roles in which upgrading the operating system is not feasible. For example, Windows 98 serves as a platform for many kiosks and point-of-sale terminals, customized applications, and classroom student workstations. Proper configuration of security settings on these computers can ensure continued reliability of line-of-business applications without exposure of the workstation itself or other computers on the network.

This chapter discusses how to accomplish the following tasks:

- Install Windows 98 and provide a patch baseline.
- Install an Internet firewall.
- Harden the boot sequence.
- Deploy baseline configurations of Microsoft Internet Explorer.
- Install Microsoft Active Directory® directory service Client Extensions.
- Configure Server Message Block (SMB) signing.
- Choose the Windows NT LAN Manager (NTLM) authentication level.
- Define effective system policies.

# Windows 98 Security Design

Much of the design of a secure Windows 98 installation involves the identification of configuration settings that can be easily modified. In a secure environment, these settings are configured to a corporate specification and locked so that the configuration remains unchanged.

## Installing Windows 98 and Providing a Patch Baseline

A baseline deployment of Windows 98 with current security patches gives a known starting point from which to implement a secure platform policy. You can find a complete list of Windows 98 patches on the Microsoft Windows 98 download site at www.microsoft.com/windows98/downloads/corporate.asp. (Chapter 6, "Patch Management," in this guidance discusses patches in great detail.) Your initial deployment of Windows 98 workstation into a network should include remediation of known vulnerabilities. For example:

- All Windows 98 workstations should be configured using only the options necessary for proper network performance. For example, if local file and printer shares are not required, File and Print Sharing options should not be specified in network configuration.
- Make sure that you have applied all critical updates from the Microsoft Web site to workstations and that you have installed recommended updates if they apply to the local computing environment.
- Microsoft Internet Explorer 6 Service Pack 1 (SP1) incorporates the latest secure browsing enhancements and fixes. You should install it on any computer connecting to the Internet. You can install it directly from the Microsoft Internet Explorer Web site at www.microsoft.com/ie, order it on a CD-ROM, or install it in a customized configuration using the Internet Explorer Administration Kit (IEAK).
- A very efficient method of installing a recent secure baseline configuration is to use the Microsoft Windows Security Update CD, which is available from the Microsoft Trustworthy Computing Web site at www.microsoft.com/security/protect/ or by calling Microsoft Product Support Services.

Trey chose to build its computers from a common image that was applied to new workstations with the Ghost imaging tool. The common image was built by installing Windows 98, Internet Explorer 6.0, and all released patches and service packs, plus content from the Windows Security Update CD.

## Installing an Internet Firewall

A firewall is a security solution that segregates one portion of a network from another portion, allowing only authorized network traffic to pass through according to traffic filtering rules. Firewalls can be either hardware or software-based. An Internet firewall is one that exists between your local computer network and the Internet to protect against malicious attacks by denying access to incoming network traffic that is not specifically approved. In a network environment, the network itself should have a hardware firewall or similar security product such as Microsoft Internet Security and Acceleration (ISA) Server to protect against external threats. You can protect individual workstations by installing software firewall products, which are available from several vendors; see the “More Information” section at the end of this chapter.

## Boot Sequence Hardening

One potential security weakness in Windows 98 is the boot sequence, which can be interrupted to allow system access before policies are in place. To ensure that the system cannot be compromised in this manner, you must secure the boot process by editing the **Msdos.sys** system file with Notepad or another text editor (specific guidance is provided later in this chapter). System administrators should be aware that this setting will disable the ability to boot the computer into safe mode, which is desirable in order to keep malicious users from bypassing security measures. However, this setting will make system troubleshooting more difficult, and will have to be reset in order to use alternate boot sequences for system maintenance.

Because the boot process can be interrupted by booting from removable storage, computer basic input/output system (BIOS) configuration should be set to boot from the primary hard disk only. Most computers allow entry into BIOS settings with a control key during system startup. After setting the BIOS for secure booting, you can secure it further by setting an administrative password. This approach is not perfectly secure; many system BIOS settings can be reset by an “emergency” key sequence that is often published on Web sites, and nearly all can be reset by opening the system case and changing the position of a hardware jumper. In settings where physical security of the computer is uncertain, you should physically lock the system chassis.

Trey engineers modified the boot timeout value for each computer running Windows 98 in the domain. All desktop and mobile computer users were given cable locks to secure their computers. Also, on computers that support it, BIOS settings were changed to restrict alternate booting.

## Deploying Baseline Configurations of Internet Explorer

The Internet Explorer Administration Kit (IEAK) contains many tools for customizing, deploying, and maintaining Internet Explorer 6. With the IEAK, network administrators can identify secure configurations to deploy to network clients. Deployments can include custom applications, pre-built favorites lists, privacy and security settings, specification of proxy server, and almost any other customization of Internet Explorer. Maintaining such configurations allows for easy update of any security components as new Internet vulnerabilities arise.

Internet Explorer allows configuration of security zones to allow or block downloads or active content. The four zones that can be configured are:

- Internet. Contains all Web sites not included in other zones.
- Local intranet. Contains all sites within local networks.
- Trusted sites. Contains Web sites that are trusted not to contain malicious content.
- Restricted sites. Contains Web sites that have the potential to contain malicious content.

You can set customized security configurations within each of these four security zones, or you can select preset safety levels ranging from “low” to “high” from a drop-down list.

To deploy Internet Explorer 6.0, Trey built a custom configuration with the IEAK (a process described in detail in the IEAK documentation, available at www.microsoft.com/windows/ieak) that includes a complete list of Security Zone settings for trusted hosts.

Some malicious code can come from active content that invokes new Internet Explorer sessions or “pop-up” windows. Many pop-ups attempt to trick users into installing Trojan Horses (programs that appear to be useful but contain hidden code to exploit or damage computer systems), viruses (programs designed to replicate themselves on multiple computers), or spyware (programs that conduct certain activities on a computer without obtaining appropriate user consent). A simple way to stem the tide of pop-ups is to install a pop-up blocker. A good one is available for free from Microsoft as part of the MSN Toolbar, which is available at http://toolbar.msn.com. Trey deployed the MSN toolbar as part of its Windows 98 image update.

## Installing Active Directory Client Extensions

Windows 98 clients do not have all of the features of newer operating systems that were designed to take advantage of Active Directory services. Microsoft published the Active Directory Client Extension Add-On (DSClient) for Windows 98 to provide access to Active Directory networks. The following Active Directory features become available to Windows 98 clients through the use of DSClient:

- Awareness of Active Directory sites. This awareness gives the client the ability to log on to the domain controller closest to the client on the network, rather than the primary domain controller (PDC) or PDC emulator role holder. It also gives the client the ability to reset passwords against any domain controller. In Windows NT version 4.0 domains the PDC handles all password changes, but in Active Directory any domain controller can service these requests. DSClient extends this functionality to Windows 98 clients. These enhancements help to reduce network traffic and load on the PDC.
- NTLM version 2 (NTLMv2) authentication. NTLMv2 authentication gives much more secure authentication than the LAN Manager (LM) authentication that shipped with Windows 98. While not as strong as Kerberos authentication, NTLMv2 is far more secure than LM.
- Active Directory Services Interface (ADSI). ADSI provides a common application programming interface (API) to applications and allows a scripting interface for Active Directory.
- Distributed File System (DFS) fault tolerance. DSClient gives access to Windows 2000 and Microsoft Windows Server™ 2003 DFS failover shares as specified in Active Directory.
- Active Directory Windows Address Book properties. DSClient extends the Windows 98 environment to expose extended Active Directory schema elements through the **Search** command on the **Start** menu. It also allows users who have permissions to edit properties on user objects within Active Directory.

The following Active Directory features are *not* made available through the Active Directory Client Extensions:

- Kerberos support. Full Kerberos support is available only on Windows 2000 and later clients.
- Group Policy support. Group Policy participation and IntelliMirror object management are not made available to the clients with earlier-version operating systems.
- Internet Protocol Security (IPSec) and L2TP support. These advanced secure networking protocols are not available.
- Service Principle Name (SPN) or mutual authentication. These capabilities are not enabled through DSClient.

It is important to obtain the latest version of DSClient from Microsoft Product Support Services. DSClient 2003 is available as a hotfix; more information can be obtained from KB article 323455, "Directory Services Client Update for Windows 98" at http://support.microsoft.com/?id=323455. Before installing DSClient, ensure that your workstations are running Internet Explorer 6 with SP1 or later.

The Trey IT staff manually deployed DSClient to each of the computers running Windows 98 in the domain. They did this to allow integration with Kerberos authentication and to allow the computers running Windows 98 to use a higher NTLM authentication level.

## Configuring SMB Signing

SMB signing is a cryptography technique that allows each packet sent between a client and server to be digitally signed to verify authenticity. This technique prevents client or server impersonation in the network by computers that may attempt to interject themselves in the middle of communications, and it verifies the source of all network communications.

SMB signing was introduced with Windows NT 4.0 Service Pack 3 (SP3) and is described in KB article 161372, "How to Enable SMB Signing in Windows NT" at http://support.microsoft.com/?id=161372. To enable it in Windows 98, you must make a registry **DWORD** entry in the **HKLM\SYSTEM\CurrentControlSet\Services\VxD\VNetsup** key to either require signing or support it if the communications partner requires it. There are two **DWORD** values that together control the use of SMB signing:

- If you set **EnableSecuritySignature** to **1** and **RequireSecuritySignature** to **0**, SMB signing will be used if the client and server both support it. This setting allows the opportunistic use of signing but does not prevent the client from connecting to other clients or servers that do not support signing.
- If you set **RequireSecuritySignature** to **1** and **EnableSecuritySignature** to **0**, the client will only communicate with servers that support SMB signing.

SMB signing, when used, should be configured on all computers that participate on a network. Computers that do not have these registry entries will not be able to communicate with other network hosts. The overhead for SMB signing typically results in a 10 to 15 percent decrease in network performance.

Trey chose to disable the use of SMB signing for its Windows 98 clients to provide complete compatibility with their existing environment. SMB signing is enabled, but not required, for servers and domain controllers, as well as for Windows NT and Windows 2000 clients. On their Windows 98 clients, the client settings used were **EnableSecuritySignature=0** and **RequireSecuritySignature=0**, which prevents Windows 98 clients from requesting or accepting signed connections. Although this approach denies those clients the additional security against spoofing and man-in-the-middle attacks, turning off SMB signing preserves compatibility, which was evaluated as being more important to Trey business operations than the additional security. This change also required Trey to make a configuration change on their Windows Server 2003 domain controllers, since Windows Server 2003 enables SMB signing by default.

## Choosing the NTLM Authentication Level

Windows 98 uses older, less secure authentication LM encryption by default. Vulnerabilities and exploits have been published against them, and Microsoft has strengthened the authentication security protocols to mitigate these vulnerabilities. After

the DS Add-On client has been configured, you can set it to use the more secure NTLMv2 authentication method.

With DSClient installed, Windows 98 supports two levels of NTLM and NTLMv2 authentication that are controlled by the **LMCompatibility** registry value described later in this chapter. These values are:

- **0** (Send LM & NTLM responses). Offers the most interoperability. Clients may use LM or either version of NTLM to authenticate.
- **3** (Send NTLMv2 response only). Use this value only if all clients with earlier operating system versions have DSClient installed.

Trey initially deployed the registry key with a value of **0**, which mirrored the existing environment. After the Windows NT 4.0 servers at Trey had been upgraded as described in Chapter 4, "Hardening Microsoft Windows NT 4.0," Trey reset the **LMCompatibility** value on computers running Windows 98 to **3**.

The default installation of NTLMv2 encryption provides 56-bit key lengths on systems where the 56-bit version of Internet Explorer is installed. Systems where Internet Explorer was installed after 1999 probably have the 128-bit version; older clients can be upgraded to 128-bit encryption as described in the preceding chapter. If the 128-bit version of Internet Explorer is installed before you install the DSClient, 128-bit NTLMv2 authentication will be enabled. As described in Chapter 4, “Hardening Microsoft Windows NT 4.0,” Trey installed the 128-bit update on its computers running Windows 98 and deployed NTLMv2 authentication support for all its computers. After these changes were in place, Trey was able to proceed with enabling NTLMv2 support on servers and domain controllers.

## Defining Effective System Policies

System policies let you centrally apply security policies to overwrite default settings in the local computer registry. Network administrators can identify areas of vulnerability within computers running Windows 98 and configure many of the settings to be as secure as possible.

Windows 98 clients apply policies in the **Config.pol** file located in the PDC’s Netlogon share (because computers running Windows 98 can only enumerate a domain user’s group memberships from the PDC, not from any backup domain controllers (BDCs)). Because the domain controllers at Trey are running Windows Server 2003, this is not a problem for the environment; sites that are still running Windows NT® 4.0 domain controllers may be able to use the recommendations in KB article 150687, “Group Policies Not Applied on Windows NT Domain” at http://support.microsoft.com/?id=150687 to deploy user-specific policies.

Most of the Windows 98 policies are oriented toward restricting the user’s ability to change the desktop environment. Trey elected not to use these policies because they add little effective security. Instead, Trey chose to apply policy settings that would make it more difficult for a malicious attacker to cause damage, or for an innocent but untrained user to accidentally break needed functionality. Trey chose to apply policies to:

- Require logon security, and to present a logon banner that describes organizational policies.
- Set a password policy that hides user passwords while they are typed and requires long alphanumeric passwords.
- Disable file and print sharing and remote dial-in access.
- Prevent users from running the registry editing tools.

Remember that a mistake can quite easily secure the computer *too* much and lock out functionality necessary for using or administering the computer. Therefore, as a best practice, use a computer that is not a primary workstation and can be reconfigured. It is also a good idea to make group-specific or user-specific policies for administrators that relax some of the restrictions so that administrators can easily access registry editing and troubleshooting tools. Trey developed a test schedule that deployed proposed policy settings in a lab environment with computers built to accurately represent its production hosts; this test schedule was used to verify that the policies worked as intended without unintended side effects.

# Implementation

Windows 98 lacks most of the configuration and management tools introduced in later versions of Windows. This limitation meant that the Trey Research staff was forced to choose between building a secure Windows 98 configuration and deploying it via imaging to all their current workstations, or manually applying security settings. Because they were already planning a deployment of Windows XP Service Pack 2 for all computers as part of their IT modernization plan, they elected to use manual configuration to avoid having to rebuild affected computers twice, even though it means an additional degree of complexity in the short term.

## Implementation Prerequisites

For these implementation details to work correctly, you must have a basic Trey Research infrastructure implemented as introduced in Chapter 2, "Applying the Security Risk Management Discipline to the Trey Research Scenario."

## Implementation Overview

Implementing this solution scenario will involve performing the following activities:

- Installing Windows 98 and providing a patch baseline.
- Installing an Internet firewall.
- Boot sequence hardening.
- Deploying Internet Explorer.
- Installing the Active Directory Client Extensions for Windows 98.
- Configuring SMB signing.
- Choosing the NTLM authentication level.
- Defining effective system policies.

## Installing Windows 98 and Providing a Patch Baseline

Trey developed a standard Windows 98 configuration by reinstalling Windows 98 with its standard defaults on a test computer, installing patches from the Security Update Kit, and then adding the most up-to-date set of patches from Windows Update. After the installation was complete, Trey used the Windows Update catalog tool to analyze the test system and print a report listing the patches that were applied. This list of patches was then used to update other Windows 98 systems to the same baseline.

## Installing an Internet Firewall

After evaluating a number of personal firewall products, the Trey IT director chose one that allowed centralized configuration and reporting of attempted attacks and penetration. This firewall product was then deployed to all desktop and mobile computers running Windows 98. Microsoft does not recommend or endorse specific firewall products, but some available products are listed in the “More Information” section at the end of this chapter.

## Boot Sequence Hardening

To secure a computer running Windows 98 against interruption of operating system startup, you must modify the Msdos.sys system file and configure the computer BIOS to remove access to removable media as boot devices.

## Securing the Boot Process

To keep the boot sequence from being interrupted before security policies are enforced, you should edit the **Msdos.sys** system file to disable the ability to change startup behavior and circumvent policies. KB article 118579, "Contents of the Windows Msdos.sys File" at http://support.microsoft.com/?kbid=118579 explains how to locate and edit this file.

Because the **Msdos.sys** file is hidden and marked Read-Only, you should modify it to remove these attributes until the file is edited.

▶ **To modify Msdos.sys with Notepad**

1. Click **Start**, point to **Find**, and then click **Files Or Folders**.
2. In the **Named** box, type **msdos.sys**.
3. In the **Look In** box, click your boot drive (usually drive C).
4. Click the **Find Now** button.
5. Right-click the **Msdos.sys** file and select **Properties**.
6. Clear the **Read-Only** and **Hidden** check boxes to remove these attributes from the file, and then click **OK**.
7. Right-click the **Msdos.sys** file and select **Open With**.
8. In the **Choose the program you want to use** box, click **Notepad**, and then click **OK**.
9. Add the following two lines to the [**Options**] section:

   BootKeys=0

   BootSafe=0

   The **BootKeys** Boolean value specifies whether keyboard function keys are allowed on system startup. Because several of these keys can be used to interrupt the boot process, a secure system disables the keys by assigning a value of **0**.

   **BootSafe** is another Boolean setting that allows safe-mode booting. Setting **BootSafe** to a value of **0** locks the computer from a safe-mode boot.
10. Save the file and close Notepad.
11. Right-click the **Msdos.sys** file and select **Properties**.
12. Select the **Read-Only** and **Hidden** check boxes to set the attributes for the file, and then click **OK**. Close the **Find** dialog box.
13. Reboot the computer for the changes to take effect.

---

**Note:** Further discussion of the contents of Msdos.sys is available in the KB article referenced earlier.

---

### Removing Access to Removable Media As Boot Devices

If a computer can be booted from removable media, system security settings can be completely bypassed and reconfigured. Consult your system manufacturer's guidance for instructions for accessing the system BIOS.

▶ **To disable booting from removable media**

1. Set the primary hard disk as the first boot device.
2. Disable booting from floppy disk and CD-ROM devices.
3. Consider disabling universal serial bus (USB) and FireWire ports if not needed in your business environment.
4. Set the BIOS password (if available) to prevent these security measures from being reset.

## Deploying Internet Explorer

Administrators of large networks can build custom installations of Internet Explorer 6.0 SP1 with the Internet Explorer Administration Kit (available at www.microsoft.com/windows/ieak) to ensure that workstations are running the latest secure build of Internet Explorer. IEAK allows administrators to set administrative profiles to preconfigure Internet Explorer security settings, lock down Microsoft NetMeeting® and Microsoft Outlook® Express, and control which features users are able to change.

---

**Note**: The most up-to-date version of the Active Directory Client Extensions requires Internet Explorer 6.0, as described in KB article 555038 "How to enable Windows 98/ME/NT clients to logon to Windows 2003 based Domains" at http://support.microsoft.com/?kbid=555038.

---

## Installing Active Directory Client Extensions for Windows 98

Microsoft has created extensions for Windows 98 to allow participation in Active Directory domains. This client should be installed on all Windows 98 workstations in these environments. Although the Active Directory Client Extensions for Windows 98 were distributed with Windows 2000, a new update is available from Microsoft Product Support Services as a free hotfix.

## Configuring SMB Signing for Network Communications

SMB signing ensures that each packet transmitted across a network is digitally signed, which provides a high level of security but may incur a network performance cost of 10-15 percent. If SMB signing is configured, *all* systems in the network should be configured to use SMB signing. However, to ensure maximum compatibility at the expense of some security, Trey elected to force its Windows 98 clients to disable SMB signing. Their configuration can be implemented as follows.

▶ **To disable SMB signing on the Windows 98 client**

1. Start the Registry Editor by typing **Regedit.exe** at a command prompt and pressing **ENTER**.
2. Find the key **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\VxD\VNetsup**

3. Add two values to this key:
   - Value Name: EnableSecuritySignature
     Data Type: **REG_DWORD**
     Value: **0** (disables the use of signing when the server supports it)
   - Value Name: RequireSecuritySignature
     Type: **REG_DWORD**
     Value: **0** (allows communication even when the server cannot support signing)
4. Exit the Registry Editor.
5. Restart the computer.

▶ **To change the Windows Server 2003 default setting that requires SMB signing**

1. Log on to the Windows Server 2003 domain controller using an account with administrative privileges on the domain.
2. Launch the Microsoft Management Console (MMC.exe) and add the Group Policy Object Editor snap-in. Target the Group Policy Object Editor snap-in at the **Default Domain Controllers Policy** object for the domain, and then click **Finish**.
3. Expand the Default Domain Controllers Policy object, then expand **Computer Configuration \ Windows Settings \ Security Settings \ Local Policies \ Security Options**.
4. Double-click **Microsoft network server: Digitally sign communications (always)**.
5. Select the **Define this policy setting** checkbox, and then ensure that the **Disabled** button is selected. Click **OK**.
6. Close the Microsoft Management Console window.

## Choosing the NTLM Authentication Level

After the Active Directory Client Extensions have been installed, you should enable NTLMv2 authentication. KB article 239869, "How to enable NTLM 2 authentication," provides guidance on enabling these settings.

▶ **To set the NTLMv2 authentication level**

1. Start the Registry Editor by typing **Regedit.exe** at a command prompt and pressing **ENTER**.
2. Locate and click the following key in the registry:
   **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control**
3. Create a new subkey of **Control** named **LSA**.
4. Create a **DWORD** value called **LMCompatibility** and set it to 3 (see “Choosing the NTLM Authentication Level” earlier in this chapter for a discussion of permissible values).
5. Restart the computer.

---

**Note**: You should not set the value of LMCompatibility to 3 *until after you have updated your Windows NT servers to allow the use of NTLMv2*. If the environment contains computers that are not configured to use NTLMv2, communications will fail. Ensure that you plan these settings across your enterprise to ensure uninterrupted communication. For more information about restricting the LM compatibility level on Windows NT, see Chapter 4, "Hardening Microsoft Windows NT 4.0."

---

# Configuring System Policies for Security

System Policy Editor is a powerful tool that can be used to limit access to a computer running Windows 98 in a precise manner by preventing users from changing security settings while allowing them to accomplish their work. You should create system policies to help protect your workstations from tampering and enforce custom security settings.

## Installing System Policy Editor

You must install the System Policy Editor on the same platform for which you want to create policies. That is, if you want to make policies for Windows 98 systems, you must edit and save the policy file from a Windows 98 workstation.

▶ **To install the System Policy Editor from the Windows 98 CD-ROM**

1. Click **Start**, point to **Settings**, and then click **Control Panel**.
2. Click **Add/Remove Programs**.
3. Click the **Windows Setup** tab, and then click **Have Disk**.
4. Insert the Windows 98 disk into your CD-ROM drive.
5. In the **Install From Disk** dialog box, browse to the \Tools\Reskit\Netadmin\Poledit folder on the CD, select the **Poledit.inf** file, and then click **OK**.
6. In the **Have Disk** dialog box, select the System Policy Editor component and then click **Install**.

This installation copies Poledit.exe to the Windows folder and Windows.adm, Common.adm, and Poledit.inf to the \Windows\Inf folder. It also makes the necessary changes to your registry and adds a start menu item to your Programs\Accessories\System Tools folder.

---

**Warning:** Before you edit the registry, you should first make backup copies of the registry files (System.dat and User.dat), which are hidden files in the Windows system folder.

---

## Recommended Policy Settings

After installing the System Policy Editor, you can use it to create a policy for local installation or distribution. Notice that as you change policy settings, the check boxes for each policy setting have three states: selected, cleared, and disabled (dimmed). If the check box is selected, the policy will apply as specified. If the check box is dimmed, the policy setting will be ignored. If the check box is cleared, that policy's registry settings might be deleted unintentionally.

▶ **To create a policy using recommended computer-level settings**

1. Double-click C:\Windows\poledit.exe.
2. Click **File**, and then **New Policy**.

3. Double-click the **Default Computer** icon. The **Default Computer Properties** dialog box will display.
4. Expand the **Windows 98 Network** node. The following settings are recommended for the network properties.

   Under **Logon,** select the following check boxes:

   - **Logon Banner.** Adds a start-up banner that describes organizational policy about computer usage.
   - **Require validation from Network for Windows Access**. Requires users to authenticate over the network rather than at their local computers.
   - **Do not show last user at logon**. Forces users to type in a valid user name, instead of displaying the previous value.
   - **Do not show logon progress**. Hides progress of the logon session.

   Under **Password**, select the following check boxes:

   - Hide share passwords with asterisks. Masks passwords as they are typed.
   - **Disable password caching**. Forces the client to authenticate against a more secure network controller rather than storing passwords locally in a less-secure cache.
   - **Require alphanumeric Windows password**. Forces a higher level of complexity for local passwords.
   - **Minimum Windows Password length**. Lets you specify the number of characters required for a password. A setting of **8** is a relatively strong password length.

   Under **Microsoft Client for Windows Networks**, select the following check boxes:

   - **Log on to Windows NT**. Lets the administrator hard-code the domain to which the workstation is authorized to log on.
   - **Disable caching of domain password**. Minimizes the exposure of locally cached passwords.
   - **Workgroup.** Allows the administrator to hard-code the domain to which the workstation will log on. This setting is required for domain logon.
   - **Clear the Alternate Workgroup** check box so that the user cannot log on to specified workgroups.

   Under **File and Print Sharing for Microsoft Networks**, clear all check boxes. File and print sharing capabilities should generally be disabled on local workstations. If users need to share files or printers, use dedicated servers and secure them properly.

   Under **Dial-Up Networking**, select the **Disable dial-in** check box to ensure that the computer cannot be accessed remotely.
5. Click **OK**.

▶ **To create a policy using recommended user-level settings**

This procedure requires that you complete the previous procedure and that the System Policy Editor is already open. Windows 98 only supports downloading a single policy from the domain controller, so user-level and computer-level settings must be combined.

1. Double-click **Default User**.
2. In the **Default User Properties** dialog box, double-click **Windows 98 System**.
3. Double-click **Control Panel**.

4. Expand **Network**, and then click **Restrict Network Control Panel**.
5. Expand **System**, and then click **Restrict System Control Panel**.
6. Expand **Restrictions**, and then click **Disable Registry editing tools**.
7. Click **OK**.
8. Save the policy to the appropriate location.

## Deploying Policies

After you configure the policies for your organization, complete the following steps to name the policy file Config.pol and save it to the correct network location so that the client workstations can automatically download and apply the settings.

▶ **To deploy policies**

1. On the **File** menu, select **Save As**.
2. Name the file **Config.pol** and store it in one of the following locations:
   - For Windows NT 4.0 domain controllers, save the file as %systemroot%\WINNT\System32\Repl\Import\Scripts\Config.pol
   - For Windows 2000 and Windows Server 2003 domain controllers, save the file as %systemroot%\sysvol\sysvol*domainName*\scripts\Config.pol

▶ **To enable automatic policy downloading for the Windows 98 client**

1. Log on to the computer running Windows 98.
2. Open the Control Panel.
3. Double-click **Network**.
4. Ensure that in the **Primary Network Logon** drop-down list, **Client for Microsoft Networks** is selected.
5. Click the **Identification** tab, ensure that the value in the **Workgroup** field matches the name of the domain, and then click **OK**.

# Testing the Solution

After the scenario implementation is complete, you are ready to validate your implementation to ensure that it meets the requirements.

## Validation

You can use the information in the following table to test the Trey scenario and validate your implementation of this guidance.

**Table 5.1: Validation Tests**

| Description | Test steps | Expected result |
|---|---|---|
| Validate installation of hotfixes | Run the QFECheck.exe utility (located in the Windows installation folder). | You are able to obtain a list of current hotfixes that match established baseline. |
| Validate installation of Internet Explorer 6 SP1 | Start Internet Explorer, and from the **Help** menu click **About Internet Explorer**. | Version information should show 6.0.2800.xxxx. |
| Validate that DSClient successfully installed | Click **Start**, and then **Search** and **For People**. | Ability to search Active Directory indicates successful installation of DSClient. |
| Validate NTLMv2 authentication | Set the domain controller to require NTLMv2 authentication. | Client can successfully log on. |
| Validate SMB signing | Set network resources to require SMB signing for communications. | Client can successfully access network resource. |
| Validate system policies | Attempt to access resources restricted by system policies. | You cannot access forbidden resources. |
| Validate system BIOS security | Attempt to access system BIOS with manufacturer-specified escape sequence. | You are presented with a password challenge. |
| Attempt to circumvent boot device | Insert a bootable floppy disk and CD. | Client does not boot from removable media. |
| Attempt to circumvent boot sequence | Press F5 or F8 during system startup. | Client does not present a menu allowing alternate boot. |
| Attempt to bypass network logon | Attempt to press ESC when presented with logon. | Windows 98 desktop not accessible until successful domain credentials are presented. |

# Summary

Many organizations have significant investments in Windows 98 systems and, for various reasons, are not able to upgrade them to newer operating systems with stronger built-in security. Given proper attention, these systems can be made relatively secure from vulnerabilities that otherwise could impact an organization. Best practices dictate a proactive security management stance and threat mitigation by staying current with security patches for operating system and applications after a well thought-out initial installation and configuration. The Active Directory Client Extensions for Windows 98 allow these workstations to obtain some of the benefits of participating in an Active Directory domain by providing secure authentication with NTLMv2. Also, man-in-the-middle attacks from computers posing as valid clients can be stopped using SMB signing. In addition, you can use Windows 98 System Policy Editor to implement strong policies that allow knowledge workers to perform their job functions while securing their computers from the chance of accidental or intentional misconfiguration.

After a strong foundation is in place, you can protect your investment with strong patch-management practices and antivirus technology, covered in Chapters 6, "Patch Management," and Chapter 7 "Antivirus Protection," respectively, of this guidance.

## More Information

- For a current list of security fixes for Windows 98, see the Microsoft Windows 98 download site at www.microsoft.com/windows98/downloads/corporate.asp.
- The Windows Security Update CD is available from the Microsoft Trustworthy Computing Web site at www.microsoft.com/security/protect/.
- For more information about Internet Explorer, see the Microsoft Internet Explorer Web site at www.microsoft.com/windows/ie/.
- For more information about IEAK, see the Internet Explorer Administration Kit Web site at www.microsoft.com/windows/ieak/.
- For more information and to download the Active Directory Client Extensions, see "Active Directory Client Extensions for Windows 95/98 and Windows NT 4.0" at www.microsoft.com/windows2000/server/evaluation/news/bulletins/adextension.asp.
- For a list of Internet firewall software vendors, see "Use an Internet Firewall" at www.microsoft.com/security/protect/windows2000/firewall.asp.
- For more information on system policies, see "Chapter 8, System Policies" of the *Microsoft Windows 98 Resource Kit* at www.microsoft.com/resources/documentation/windows/98/all/reskit/en-us/part2/wrkc08.mspx.

# 6

# Patch Management

One of the main ways you can guard against attack is to ensure that your environment is kept up to date with all of the necessary security patches. Patches are required at the server and client levels. This chapter shows you how to ensure that you find out about new patches in a timely manner, implement them quickly and reliably throughout your organization, and monitor to ensure that they are deployed everywhere. It describes the compromises in patch management implementations and concludes with a detailed description of the patch management system at Trey Research.

## Background

Patch management is a critical component of information security. When new vulnerabilities are identified in existing code or when new threats emerge, vendors release patches to fix vulnerabilities or add security features. Organizations like Trey Research must be able to quickly identify which computers need which patches, and then deploy the patches as necessary. They must be able to do this in a consistent and repeatable manner, because failing to patch even a few computers means that the overall network is still vulnerable.

# Solution Design

Exactly how you implement patch management will depend on the size and complexity of your organization. However, it is vital that you understand the importance of patch management and how it fits into the overall risk management strategy for your organization.

For example, if you decide that risk must be minimized at all costs, you could follow a strategy of shutting down all production systems every time a new vulnerability appears in your software. You may then choose not to start the systems again until you have done extensive testing on the security patch and deployed it throughout your organization. This is a very time-consuming and expensive process and would be impractical for most organizations.

Throughout the patch management process, you will need to evaluate the risks against the costs of deploying the appropriate countermeasures. After a security vulnerability has been disclosed, there may be a short period of time before a patch is released. You will need to evaluate the increased risk caused by the vulnerability and determine what must be done prior to testing and deploying a patch.

You may need to disable services, take systems offline, or restrict access to internal users or other groups as necessary. After a patch is released, you need to determine the risk of deploying it immediately against the cost of keeping services down or unprotected while you test and make sure that the patch does not negatively affect the system. If you decide to test, you need to determine how much testing you can afford to do before the risks of not deploying outweigh those of deploying.

---

**Note:** Your organization should implement a change management process. Microsoft Operations Framework (MOF) includes a change management process that can serve as the foundation process for your organization. For details on MOF, see the "More Information" section at the end of this chapter.

---

## Solution Prerequisites

The Trey Research patch management solution must allow the IT department to:

- Automatically inventory computers to identify which ones have been patched and which ones have not.
- Quickly perform on-demand inventories to immediately identify computers that are missing specified patches.
- Automatically perform scheduled scans for patch compliance, generating reports that can be maintained and tracked over time.
- Install single or multiple patches on selected computers without user intervention, minimizing the number of reboots whenever possible.

Because Trey maintains offices in several locations, each with its own administrative staff, the company needs to adopt a uniform set of tools and processes for patch management. Although this is primarily a process issue (addressed by the solution design described in a subsequent section), there are some technical issues that impact the Trey solution design:

- The Microsoft® Software Update Service (SUS) installation used by Trey computers that run the Microsoft Windows® XP or Windows 2000 operating system does not support Windows 98 or Microsoft Windows NT® version 4.0. The existence of these older operating systems limits their ability to automatically download new security updates and push them to client computers.
- Computers running Windows 98 and Windows NT can use the public, Microsoft-hosted Windows Update service to download patches and updates, but there is no way to use system policy settings to restrict the set of updates that users download. In addition, it is not possible to make these clients use an internal SUS or Windows Update Services server.

## Solution Architecture

Patch management architectures vary widely among organizations. Some organizations opt for a highly centralized and tightly managed system using as many off-the-shelf tools as possible, while others build customized patch inventory and deployment tools. However, all of these architectures share some common features and requirements. The basic patch management process consists of four phases:

- Assess the environment to determine which computers exist, how patches can potentially be distributed to them, and what business processes influence how and when patches are applied. To do this, you must look at the current environment and evaluate potential threats (as described in Chapter 2 of this guide). When you complete this evaluation, you will be able to determine the patches that you must deploy to reduce the threats to your environment.
- Identify which specific computers need which specific patches. This process may be done automatically with a tool like the Microsoft Baseline Security Analyzer or Systems Management Server, manually, or with environment-specific scripts that check patch revisions.
- Evaluate and plan the patch deployment, including testing the patch, making provisions for rolling back failed patch installations, and deciding which patches are important enough to merit immediate application. In addition, you must identify who will perform the patch testing and deployment.
- Deploy the patches to systems that need them, verifying that all needed patches were applied and that systems were rebooted when necessary.

---

**Important:** It is highly recommended that you back up all production systems prior to deploying patches.

---

## Assessing the Environment

For patch management purposes, your IT staff needs to know at least the following information:

- What systems are in the environment, including:
  - Operating systems and their versions.
  - Patch levels in use (service pack versions, hotfixes, and other modifications).
  - Functions performed by the systems.
  - Applications in use throughout the environment.
  - Contact information on the individuals or groups responsible for maintaining each system.
- What assets are present and their relative values.
- What the known threats are, and the processes in place for identifying new ones or changes in threat level.
- What the known vulnerabilities are, and the processes in place for identifying new ones or changes in vulnerability level.
- What countermeasures have been deployed.

It is highly recommended that you keep this information available to all those involved in your patch management process and ensure that it is kept up to date. After you know your assets, vulnerabilities, threats, and how your environment is configured, you can determine and prioritize which of the threats and vulnerabilities are going to be of greatest concern to you.

If you follow the guidance given in Chapter 2, "Applying the Security Risk Management Discipline to the Trey Research Scenario," much of this information will already be available as you plan your patch management deployment. Information about the specific patches in use can be gathered using the steps described in the following sections. Data about applications, assets, risks, and countermeasures can be adapted from the SRMD materials and the recommendations in previous chapters of this guidance.

## Identifying Patching Requirements

As an ongoing process, you need to ensure that your computer patches are up to date. In some cases, a new patch will be released that you will need to install on all your servers. In others, a new server that is brought online will need to be patched appropriately. You should continue to analyze all of your servers to ensure that they are completely up to date with all of the latest patches. To efficiently keep the computers in your organization up to date, you need to know what vulnerabilities exist and what protection is already in place. Tools you can use to assist in this process include the Microsoft Baseline Security Analyzer (MBSA), Microsoft Systems Management Server (SMS) version 2.0 and SMS 2003, and the Office Update Inventory tool.

### Using the Microsoft Baseline Security Analyzer

Although it is critical to know which patches have been applied to your system, it is more important to know which patches have *not* been applied. MBSA was designed to scan computers that are running Windows NT 4.0, Windows 2000, Windows XP Professional, and Windows XP Home Edition, and to generate reports about which patches are present and which ones are required.

**Note:** MBSA can be executed from any computer that is running Windows 2000 Professional, Windows 2000 Server, Windows XP Home, or Windows XP Professional. It cannot run on Windows 98 or Windows NT 4.0, and it does not scan computers running Windows 98.

The MBSA tool performs its scan by referring to an Extensible Markup Language (XML) database that Microsoft constantly updates. It also uses the popular “HFNetChk” tool that Microsoft released in August 2001. The XML file contains security bulletin names and titles and detailed data about product-specific security hotfixes, including files in each hotfix package and their file versions and checksums, registry keys that were applied by the hotfix installation package, information about which patches supersede other patches, related Microsoft Knowledge Base (KB) article numbers, and much more.

When you run the MBSA tool for the first time, it must obtain a copy of this XML file so that it can find the hotfixes that are available for each product. The XML file is available on the Microsoft Download Center Web site in compressed form (digitally signed .Cab file). MBSA downloads the .Cab file, verifies the signature, and then decompresses the .Cab file to the local computer on which MBSA is running. Note that a .Cab file is a compressed file that is similar to a WinZip (.Zip) file.

**Note:** Each time you run MBSA it attempts to connect to the Internet to download the XML file from Microsoft. If an Internet connection is not available, the tool will look for a local copy of the XML file in the tool installation folder. Each time the file is successfully downloaded during a scan, a local copy will be stored on the computer in the event that subsequent scans fail to connect to the Internet. Otherwise, for computers that never connect to the Internet, users can separately download this file from the Microsoft Download Center site and copy it onto the computers running the tool.

After the .Cab file is decompressed, MBSA scans your computer (or the selected computers) to determine the operating system, service packs, and programs that are running. MBSA then parses the XML file and identifies security patches that are available for your combination of installed software.

For MBSA to determine whether a specific patch is installed on a given computer, three items are evaluated: the registry key that is installed by the patch, the file version, and the checksum for each file that is installed by the patch.

In the default configuration, MBSA compares both file details and registry keys from the resulting XML subset to the files and registry details on the computer that is being scanned. If any of the file or registry key details on the computer do not match the information that is stored in the XML file, the associated security patch is identified as not installed, and the results are displayed in the security report. The specific KB article number that relates to the patch is also displayed on the screen.

In general, the MBSA tool scans for security issues in the Windows operating systems (Windows NT 4.0, Windows 2000, and Windows XP), such as Guest account status, file system type, available file shares, members of the Administrators group, and so on. Descriptions of each operating system check are shown in the security reports, along with instructions about how to fix any issues that are found.

**Note:** To use MBSA, you must have either local administrator or domain administrator access to the computer being checked for patches.

The MBSA tool has a number of command line switches that can be used in two modes: MBSA mode and HFNetChk mode. The MBSA-style scan will store results (as was done in MBSA version 1.0) in individual XML files to later be viewed in the MBSA user interface (UI). MBSA-style scans include the full set of available Windows, Internet Information Services (IIS), Microsoft SQL Server™, desktop applications, and security update checks.

The HFNetChk-style scan will check for missing security updates and will display scan results as text in the command line window, as is done in the standalone HFNetChk tool. MBSA 1.1 includes the "/hf" flag that will indicate an HFNetChk scan to the MBSA engine. In addition, in HFNetChk mode, MBSA can scan a list of computers supplied in a text file, and it can check systems to see whether they have all the patches published by a named Software Update Service (SUS) server.

If you are using MBSA to verify your patch status, you should ensure that it is run regularly. In most environments, the best way to do this is to schedule it to run at pre-set intervals.

---

**Note:** For more information about using the MBSA tool, see the Microsoft Baseline Security Analyzer page on Microsoft TechNet at www.microsoft.com/technet/security/tools/mbsahome.mspx.

---

## Office Update Inventory Tool

With the powerful application programming capabilities built into Microsoft Office, it has become important to pay attention to vulnerabilities within the applications themselves. Many viruses and Trojan horses take advantage of the ability of today's productivity applications to act upon active content within documents, spreadsheets, and e-mail communications. To help keep Office deployments updated and secure, Microsoft released the Office Update inventory tool. This utility can be executed on computers running the Windows 98 operating system or later and Office 2000 or later. It allows administrators to accurately assess the patch levels of their Office deployments.

The Office Update Inventory Tool is available at http://www.microsoft.com/office/ork/2003/journ/offutoolv2.htm.

## Other Methods for Determining Hotfix Levels

If you do not want or are unable to use the MBSA tool in some parts of your environment, there are other ways that you can determine whether hotfixes have been installed.

The easiest way to do this is to look in the computer's registry under the **HKLM\Software\Microsoft\Windows NT\Currentversion\hotfix** key. Every new hotfix installed should have a key with a Q name that corresponds to the KB article that discusses the hotfix. However, this is not the case for some older hotfixes and for hotfixes for some particular applications.

### Other Tools for Determining Hotfix Levels

There are two other free tools from Microsoft that you can use to gather this information. These tools are:

- Qfecheck.exe used with the /v switch. This tool is discussed in Microsoft Knowledge Base article 282784, "Qfecheck.exe Verifies the Installation of Windows 2000 and Windows XP Hotfixes" at http://support.microsoft.com/?kbid=282784. The tool informs you of service pack levels and hotfix versions installed on your servers. Qfecheck will also advise if a patch was not correctly installed in your environment.
- Hotfix.exe used with the –l switch. This tool is discussed in the Microsoft Windows NT and Windows 2000 Hotfix Installation and Deployment Guide at www.microsoft.com/technet/archive/security/tools/tools/hfdeploy.mspx. The tool displays the number and versions of hotfixes installed on your computers.

## Evaluating and Planning Patch Application

Not every threat or vulnerability poses a significant risk to your environment. As you read notifications of potential new operating system or application vulnerabilities, you should assess whether these vulnerabilities apply to your particular environment.

For example, if the vulnerability applies to the File Transfer Protocol (FTP) service in Windows 2000 and you never enable this service, then the vulnerability does not apply to you. Similarly, if you learn that there is an increased chance of hurricanes this year but your IT environment is inland, then this threat is minimal. If you respond to threats and vulnerabilities that do not apply to your environment, you will use up valuable resources and, potentially, adversely affect the stability of your environment with no corresponding benefit.

As new threats and vulnerabilities emerge, you should read any supporting information about them. This will allow you to make an informed decision about the level of significant risk to your environment and then determine the appropriate response. Your primary alternatives will be to take no action, to disable the service at risk, or to deploy a patch.

---

**Important:** When creating the plan for deploying a new patch, you should also create a rollback plan describing how to remove the patch or mitigate failures in patch installation.

---

To ensure that you stay current on new patches, make sure that you receive regular security bulletins from Microsoft. To sign up to receive the update bulletins, go to the Microsoft Worldwide Web site at www.microsoft.com/worldwide and then reference your country to then access the Microsoft Security Home page at www.microsoft.com/security/.

### Categorizing Patches

As each new patch becomes available, you should determine its importance to your environment, which in turn will help you decide how soon you need to deploy the patch and how much testing you can afford.

Microsoft provides ratings for each vulnerability that is the subject of a security bulletin. The rating levels are shown in the following table.

**Table 6.1: Vulnerability Ratings, as Defined by Microsoft**

| Computer type | Critical rating | Moderate rating | Low rating |
|---|---|---|---|
| Internet servers | Web site defacement, denial-of-service (DoS), or full control. | Difficult to exploit, unusual configuration, or transient effect. | Limited impact such as disclosure of scripts. |
| Internal servers | Elevation of privilege, data disclosure, or modification. Auditing difficult. | Auditable data disclosure, modification, or DoS. | Untargeted or fragmentary data theft or modification, limited DoS. |
| Client systems | Running of arbitrary code without user action; remote escalation of privilege. | Local escalation of privilege, untargeted data disclosure or DoS, use action exploitation. | Limited or fragmentary data theft or modification, hostile Web site attacks. |

The rating system categorizes vulnerabilities according to their potential impact and likelihood.

You can use this rating system as a guide for categorizing patches. However, the Microsoft rating system is just an overall estimate of potential impact in the context of millions of customers worldwide. The severity ratings are based on past experience and subjective judgment. For these reasons, they may not be accurate predictors of impact for your environment. Ultimately, you will need to categorize patches based on your own environment.

## Assessing the Patch

At a minimum, your patch assessment should consist of the following steps:

1. Identifying the patch owner. For all patches, you should have an identified owner who is responsible for evaluation of the patch.
2. Reviewing all documentation. Before applying any service pack, hotfix, or security patch, all relevant documentation should be read and peer reviewed. The peer review process is critical because it mitigates the risk of a single person missing critical and relevant points when evaluating the update.
3. Verifying the patch category. After further assessment of the patch, you may need to change its category. This will affect other aspects of your testing.

As you read the documentation, look for answers to the following questions:

- Is the update relevant, and will it resolve an outstanding issue?
- Will adopting the update cause other problems, resulting in a compromise of the production system?
- Are there dependencies relating to the update? (For example, do certain features need to be enabled or disabled for the update to be effective?)
- Do you need to perform any actions prior to deploying the update?

In addition to examining the documentation released with the updates, you should search the Microsoft support Web site for any additional post-release information on the update. TechNet also provides security bulletins in a searchable (by product name and service pack) database on its Web site. These materials supply critical information that must be referenced.

## Testing the Patches

As with any software, patches may not work perfectly in every environment. Ideally, you should thoroughly test any patches that you are going to install in your environment. However, many security patches need to be installed quickly in order to fix potentially serious problems.

In many cases, you will find that your testing procedure ends up being a compromise between the need to solve a security issue and the need to ensure that your patch is stable in your environment.

The amount of testing that is appropriate will depend on how you have categorized the patch. Using the Microsoft categorizations, the following table shows the minimum level of testing that you should perform for each patch type. In the Trey Research scenario, each of the server roles still had to function properly after the recommended hotfixes were installed. This was confirmed by verifying that various client computers could still connect to the network services running on each server role and performing other basic test procedures to confirm that everything was still working as expected.

**Table 6.2: Minimum Testing for Patches**

| Patch type | Testing phases |
|---|---|
| Critical severity | Assessing the patch<br>Assessing server operations (Limited) |
| Moderate severity | Assessing the patch<br>Installing the patch in a test environment<br>Assessing server operations (Full)<br>Checking the uninstall procedure |
| Low severity | Assessing the patch<br>Installing the patch in a test environment<br>Assessing server operations (Full)<br>Assessing application operations<br>Checking the uninstall procedure |

As part of your risk management procedure, you will need to determine how thoroughly to perform each step. If you skip some phases due to urgency, you should still continue to complete them in a test lab to find potential problems before they occur on already deployed systems.

All testing should occur on servers that resemble your production servers as closely as possible.

## Installing the Patch

You should make sure that the patch installs correctly, understand whether the patch requires a restart, know how much space it takes up (including an uninstall folder), understand what options are available to you, and so on. You also should read any supporting documentation for additional information to evaluate the benefits and drawbacks of applying the patch.

### Testing Server Operations

After the patch is installed, you need to make sure that the server continues to work normally. It is also a good idea to monitor the Event Log and System Monitor for any unexpected results.

Test all of the server functions and make sure that everything operates normally. How much risk you can handle on a particular server that has a particular vulnerability will determine how long you should allow the server to run before concluding that everything is running normally. If there are any problems, you need to make sure they are documented and reported to Microsoft as soon as possible.

---

**Note:** You can use Microsoft Operations Manager (MOM) to collect Event Log and System Monitor information from Windows NT 4.0 servers.

---

### Testing Application Operations

As part of your testing procedure, it is important to test the patch with any applications that coexist on the servers and to make sure that you identify any issues with dependencies. After installing the patch, you should check that all applications continue to function as they did before.

### Preparing for Uninstallation

It is possible that, despite testing, you will encounter problems after installing the patch that result in your having to uninstall the patch. It is important, therefore, to test that the uninstallation works. After uninstalling, you should check that the server continues to run as expected and continue to watch the Event Log and System Monitor counters.

### Creating a Rollback Plan

Even if testing proceeds entirely without incident, it is still possible that there will be challenges as you deploy the patch throughout your organization. You need a plan of action to restore the system to its original state before the patch was deployed.

In some cases, this plan will consist of taking a snapshot backup of a server before the install occurs, so that the server can be restored quickly if problems occur. You should thoroughly test the rollback plan that your organization designs.

## Deploying the Patches

If the testing has proceeded without incident, you are ready to deploy the patch across your organization. There are a number of ways to do this, including the following methods and processes:

- Manual deployment.
- Automated deployment tools from Microsoft or other vendors.
- Scripts built in, and customized for, your environment.

---

**Note:** For additional information about deploying patches, see the Microsoft TechNet article "Best Practices for Applying Service Packs, Hotfixes and Security Patches" at www.microsoft.com/technet/security/bestprac/bpsp.mspx.

---

## Manual Deployment

Manual hotfix installation is the most common installation method. This method consists of simply running the executable that corresponds to the hotfix on each server. If your organization has many servers, though, this may be impractical. Because Trey Research has a moderate number of servers, and because the company's servers are spread among multiple locations, manual deployment is the current patch deployment method.

The name of most hotfixes will tell you important information about the fix. For example, a typical name for a hotfix is Q292435_W2K_SP3_x86_en.exe. In this case:

- Q292435 is the KB article number that provides more information about the hotfix.
- W2K is the product it is intended for (Microsoft Windows 2000).
- SP3 is the Service Pack in which it will be included.
- x86 is the processor architecture for which it is designed.
- en indicates the language in which the hotfix is released (English, in this case).

---

**Note:** Hotfixes that have a file name of QXXXXXX.exe and do not have W2K_SP3_x86 appended to the file name are specific to applications such as Microsoft Internet Explorer.

---

Hotfixes also support several command line switches that you can use to control the behavior of the hotfix installation process. These switches are listed in the following table:

**Table 6.3: Switches for Hotfix Executables**

| Switch | Description |
|---|---|
| -y | Perform uninstall |
| -f | Force applications to close at shutdown |
| -n | Do not create an uninstall directory |
| -z | Do not restart when update completes |
| -q | Use Quiet mode—no UI |
| -m | Use Unattended mode |
| -l | List installed hotfixes |

---

**Note:** Application-specific hotfixes with file names of QXXXXXX.exe typically do not support all of the switches listed in the previous table.

---

If you script the installation of multiple hotfixes, you will want to use the -q and -z switches so that the hotfix is installed without a UI and does not force a restart.

Typically, when you install multiple hotfixes you need to restart the computer between each one. This is because any files that are locked or in use cannot be replaced, so they are placed in a queue to be replaced after the system restarts.

QChain is a tool that allows you to string together several Windows NT 4.0 hotfixes with only a single restart, instead of restarting between each install. To use QChain, run the hotfix installer with the -z switch to instruct the installer not to restart after the installation. Then run QChain.exe and restart the computer.

If additional Windows components, such as the Domain Name System (DNS) service, are added after applying a service pack and patches, it is necessary to reapply the service pack and patches to ensure that the new component is properly patched.

## Scripted Deployment

You may want to create your own scripts using the Microsoft Visual Basic® Scripting Edition (VBScript) language or batch files to roll out patches. These scripts could be in the form of logon or startup scripts that check the current patch status and then check updates from a centralized server. Your scripts can include QChain to ensure that only a single restart is needed if more than one hotfix is required.

## Deployment Monitoring and Reporting

After you have installed the patches into your production environment, you need to continue to monitor your servers. Make sure that you watch the Event Log and System Monitor counters for problems. If you see any other errors on the computers for the next couple of weeks, you should test to make sure that they are not related to the patch that you deployed. Also, if you implemented a patch without thorough lab testing because it was time-critical, you should continue testing the patch in a lab environment afterwards to make sure that nothing was missed.

In addition to monitoring existing servers, it is very important that you monitor the environment as a whole to ensure that new servers are not brought onto the network without installing the current patches. New servers should always receive the latest build, and your organizational monitoring policy should ensure that this happens.

The only way to be sure that any process is working properly is to review it. When you complete the patch management process for each patch, you should review the process to ensure that each one was deployed correctly, and that all procedures ran correctly. This review will help ensure that your patch management process continues to function as it should. When you review the process, you should continue to analyze the environment for further changes. If any occur, you will need to start the patch management process again.

# Implementation

The Trey Research patch management solution includes a number of separate components that work together to deliver reliable, robust inventory and delivery services. The first step Trey performed was a comprehensive analysis of the company's existing network and systems. Patches are often applied inconsistently throughout an organization, and there is no documentation on why, when, and where they have been deployed. Trey wanted to build a unified patch management process that they could use to provide repeatable processes for consistently applying patches to the right computers at the right time.

## Building Update Staging Servers

In many environments, it can be beneficial to have specialized computers from which you perform many of the steps of the patch management process. These systems provide specialized locations for storing security tools, patches, hotfixes, service packs, and documentation. You can use these systems as a place to perform patch analysis, retrieval, and deployment. Microsoft Software Update Service (SUS) and Windows Update Services (WUS) are products that perform all these functions for Windows networks. However, SUS could not be used in the Trey environment because Windows 98 and Windows NT 4.0 do not support it, and Trey chose to build its own staging servers for testing and delivering patches. These servers are actually shares hosted on Microsoft Windows Server™ 2003 domain controllers; Trey downloaded the Microsoft Security Tool Kit, placed it on the staging servers, and used the patches contained therein to harden existing computers. In addition, Trey updated its server and workstation build process to include application of patches from the Security Tool Kit when new computers are built.

Trey chose to use its domain controllers as staging servers to ensure that the security update systems are on one or more dedicated computers that can be tightly controlled and secured. Trey made this decision because these systems will be used to deploy and maintain security patches for all systems in their environment.

Although security update systems do not generally need to be high-powered servers (because the load on them will typically be very light), maintaining high availability is very important.

To properly deploy a security update system, the computer will need direct or indirect Internet access to download the latest patch information from trusted sources, as well as access to each computer that it is responsible for keeping current.

---

**Note:** MOF discusses update systems as part of the release management process.

---

## Identifying Missing Patches

An ongoing analysis process is required to ensure that all servers and workstations remain current with all of the latest patches. To efficiently keep the computers on its network up to date, the Trey IT staff used a combination of the tools described previously in this chapter.

## Scanning the Base Operating System

MBSA allows the Trey IT staff to regularly scan their Windows NT servers and workstations and catalog the results. At first, the IT director used MBSA to perform a quick baseline scan to identify systems that were missing patches; once those systems were patched, MBSA was scheduled to run regularly on different sets of systems.

▶ **To perform a baseline scan using MBSA**

1. Log on to a computer running Windows 2000, Windows XP, or Windows Server 2003 with an account that has administrative privileges for the computer or computers you want to scan.
2. Open a command prompt.
3. Launch the command-line version of MBSA with the –b switch, as follows:

   `mbsacli.exe /hf –d <yourDomain> -b`

   The –b switch runs MBSA in its baseline-scan mode, which will only check for patches identified as baseline by the Microsoft Security Response Center.
4. Review the text report produced, which lists each system found, whether it was scanned, and information about any missing patches.

▶ **To perform a standard scan using MBSA**

1. Log on to a computer running Windows 2000, Windows XP, or Windows Server 2003 with an account that has administrative privileges for the computer or computers you want to scan.
2. Launch MBSA and click **Scan more than one computer**.
3. In the **Domain Name** field, type the name of the domain that you want to scan.
4. Click the **Start scan** button and wait while MBSA enumerates and scans the systems on your network.
5. Click **Pick a security report to view**.
6. Select a computer from the list of security reports, and then double-click the associated computer name.
7. Review the report.

---

The "Microsoft Baseline Security Analyzer V1.2" white paper, available at www.microsoft.com/technet/security/tools/mbsawp.mspx, contains a complete description of the scanning modes and vulnerability checks that MBSA performs.

---

## Scanning Office Installations

You can use the Office Update Inventory Tool described previously in this chapter to scan individual workstations for critical Office updates. However, it requires that a client executable be installed and run for each system. Because Trey Research will be migrating to Office 2003 as a part of its IT modernization plan, the company has a mix of Office 2000 and Office XP installations, and administrators have already installed the current service packs as part of normal system maintenance. Trey assessed the risk of a network compromise due to an Office security flaw as relatively low, so it has adopted a strategy that allows users to directly visit the Office Update Web site at http://office.microsoft.com/officeupdate to check for updates. Critical systems, and those belonging to executives, are scanned in one of two ways:

- Systems running Windows NT, Windows 2000, and Windows XP are scanned using MBSA in local mode, which checks for Office security patches on the local computer only.
- Systems running Windows 98 are manually scanned using the Office Update Inventory Tool.

## Planning for Patch Application

The Trey Research IT director wrote a patch application plan template that is used as the base for the company's monthly patch application plans. Because Microsoft releases security patches on a predictable schedule, having a template allows Trey administrators to follow a standardized process to evaluate, assess, and deploy each set of security patches when Microsoft releases it. Template guidelines include:

- What services, systems, and applications are considered to be business-critical and which ones are optional or non-essential.
- How to determine whether a patch applies to a particular system or set of systems.
- How to determine whether the patch may cause other problems.
- Instructions for identifying dependencies for each patch.
- Criteria for deciding whether a vulnerability is sufficiently important to require emergency installation or scheduling of an additional patch deployment.
- Who must review and approve patch deployment.
- How to roll back patch installation in the event of a failure.

In addition to developing this template, the Trey IT staff signed up for the Microsoft security notification service, and they regularly review the Microsoft Security Home page at www.microsoft.com/security/.

### Categorizing Patches

Trey Research uses the standard Microsoft categorization scheme for patch severity, as shown earlier in Table 6.1. However, the company adds an additional parameter to the categorization: an indication of whether the patch is applicable to its environment. For example, because Trey is not currently using Office 2003, patches for it would receive a rating of “not applicable,” regardless of the severity that Microsoft associates with the vulnerability. Alternatively, because Trey uses Microsoft Data Engine (MSDE) and SQL Server 2000 heavily, all SQL Server patches are rated as “applicable." This approach requires more categorization effort, but it allows Trey to focus their effort on patches that are most pertinent to the company's environment.

### Testing the Patches

The lead IT administrator at Trey Research developed a patch test plan that establishes the conditions for testing patches for wide deployment based on their severity, their applicability, and the nature of the systems being patched. Trey chose to adopt the Microsoft recommendations for the level of testing (outlined in Table 6.2) and built a model of its production network for lab use that contains representations of workstations and servers. New patches are first deployed to the test lab to verify that the patches install properly and that they do not break anything critical. If initial testing is successful, patches are deployed according to the criteria in the patch application plan.

After the patch is installed in the test lab, Trey performs a standard set of tests. These tests include adding and removing resources in Microsoft Active Directory® directory service and running a standard set of the company's line-of-business applications. In addition, test pass criteria require checking the Event Log and System Monitor for any unexpected results.

**Note:** You can use Microsoft Operations Manager (MOM) to collect Event Log and System Monitor information from Windows NT 4.0 servers.

## Deploying the Patches

If the testing has proceeded without incident, Trey Research deploys the patches to systems that require it. They use a combination of two methods: manual deployment by administrators or affected users and automated deployment using custom scripts.

### Manual Deployment

Manual hotfix installation is the most common installation method. This method consists of simply running the executable corresponding to the hotfix on each server. If your organization has many servers, though, this method may be impractical. Because Trey has a moderate number of servers, and because the company's servers are spread among multiple locations, manual deployment is the standard patch deployment method for servers. To avoid unnecessary downtime, the IT staff installs fixes using the –z switch to prevent reboots until after the last patch is installed; the last step in the installation process is to run Qchain.exe to unify all files installed during the patch operations.

### Scripted Deployment

Trey Research built a custom script that uses QChain to install multiple hotfixes and reboot the computer after the last fix is installed. The script is described in KB article 296861, "How to Install Multiple Windows Updates or Hotfixes with Only One Reboot" at http://support.microsoft.com/?kbid=296861. The following script sample shows how Qchain was used:

```
@echo off
setlocal
set PATHTOFIXES=some path
%PATHTOFIXES%\Q123456_w2k_sp2_x86.exe -z -m
%PATHTOFIXES%\Q123321_w2k_sp2_x86.exe -z -m
%PATHTOFIXES%\Q123789_w2k_sp2_x86.exe -z -m
%PATHTOFIXES%\qchain.exe
```

# Summary

The majority of breaches in IT security come from exploitation of system environments that are not fully up to date with security patches. Good patch management is essential to minimizing the security risks that you face. If you treat patch management seriously you will likely be able to dramatically reduce the costs associated with security breaches. For Trey Research, as in most organizations, patch management is an ongoing process; servers must be kept up to date with security-related hotfixes.

## More Information

- The Microsoft Security Tool Kit can be useful for obtaining the service packs and hotfixes needed to keep your servers current. The toolkit contains important security information, current service packs, critical security patches for Windows NT 4.0, Windows 2000, IIS, and Microsoft Internet Explorer. It also includes the Critical Update notification tool. This tool connects to the Windows Update site to ensure that all the latest patches are installed on your computers. The Security Tool Kit is available from Microsoft TechNet at www.microsoft.com/technet/Security/tools/stkintro.mspx.
- The Microsoft Solution Accelerator for patch management using SMS 2003 should be considered the master guide for deploying SMS 2003 patch management systems. *Patch Management Using Microsoft Systems Management Server 2003* is available at www.microsoft.com/technet/itsolutions/techguide/msm/swdist/pmsms/2003/pmsms031.mspx.
- *Systems Management Server 2003 Concepts, Planning, and Deployment Guide*, a complete reference to SMS deployment issues, is available at www.microsoft.com/resources/documentation/sms/2003/all/cpdg/en-us/default.mspx.
- Comprehensive information about SMS 2003 security patch management capabilities is available at www.microsoft.com/smserver/evaluation/capabilities/patch.asp.
- The Office Update Inventory Tool is available at www.microsoft.com/office/ork/2003/journ/offutoolv2.htm.
- The article "Best Practices for Applying Service Packs, Hotfixes, and Security Patches" is available at www.microsoft.com/technet/security/bestprac/bpsp.mspx. This article describes a set of recommended practices for managing the installation, testing, and deployment of hotfixes, service packs, and security patches.
- Organizations that use SMS 2.0 may want to obtain the SMS 2000 Software Update Services Feature Pack, which adds advanced inventorying and patch management capacity to SMS 2.0. More information about the SUS Feature Pack for SMS 2.0 is available at www.microsoft.com/technet/prodtechnol/sms/sms2/confeat/smsfpdep.mspx.

## Third-Party Tools

A number of third-party tools are available to help with patch management. These tools offer some features not currently available with the free tools from Microsoft, such as the ability to deploy fixes and have the status reported back, create computer groupings with similar update needs, support other products not covered by the tools described previously, and use a graphical user interface (GUI) for administrative tasks. You should evaluate these features and determine whether they are appropriate for your environment. Available third-party tools include:

- Shavlik HFNetChkPro. Built on HFNetChk technology, HFNetChkPro offers a number of command-line features not included in MBSA. More information on HFNetChkPro is available at www.shavlik.com/pHFNetChkPro.aspx.
- Bindview bv-Control. This product offers a GUI to help simplify the process of checking for noncompliant computers. It also has an update service that lets you know when new patches have been released. More information is available at www.bindview.com/Products/VulnMgmt/index.cfm.
- Pedestal Software Security Expressions. This product allows administrators to implement security lock down policies on Windows-based and UNIX-based computers. It also has the capability to check for hotfixes and automatically download and install them if needed. More information is available at www.pedestalsoftware.com/products/se/.

# 7

# Antivirus Protection

Previous chapters in this guidance have referred to the risks that viruses and other forms of malicious software or *malware* present to the servers at Trey Research, the company in the guide scenario. This chapter delves more deeply into the subject of malware and discusses how Trey can protect its Microsoft® Windows®-based clients and servers from this risk.

## Introduction

In a relatively short period of time, viruses and other forms of malware have gone from being minor, infrequent nuisances to major security threats. Their proliferation has caused considerable damage on a global scale and even resulted in the demise of some companies that were caught unprepared. Protecting against these threats is therefore a critical issue for any organization, regardless of whether it runs older systems or the latest technology.

There are numerous paths that viruses and worms can follow to infect an organization's computing systems, and you must effectively protect each path to avoid penetration by malicious code. This chapter explores virus prevention and reaction steps that Trey should take to protect file, mail, and Web servers, workstations, and other network devices from these threats. It also examines the potential consequences of viruses and worms in order to provide rationale for implementation of broad scale solutions.

## Background

To develop an adequate defense against any threat requires that you first understand the nature of the threat. The better your understanding, the better equipped you will be to design and implement an effective response. Threats fall into three general categories: Trojan horses, viruses, and worms. It is critical that you understand how these threats differ from one another:

- Trojan horse. A program that appears to be useful or harmless but that contains hidden code ("malicious payload") designed to exploit or damage the system on which it is run. Trojan horse programs are most commonly delivered to users through e-mail messages that misrepresent the program's purpose and function. Trojan horses are also called "Trojan code."
- Worm. A worm uses self-propagating malicious code that can automatically distribute itself from one computer to another through network connections. A worm can take harmful action, such as consuming network or local system resources, even possibly causing denial of service (DoS) attacks. Some worms can execute and spread without user intervention, while others require users to directly execute the worm code in order to spread. Worms may also deliver a payload in addition to replicating themselves.
- Virus. A virus uses code written with the express intention of replicating itself. Viruses attempt to spread from computer to computer by attaching themselves to host programs. They may damage hardware, software, or data. When the host is executed, the virus code also runs, infecting new hosts and sometimes delivering an additional payload.

---

**Note**: For more background information on antivirus terminology and defensive strategies, see *The Antivirus Defense-in-Depth Guide* referenced in the “More Information” section at the end of this chapter.

---

## Business Issues

Many organizations underestimate the destructive power of malware. Some have never been infected and therefore do not understand the potential consequences. Others have had only minimal exposure to the problem. However, even a virus with a relatively minor payload can have a tremendous overall impact. Possible damages are not limited to lost files or systems that must be restored or reinstalled. but also include:

- Loss of timely access to data. If a network-wide infection occurs, documents, databases, e-mail messages, or other important data items could potentially be lost if no recent or usable backup is available. The result would likely be a significant investment in time to recreate the lost data, affecting production schedules and employee morale. Even the loss of one critical document can have a significant impact on a business. An engineer might make considerable progress on an analysis report in a day's time, only to have that effort be completely lost when a virus causes the document to be corrupted or destroyed.

- Lost productivity. Recreating documents and other files that are lost through a malware attack can be a disaster for an organization. The most common circumstance in which a backup is found to be unusable is when an emergency restore is attempted. In some cases, just identifying what has been lost can be a major undertaking. Actually recreating files is just one aspect of the recovery. Invariably, deadlines must be delayed because users must concentrate on recreating data rather than handling their normal tasks. Recovery can also make it necessary to bring in temporary workers to rebuild the data. This lost productivity and added expense can easily put organizations out of business, particularly if they are working on small margins or in very competitive markets.
- Exposure of proprietary or customer data. Loss of proprietary information such as source code, trade secrets, or even confidential memos can have dire consequences for an organization. In the case of Trey Research, much of the data gathered by its engineers has commercial value. In addition, some of the analysis methods and algorithms used to compute results from raw data are proprietary. Loss or exposure of this data can have serious financial consequences.
- Contributory liability. In environments where legally or commercially sensitive data is in use, liability issues arise. For example, imagine that a company hired Trey to assess contamination at the site of a school. If that customer's data were to be prematurely exposed, the customer might become liable—and either the plaintiff or the defendant in any resulting suit could decide to include Trey in its liability.

## Technical Issues

The chief technical issues for antivirus protection involve where and how files, e-mail messages, and network traffic is scanned to determine whether viruses are present. Consider the following strategies:

- Server-based scanning should be implemented first. Before you consider any options for protection, however, you should take the time to evaluate—and, if necessary, improve—disaster recovery policies and procedures for each server. You should evaluate not only the existing capabilities for backup, but also the recovery procedures in terms of time to restore. Compare the amount of time that users will be idle with the cost of revamping your backup procedures with newer, faster hardware. In addition, you need to have a mechanism in place for performing selective test restores as part of your backup procedures. A backup that cannot be successfully restored is virtually worthless.
- File servers should be protected with antivirus solutions on each server that actively monitor the file system. When existing files are modified or new files are added, the antivirus application scans the files and can quarantine or repair an infected file before the infection can spread.
- E-mail servers should use scanners that understand the proper way of scanning queues, transaction logs, and message databases. Antivirus programs that do not properly scan and disinfect these items can cause interruptions to e-mail service and data loss.
- *All* clients should be protected with antivirus software. Whatever scanning tool is selected must provide proactive, real-time scanning of a client's file system to catch infections as they occur. Support for one-time and scheduled scans is less important if the application supports real-time scanning, but *only* if you force configuration of the clients to enable real-time scanning. Without real-time scanning, start-up scans and frequent-scheduled scans are a practical necessity.

- The frequency of updates to antivirus scanner signatures is very important, because most vendors release signature updates quickly after new viruses are identified. An antivirus solution without an up-to-date antivirus signature database is only partially effective. Since new viruses are created every day, you must keep the antivirus signature database as current as possible to ensure effective scanning. When planning an antivirus strategy, give consideration to the accessibility of the antivirus updates and how frequently your servers will need to be updated.
- How an antivirus solution reacts to a virus outbreak on the network and provides notification to administrators is crucial. Extended logging and notification by e-mail, pager, or other methods are very important features that can notify the administrative team quickly when an outbreak occurs.

## Security Issues

Antivirus coverage is obviously a key part of security management and improvement. However, Trey faces several additional issues related to other parts of its security improvement plan:

- Consistent attention to patch management will help reduce the likelihood of infection by many viruses and worms. Trey is moving aggressively to implement a patch management plan (as described in Chapter 6, “Patch Management”) to ensure that needed patches are quickly applied to vulnerable computers.
- In addition to deploying antivirus scanners, it may be necessary to use intrusion detection or network monitoring software to catch anomalous network usage patterns that may indicate a newly emerging virus or worm.
- In environments where Windows 2000 or Windows XP is the primary desktop operating system, it is possible to use Group Policy objects (GPO) to force the installation of antivirus software and to ensure that it remains active. There is no good way to do this in Microsoft Windows NT® or Windows 98.
- Reducing the attack surface of clients may require additional measures, including upgrading the installed version of Microsoft Internet Explorer and using the Internet Explorer Administration Kit, Group Policies, or Windows NT system policies to set appropriate security policies for the browser.

## Solution Requirements

The solution requirements for most antivirus deployments are fairly simple. In the case of Trey Research, the requirements are to:

- Provide automatic, file-level scanning for all clients.
- Provide automatic, file-level scanning for all file servers.
- Provide automatic and on-demand scanning (using Microsoft Exchange-aware scanning tools) for all Exchange mailbox servers.
- Receive automatic updates for scanning signatures for all scanners within 12 hours of vendor release.
- Investigate restrictions on the use of removable media (floppy disks, CD-ROMs, USB thumb drives) to help prevent the spread of viruses.

Because of the current composition of the Trey network, several additional requirements have been explicitly postponed until the desktop environment is upgraded:

- Automatic initial installation of antivirus software.
- Enforcement of the Trey security policy that requires antivirus software to be installed and running at all times.
- Use of a separate perimeter antivirus/anti-spam filter that screens e-mail before initial delivery to the Exchange servers.

# Solution Design

The Trey IT director wanted quick deployment of an antivirus solution to help reduce the company's exposure to viruses and other forms of malware.

## Solution Concept

The Trey antivirus solution has two tiers: client-side protection and server-side protection. Trey purposefully chose scanners from two different vendors to maximize the likelihood that signatures would become available quickly during an outbreak.

## Solution Prerequisites

There are no prerequisites for this solution.

## Solution Architecture

The Trey antivirus solution has several components that must work together to deliver effective antivirus protection.

### Server Protection

Virus protection is not a one-step process in which installation of a single application suddenly delivers instant and perfect security. The first step for Trey involved patching all of its servers to a solid baseline. Because many forms of malware exploit known vulnerabilities in the operating system, installing the latest service pack and any post-service pack patches should be your first step in protecting against threats. Microsoft Windows Server™ 2003 often allows updates without a server reboot, but earlier Windows versions usually do not. Upgrading a server, therefore, requires careful planning to ensure that the server is adequately backed up prior to initiating the upgrade and that the upgrade can be rolled back if the upgrade fails. Consideration must also be given to when the upgrade can best be accomplished to minimize downtime.

To minimize the amount of time required to update the server, the Trey IT administrators used QChain.exe, available from Microsoft Knowledge Base article 815062: "The correct file is not installed when you chain multiple hotfixes" at http://support.microsoft.com/?kbid=815062, to install multiple updates with a single reboot. Additional information about using QChain.exe is available in Microsoft Knowledge Base article 296861: “How to Install Multiple Windows Updates or Hotfixes with Only One Reboot” at http://support.microsoft.com/?id=296861. This patching was conducted during non-business hours.

Trey also enhanced its perimeter protection (as described in Chapter 3, ”Network Hardening and Security”) to provide additional security for computers inside the network.

In addition to providing scanning at the file server, you must also analyze the server's existing file system security to identify means of minimizing the server's exposure. First, the servers should all use the NTFS file system (NTFS), because FAT provides essentially no security. Using NTFS enables administrators to tune permissions to ensure that only the operating system itself has the capability to write to core directories and files; this helps restrict the damage that a virus can do if it does infect the servers.

Additionally, the Trey IT administrators reviewed all of the existing file server shares to eliminate unnecessary shares. They also added appropriate NTFS and share permissions to the shares to prevent anonymous access. These steps can protect against worms and viruses that exploit unprotected shares to propagate. In addition, you can use hidden shares to further reduce server exposure.

All file/print, application, and member servers are protected with the same scanning product. Exchange servers are protected with an Exchange-aware product from the same vendor. This approach allows all of the servers to be managed as a unit with the vendor's enterprise management tools. Local administrators do not have the ability to configure scanning settings, although they can initiate manual scans.

## Mail Server Protection

The most common entry point for malware in a network is the e-mail system, so in order to be effective, any protection scheme must include e-mail server protection. Although antivirus solutions that scan the file system can often catch viruses as they arrive in the e-mail system, a much better approach is to install an antivirus solution that actively and specifically scans messages and attachments. In networks in which Exchange Server is installed, you can choose from several antivirus solutions that are designed to work in conjunction with Exchange to proactively scan incoming and outgoing mail. These antivirus solutions typically use the antivirus application programming interface (API) built into Exchange Server to access and scan messages and attachments.

In situations in which other e-mail servers are used within the network, or where users connect to an external e-mail server, antivirus solutions that scan Simple Mail Transfer Protocol (SMTP) traffic at the gateway are an effective means of preventing infection from incoming messages, as well as preventing viruses from leaving the network in outgoing messages. Rather than integrating with the e-mail server software, these gateway-based solutions scan the SMTP traffic as it enters the network.

In some scenarios, using multiple antivirus engines can add an extra layer of protection. Some antivirus solutions such as GFI MailSecurity (available through the GFi Security and Messaging Software Web site at www.gfi.com) support multiple, concurrent antivirus engines from multiple antivirus vendors. Using multiple engines helps guard against the possibility that one particular engine will miss an infected message because the vendor has not yet provided an updated antivirus signature database file, or because a vendor's update site is down due to a DoS attack or network outage.

If you are not able to use multiple scanning engines, consider a combination of solutions. For example, you might deploy an Exchange Server-based solution to scan at the server with one or more antivirus engines, along with a gateway-based solution that uses one or more other antivirus engines.

Exploit detection is another important consideration when evaluating and deploying antivirus solutions for your e-mail system. Scanning for known viruses is essential, but scanning for e-mail exploits is equally important. An e-mail exploit is a script, an executable file, a malformed Multipurpose Internet Mail Extensions (MIME) header, or other mechanism that is used to exploit a vulnerability in the client e-mail application or operating system. The Nimda and BadTrans.B viruses are examples of viruses that used e-mail exploits to propagate and infect target systems.

An antivirus solution that provides e-mail exploit detection scans messages for methods that exploit the operating system or e-mail client. In effect, scanning for e-mail exploits enables the antivirus solution to scan for a broad category of potential threats. Although each virus has an individual signature and requires identification of that specific signature, a single exploit might be used by many viruses, including new viruses that attempt to take advantage of existing exploits. By blocking the exploit, you effectively block a range of viruses.

Whatever antivirus solutions you choose for your e-mail servers and network, there are two main reasons to consider scanning outgoing messages as well as incoming messages. First, the presence of an outgoing infected message is a sure indicator of an infected client system and can serve as early warning against a network outbreak.

Scanning outgoing messages for e-mail exploits is important for the same reason. Second, outgoing viruses, even if they fail to have a major impact on your internal network, can have disastrous consequences for your organization's reputation and customer relationships. Customers whose own networks become infected by a virus sent from your e-mail server will have their confidence in your company shaken, which could very well damage the business relationship.

Even in the absence of an e-mail solution, you can take steps to reduce the likelihood of an e-mail-borne virus infection. You can implement an attachment-blocking scheme even if you have an e-mail scanning solution in place. Blocking attachments not only helps to eliminate virus infections by preventing users from receiving executables and other types of files most susceptible to infection, but it can also help block e-mail exploits that rely on scripts or other file types that enable the exploit.

Microsoft Knowledge Base article 235309, "Outlook E-Mail Attachment Security Update" at http://support.microsoft.com/?kbid=235309 provides information about enhanced security protection for Microsoft Outlook®, such as attachment blocking and other features to prevent specific types of attachments from being opened by users and thereby infecting the local system and eventually the network. The security update is also available for Outlook 98 and is incorporated into Outlook 2002 and later versions. Exchange Server administrators can configure attachment blocking and other e-mail security options at the server, including adding or removing specific file types from the blocked lists. The Outlook Security Features Administrative Package (AdminPak) is included on the Microsoft Office 2003 Resource Kit CD–ROM and is available from Microsoft Office Online at http://office.microsoft.com/officeupdate/. In addition to providing the means for Exchange Server administrators to configure attachment blocking options, the Administrative Package also enables administrators to configure which applications can access a user's address book, send messages programmatically, and perform other actions.

If you do not use Exchange Server or prefer not to control security options at the server, you can configure Outlook locally to modify attachment blocking. Modifying a handful of registry settings for Outlook modifies the attachment blocking behavior. For information about modifying attachment blocking options for Outlook at the local level, see Microsoft Knowledge Base article 829982, “Cannot open attachments in Microsoft Outlook” at http://support.microsoft.com/?id=829982. If you use Outlook Express as an e-mail client, either in place of or in conjunction with Outlook (or another client), you need to patch Outlook Express to help secure against e-mail-borne viruses and exploits. The Outlook Express Security Patch provides many of the same attachment blocking features as the Outlook security update and also fixes other problems, including a buffer overflow exploit for Outlook Express mail headers.

## Client Protection

All clients—Windows 98, Windows NT 4.0, Windows 2000, and Windows XP—are protected using the same scanning product. A Group Policy was configured to apply consistent settings to Windows 2000 and later clients, and the Windows NT system policy mechanism described in Chapter 4, “Hardening Windows NT 4.0,” was used to apply consistent settings to Windows NT and Windows 98 clients. Trey Research also updated its written security policy to require the use of antivirus software on all computers used to connect to the corporate network, including home computers used with the company’s virtual private networking (VPN) capability.

### Updates

Antivirus solutions typically offer scheduling of dynamic updates, either directly from the vendor across the Internet or from a local network server where the updates have been posted. You can choose to have all servers pull their updates from the vendor across the Internet or to have a single server (or selection of servers) pull the updates from the vendor and have the remaining servers pull the updates from those local servers. The method that you choose depends on whether you want or need to minimize Internet traffic. If you do, pulling from a local server is a way to reduce that traffic.

The availability of updates is another factor that can influence your decision whether to choose a single-vendor or single-engine solution or to opt for a solution that integrates scanning engines from multiple vendors. A distributed DoS attack might succeed in taking down a single vendor for a period of time, and using solutions from multiple vendors can help avoid that potential problem.

Trey chose to set their servers to update virus definitions every two hours; clients update daily at 4:00 A.M. Eastern Time. This time was chosen to provide an opportunity to stop viruses whose propagation begins in Europe (a pattern that has held for several major outbreaks in the past).

## How the Solution Works

In addition to implementing antivirus tools on clients, servers, and mail servers, Trey took some additional steps as part of its overall security modernization plan. These steps are relevant to virus protection because they strengthen the defense offered by antivirus tools alone.

Trey took the following additional measures, described elsewhere in this guidance:

- Updated client and server operating systems with all updates and patches.
- Updated Internet Explorer, Outlook Express, and Outlook with the latest versions and patches.
- Developed and implemented programs to educate users about the methods of infection and steps they must take to limit the network's exposure.
- Hardened servers and clients by removing or disabling unnecessary services that can be subjected to attack, removing shares, and locking down services such as IIS.
- Physically secured servers to prevent access by unauthorized or untrained users.
- Implemented perimeter protection with firewalls.
- Deployed security options for Internet Explorer through the *Internet Explorer Administration Kit (IEAK)* available at http://www.microsoft.com/windows/ieak/evaluation/default.asp, system policies, or Group Policies to prevent install-on-demand, Microsoft ActiveX® scripting, and other potential threats.

# Summary

Antivirus protection is a necessary part of any security effort that is targeted at networked computers. When you choose and deploy the right antivirus solution appropriately, it adds a valuable degree of protection to other security measures that are already in place.

This chapter has described the fundamentals of how Trey Research chose to implement antivirus protection. The Trey approach calls for client-based and server-based antivirus scanners from different vendors, coupled with aggressive scanning and signature update policies. This deployment is backed by changes to the security and computer use policy that require all users to use antivirus software at all times.

## More Information

- Additional information about Exchange Server and antivirus solutions is available in Microsoft Knowledge Base article 823166: "Overview of Exchange Server 2003 and antivirus software" at http://support.microsoft.com/?kbid=823166.
- Additional information about the Outlook Express Security Patch is available in Microsoft Knowledge Base article 267580: "OLEXP: Information about the Outlook Express Security Patch" at http://support.microsoft.com/?kbid=267580.
- *The Antivirus Defense-in-Depth Guide* from Microsoft provides comprehensive information about designing a defense-in-depth antivirus solution. The guide is available at http://go.microsoft.com/fwlink/?LinkId=28732.

# 8

# Conclusion

This guidance has discussed many of the challenges involved to deploy older Microsoft® Windows NT® version 4.0 and Microsoft Windows® 98 clients in enterprise environments. The guidance referenced the experiences of Trey Research, a fictitious company that needed to identify and mitigate security vulnerabilities. Many of the issues that Trey resolved are faced by organizations that need to implement their own secure networking environments.

The information presented in this guidance will allow you to identify secure networking infrastructure, create baseline configurations for servers and workstations running older operating systems, define procedures for sound patch management, and institute a proactive antivirus strategy. Together, these techniques can help extend the life of important assets until organizations are able to upgrade to newer, more secure operating systems.

While the Trey Research scenario may not mirror your particular environment, the techniques presented here apply to organizations of all sizes and can be adapted to most environments. The most important concept is that sound security is an ongoing process that should be integrated into the daily life of an IT organization, and that security involves installation, baselining, monitoring, and updating processes.

# Threats in the Trey Environment

The threats that Trey identified as part of their risk analysis required a detailed plan of identification and mitigation. The following sections discuss the ways in which the Trey IT staff was able to plan and carry out the task of securing the company network.

## Physical Security Threats

The biggest physical security threat to Trey systems involves uncontrolled physical access to their computers. Many key servers and workstations are in relatively unprotected areas. Trey partially mitigated this risk by using physical locks and better access controls, along with moving the most important servers to more secure locations. Other physical security risks, like environmental damage, were adequately mitigated by non-technical measures.

## Denial-of-Service Threats

Most of the denial-of-service risks that affect the Trey network can be addressed by the network hardening and filtering features available in Windows NT 4.0. To bolster this protection, they installed personal firewalls on all their Windows 98 and Windows NT computers. In addition to proper server and workstation configuration, proper network segmentation and firewall configuration helped Trey specify which types of traffic were allowed to pass to specific computers. However, Trey was not able to completely mitigate the threat of network traffic tampering or spoofing. Microsoft Windows 2000, Windows XP, and Windows Server™ 2003 support the use of Internet Protocol Security Extensions (IPsec) to protect sensitive network traffic, but Windows 98 or Windows NT do not support IPsec.

## Malicious Code Threats

Trey identified three key malicious code threats: user execution of malicious code, virus outbreaks, and worm outbreaks. They addressed these in three specific ways:

- To partially mitigate the risk of user execution of malicious code, Trey upgraded its workstations to Internet Explorer 6.0 SP1 and applied more restrictive security settings. However, a more effective approach—using Windows software restriction policies to only allow trusted applications to run—cannot be applied because Windows NT and Windows 98 do not support restriction policies.
- To guard against the possibility of virus infection, Trey Research installed centrally managed antivirus software to protect all servers and workstations. Virus definitions are now updated daily, and weekly scans can be scheduled or manually initiated from the antivirus console. If a virus-infected computer, downloaded file, or media is introduced into the network, protected clients have a substantially reduced risk of infection.
- To help protect against worms, Trey segmented their network so that all older computers are on their own separate segment of the network with their own firewall. They developed patch management processes to ensure timely deployment of new patches to all systems, and reviewed and restricted their firewall configuration to ensure that no unnecessary ports were open.

Trey's effort to mitigate these threats is complicated because neither Windows NT nor Windows 98 supports the full set of patch management tools available from Microsoft and other third-party vendors. In particular, Trey cannot use the Microsoft Baseline Security Analyzer (MBSA) to scan their Windows 98 computers, so staff at the company must inventory these computers, and then manually apply patches to them. This represents a serious problem for Trey that will only be solved when the company completes its IT modernization.

Ultimately, the biggest protection against threats from malicious code is educating users to use good security practices. These practices include choosing the appropriate Internet Explorer security settings, and exercising caution in downloading and executing programs or attachments.

## Information Disclosure Threats

Trey took three primary measures to help mitigate the information disclosure threats described in Chapter 2. First, the company began requiring the use of NTLMv2 authentication on all its computers. This requirement had a significant compatibility impact, and it remains less secure than a pure Kerberos deployment. However, it is adequate as a stopgap protective measure until Trey can complete its IT modernization and deploy IPsec.

As an additional protective measure, Trey enabled the use of server message block (SMB) signing for its Windows NT, Windows 2000, and Windows Server 2003 computers. This approach helps ensure that older computers would be able to vouch for the authenticity of all network transmissions. It drafted a plan to monitor computers for performance bottlenecks. A 10 percent to 15 percent performance drop is expected; if this proves to be too much of an impact on some servers, Trey can remove the **RequireSecuritySignature** setting from affected clients. A matching Group Policy object (GPO) in the Microsoft Active Directory® directory service controls SMB signing for native Active Directory servers and workstations; Trey may also have to reset this.

For more protection against offline attacks against the Security Account Manager (SAM) database, Trey required the use of the Syskey utility on all their Windows NT, Windows 2000, and Windows 2003 servers. This helps reduce the risk that an attacker will obtain sensitive security data from these machines.

Trey was not able to effectively mitigate the risk of data theft from the company's mobile computers. The Encrypting File System (EFS), available in Windows 2000 and Windows XP, allows users to selectively encrypt critical data on their computers so that if an attacker steals one, the attacker cannot recover the encrypted data; in this way, EFS would provide the company with effective mitigation.

## Account Compromise Threats

The biggest account threat Trey faces involves the possibility for an attacker to compromise accounts due to weak passwords. Because Windows 98 supports a maximum password length of 8 characters, Trey could not adequately mitigate this risk. However, the use of NTLMv2 authentication helped to somewhat reduce the risk that password hashes could be recovered from the network. There is still some risk that an attacker could reset the local password on individual computers; Windows 2000, Windows XP, and Windows Server 2003 include features designed to prevent this type of attack.

# Summary

This guidance discussed how to identify and mitigate security risks in network environments containing computers running Windows 98 and Windows NT 4.0 Workstation and Server.

- Chapter 1 discussed the Trey infrastructure and introduced a fairly typical network environment of its size.
- Chapter 2 introduced the SRMD process and components, giving examples of how they fit into the Trey IT environment. The chapter then used the principles of SRMD to audit and identify vulnerabilities in the company's network.
- Chapter 3 discussed ways to securely design a network infrastructure to support necessary communications while denying unnecessary or harmful traffic.
- Chapters 4 and 5 demonstrated ways to mitigate the security vulnerabilities of Windows NT 4.0 servers and workstations, and Windows 98 workstations. Key topics included configuring computers to participate in Active Directory, applying security policies, and choosing reliable and secure authentication and communications protocols.
- Chapter 6 established procedures for implementing an ongoing patch management process and discussed auditing, baselining, patch installation, and process automation.
- Chapter 7 provided an analysis of the growing menace of viruses, worms, and Trojan horses, and discussed sound methods for countering these threats.

Organizations that make these types of tasks a regular part of ongoing operational procedures will reap benefits in terms of system uptime, user satisfaction, and an extended lifecycle for integral older computers.

# Index

## T



## V



## W



## Y

# Acknowledgments

The Microsoft Solutions for Security group (MSS) would like to acknowledge and thank the team that produced *The Microsoft® Windows NT® 4.0 and Windows® 98 Threat Mitigation Guide*. The following people were either directly responsible or made a substantial contribution to the writing, development, and testing of this solution.

### Authors
Devin Ganger
Tom Meunier
Paul Robichaux

### Testers
Gaurav Singh Bora
Swathi Palukuru

### Editors
Wendy Cleary
John Cobb
Steve Wacker

### Reviewers
Mert Biyikli
Eric Cameron
Tim Cole
Steve Clark
Mike Greer
Chrissy Lewis
Allen Stewart
Jessica Zahn

### Program Manager
Bomani Siwatu

### Contributors
Scott Blamey
Chase Carpenter
Tim Cole
Masoud Hoghooghi
Joanne Kennedy
Mohan Kotha
Randy McLaughlin
Geoff Morris
Jeff Newfeld
Rob Oikawa
Tessa Porterfield
Bill Reid
Brian Schafer
Dale Weiss
Bill Wesse
Jay Zhang