

Securing Network Servers

Gary Ford
Julia Allen
Christopher Alberts
Barbara Fraser
Eric Hayes
John Kochmar
Suresh Konda
Klaus-Peter Kossakowski
Derek Simmel
Dwayne Vermeulen

February 1999

SECURITY IMPROVEMENT MODULE
CMU/SEI-SIM-007



Carnegie Mellon
Software Engineering Institute

Pittsburgh, PA 15213-3890

Securing Network Servers

CMU/SEI-SIM-007

Gary Ford

Julia Allen

Christopher Alberts

Barbara Fraser

Eric Hayes

John Kochmar

Suresh Konda

Klaus-Peter Kossakowski

Derek Simmel

Networked Systems Survivability Program

Dwayne Vermeulen

Carnegie Mellon Research Institute

February 1999

This report was prepared for the

SEI Joint Program Office
HQ ESC/AXS
5 Eglin Street
Hanscom AFB, MA 01731-2116

The ideas and findings in this report should not be construed as an official DoD position. It is published in the interest of scientific and technical information exchange.

FOR THE COMMANDER

Mario Moya, Maj, USAF
SEI Joint Program Office

This work is sponsored by the SEI primary sponsor and the U.S. Army Land Information Warfare Activity (LIWA) ACERT.

The Software Engineering Institute is a federally funded research and development center sponsored by the U.S. Department of Defense.

“CERT” and “CERT Coordination Center” are registered in the U.S. Patent and Trademark Office.

Copyright © 1999 by Carnegie Mellon University.

Requests for permission to reproduce this document or to prepare derivative works of this document should be addressed to the SEI Licensing Agent.

NO WARRANTY

THIS CARNEGIE MELLON UNIVERSITY AND SOFTWARE ENGINEERING INSTITUTE MATERIAL IS FURNISHED ON AN “AS-IS” BASIS. CARNEGIE MELLON UNIVERSITY MAKES NO WARRANTIES OF ANY KIND, EITHER EXPRESSED OR IMPLIED, AS TO ANY MATTER INCLUDING, BUT NOT LIMITED TO, WARRANTY OF FITNESS FOR PURPOSE OR MERCHANTABILITY, EXCLUSIVITY, OR RESULTS OBTAINED FROM USE OF THE MATERIAL. CARNEGIE MELLON UNIVERSITY DOES NOT MAKE ANY WARRANTY OF ANY KIND WITH RESPECT TO FREEDOM FROM PATENT, TRADEMARK, OR COPYRIGHT INFRINGEMENT.

This work was created in the performance of Federal Government Contract Number F19628-95-C-0003 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center. The Government of the United States has a royalty-free government-purpose license to use, duplicate, or disclose the work, in whole or in part and in any manner, and to have or permit others to do so, for government purposes pursuant to the copyright license under the clause at 52.227-7013.

Use of any trademarks in this report is not intended in any way to infringe on the rights of the trademark holder.

This document is available through Asset Source for Software Engineering Technology (ASSET): 1350 Earl L. Core Road; PO Box 3305; Morgantown, West Virginia 26505 / Phone: (304) 284-9000 or toll-free in the U.S. 1-800-547-8306 / FAX: (304) 284-9001 World Wide Web: <http://www.asset.com> / e-mail: sei@asset.com

Copies of this document are available through the National Technical Information Service (NTIS). For information on ordering, please contact NTIS directly: National Technical Information Service, U.S. Department of Commerce, Springfield, VA 22161. Phone: (703) 487-4600.

This document is also available through the Defense Technical Information Center (DTIC). DTIC provides access to and transfer of scientific and technical information for DoD personnel, DoD contractors and potential contractors, and other U.S. Government agency personnel and their contractors. To obtain a copy, please contact DTIC directly: Defense Technical Information Center / Attn: BRR / 8725 John J. Kingman Road / Suite 0944 / Ft. Belvoir, VA 22060-6218 / Phone: (703) 767-8274 or toll-free in the U.S.: 1-800-225-3842.

Table of Contents

Preface	iii
Securing Network Servers	1
1. Develop a computer deployment plan that includes security issues.	7
2. Keep operating systems and applications software up to date.	13
3. Offer only essential network services and operating system services on the server host machine.	17
4. Configure computers for user authentication.	19
5. Configure computer operating systems with appropriate object, device, and file access controls.	23
6. Identify and enable system and network logging mechanisms.	25
7. Configure computers for file backups.	31
8. Protect computers from viruses and similar programmed threats.	35
9. Configure computers for secure remote administration.	37
10. Allow only appropriate physical access to computers.	39

Preface

This document is one of a series of publications of the Software Engineering Institute at Carnegie Mellon University called *security improvement modules*. They are intended to provide practical guidance to help organizations improve the security of their networked computer systems.

Module structure

Each module addresses an important but relatively narrowly defined problem in network and system security. The first section of the module describes the problem and outlines a set of *security improvement practices* to help solve it. Each practice is a recommended way of performing common tasks related to the secure operation of networked computer systems.

The remaining sections of the module are detailed descriptions of the practices. Each includes a rationale for the recommended actions and a description of how to perform them.

Intended audience

The practices are primarily written for system and network administrators whose day-to-day activities include installation, configuration, and maintenance of the computers and networks. Occasionally, practices are written to assist the managers responsible for network and system administration.

Revised versions

Network and system technologies continue to evolve rapidly, leading to new security problems and solutions. Modules and practices need to be revised occasionally, so to permit more timely publication of new versions, we also publish them on the World Wide Web. At the end of each section of this document is the URL of its Web version.

Implementation details

How an organization adopts and implements the practices often depends on the networking and computing technologies it uses. For some practices, technology-specific implementation details are published on the World Wide Web. The Web version of each practice contains links to the implementation details.

Acknowledgments

This report and the effort to produce it were sponsored by the SEI primary sponsor and the U.S. Army Land Information Warfare Activity (LIWA) ACERT.

The authors acknowledge contributions made to this report by Greg Gravenstreter of the Software Engineering Institute.

Securing Network Servers

The development of computer networks has resulted in an important class of computers: network servers. The primary purpose of these machines is to provide services, including both computational and data services, to other computers on the network.

Because of their service role, it is common for servers to store many of an organization's most valuable and confidential information resources. They also are often deployed to provide a centralized capability for an entire organization, such as communication (electronic mail) or user authentication. Security breaches on a network server can result in the disclosure of critical information or the loss of a capability that can affect the entire organization. Therefore, securing network servers should be a significant part of your network and information security strategy.

Many security problems can be avoided if servers and networks are appropriately configured. Default hardware and software configurations, however, are set by vendors who tend to emphasize features and functions more than security. Since vendors are not aware of your security needs, you must configure new servers to reflect your security requirements and reconfigure them as your requirements change.

The practices recommended here are designed to help you configure and deploy network servers that satisfy your organization's security requirements. The practices may also be useful in examining the configuration of previously deployed servers.

A note on terminology

The term “server” is used in this module to mean the combination of the hardware, operating system, network service, application software, and network connection. When it is necessary to be more specific, we explicitly mention one of those five components.

Although this module focuses on securing network servers, many of the practices are also applicable to securing workstations or other computers on your network. To make it easier for these practices to appear in other modules, we use the word “computer” to mean workstations, servers, or other computers.

Who should read these practices

These practices are applicable to your organization if

- you operate or plan to operate a networked system of workstations that depend on servers for information or computation services

- you operate or plan to operate a public network server connected to an external network such as the Internet

We assume that you have the following security requirements for information resources stored on those servers:

- Some or all of the information is sensitive or proprietary. Access must be limited to authorized and properly authenticated users (inside or outside your organization).
- The integrity of that information is critical. It must not be compromised; that is, not modified by unauthorized users or processes operating on their behalf.
- That information must be readily accessible by authorized users whenever they need it in the course of their work.

We also assume that you have these security requirements for the capabilities provided by those servers:

- Only authorized and properly authenticated users may use these capabilities.
- Users must be able to access these capabilities quickly.

What these practices do not cover

These practices help you appropriately configure a network server and its operating system when it is first deployed. They do not attempt to address security issues for

- the particular network service that will be provided by that server. We expect to address the major network services in other modules. See, for example, the module *Security for a Public Web Site* [Firth 97b].
- the network to which the server is attached
- day-to-day operations of servers. We address operations in other modules; for example, activities related to detecting signs of intrusion are covered in the modules *Preparing to Detect Signs of Intrusion* [Kochmar 98] and *Detecting Signs of Intrusion* [Firth 97a].
- desktop machines that may provide some kind of network service but that are primarily used as workstations. Workstation security is addressed in the module *Securing Desktop Workstations* [Simmel 99].
- security of network services provided by third parties under contract with your organization. See the module *Security of Information Technology Service Contracts* [Allen 98].

We assume that the reader is already capable of performing the initial setup of the computer: unpacking it, confirming the hardware configuration, installing the default operating system, and attaching the network connection, so we do not explicitly cover those aspects of configuring a network server. We note, however, that *some of the practices are most effective if performed during the process of installing the operating system.*

Also, these practices do not address in detail the physical security of servers (such as protection from theft and natural disasters).

Security issues

There are three major security issues related to network servers:

1. **Confidentiality** - Maintaining the confidentiality of information stored on the servers. This includes
 - ensuring that only authorized users can access the services and information

- ensuring that authorized users can access only the services for which they are authorized
2. **Integrity** - Maintaining the integrity of information stored on the servers. This includes
 - ensuring that you can recognize and recover from breaches of integrity
 3. **Availability** - Maintaining the availability of the services. This includes
 - ensuring that services are uninterrupted even when there are hardware or software failures or during routine system maintenance
 - ensuring that you can recognize and recover from security incidents in a timely manner

The common security requirements of confidentiality, integrity, and availability, described above, can be especially critical for network servers. For example:

- File servers and database servers often are used to store your organization's most important information resources, which demand strict confidentiality. They may also store information used for management decisions or customer billing, which demand a high level of integrity.
- Authentication servers store information about user accounts and passwords; disclosure of such information could compromise all the information on all the hosts on your network.
- Public servers such as World Wide Web servers can be a major component in the strategy your organization uses to represent itself to the public, so the integrity of the information on those servers is critically important.
- Servers used by customers for electronic commerce must be available to prevent loss of revenue.
- Servers that provide essential capabilities for employees of your organization must be reliably available; otherwise people may be unable to work.

There are other aspects of network servers that can make them tempting targets for intruders:

- Public servers often have publicly known host names and IP (Internet Protocol) addresses.
- Public servers may be deployed outside an organization's firewall or other perimeter defenses.
- Servers usually actively listen for requests for services on known ports, and they try to process such requests.
- Servers often do not have a human user who notices signs of unusual activity.
- Servers are often remotely administered, so they willingly accept connections from privileged accounts.
- Servers often are configured to reboot automatically after some kinds of failures, which can offer opportunities for intruders.

Security improvement approach

To secure a network server, we recommend a three-part approach. It requires implementing security practices in these areas:

1. planning and executing the deployment of servers
2. configuring servers to help prevent security incidents

3. maintaining the integrity of the deployed servers

The practices are designed to improve security in two major ways:

- They help to maximize security on each network server host, which provides a backup in case of failure of perimeter defenses. Host security is also a first-line of defense against internal threats, which generally have a higher probability of occurrence than external threats.
- They prepare you to better recognize and recover from security breaches.

Summary of recommended practices

Area	Recommended Practice
Planning deployment	1. Develop a computer deployment plan that includes security issues.
Configuring servers	2. Keep operating systems and applications software up to date. 3. Offer only essential network services and operating system services on the server host machine. 4. Configure computers for user authentication. 5. Configure computer operating systems with appropriate object, device, and file access controls. 6. Identify and enable system and network logging mechanisms. 7. Configure computers for file backups. 8. Protect computers from viruses and similar programmed threats. 9. Configure computers for secure remote administration.
Maintaining server integrity	10. Allow only appropriate physical access to computers.

Abbreviations used in these practices

DNS	Domain Name Service
ftp	file transfer protocol
http	hypertext transfer protocol
IP	Internet Protocol
LAN	local area network
NFS	Network File System
NIS	Network Information System
NTP	Network Time Protocol
smtp	simple mail transfer protocol
TCP	Transmission Control Protocol
WORM	Write Once, Read Many
WWW	World Wide Web

References

- [Allen 98] Allen, Julia, et al. *Security of Information Technology Service Contracts* (CMU/SEI-SIM-003, ADA351646). Pittsburgh, Pa.: Software Engineering Institute, Carnegie Mellon University, 1998. Available at <http://www.cert.org/security-improvement/modules/m03.html>.

- [Firth 97a] Firth, Robert, et al. *Detecting Signs of Intrusion* (CMU/SEI-SIM-001, ADA329629). Pittsburgh, Pa.: Software Engineering Institute, Carnegie Mellon University, 1997. Available at <http://www.cert.org/security-improvement/modules/m01.html>.
- [Firth 97b] Firth, Robert, et al. *Security for a Public Web Site* (CMU/SEI-SIM-002, ADA329626). Pittsburgh, Pa.: Software Engineering Institute, Carnegie Mellon University, 1997. Available at <http://www.cert.org/security-improvement/modules/m02.html>.
- [Kochmar 98] Kochmar, John, et al. *Preparing to Detect Signs of Intrusion* (CMU/SEI-SIM-005). Pittsburgh, Pa.: Software Engineering Institute, Carnegie Mellon University, 1998. Available at <http://www.cert.org/security-improvement/modules/m05.html>.
- [Kossakowski] Kossakowski, Klaus-Peter, et al. *Responding to Intrusions* (CMU/SEI-SIM-006). Pittsburgh, Pa.: Software Engineering Institute, Carnegie Mellon University, 1999. Available at <http://www.cert.org/security-improvement/modules/m06.html>.
- [Simmel 99] Simmel, Derek, et al. *Securing Desktop Workstations* (CMU/SEI-SIM-004). Pittsburgh, Pa.: Software Engineering Institute, Carnegie Mellon University, 1999. Available at <http://www.cert.org/security-improvement/modules/m04.html>.

Where to find updates

The latest version of this module is available on the Web at URL

<http://www.cert.org/security-improvement/modules/m07.html>

1

Develop a computer deployment plan that includes security issues.

Most deployment plans address the cost of the computers, schedules to minimize work disruption, installation of applications software, and user training. In addition, you need to include a discussion of security issues.

Why this is important

You can eliminate many networked systems vulnerabilities and prevent many security problems if you securely configure computers and networks before you deploy them. Vendors typically set computer defaults to maximize available functions so you usually need to change defaults to meet your organization's security requirements.

You are more likely to make decisions about configuring computers appropriately and consistently when you use a detailed, well-designed deployment plan. Developing such a plan will support you in making some of the hard trade-off decisions between usability and security.

Consistency is a key factor in security as it fosters predictable behavior. This will make it easier for you to maintain secure configurations and help you to identify security problems (which often manifest themselves as deviations from predictable, expected behavior). Refer to the practice, "Keep operating systems and applications software up to date."

How to do it

Make the decisions described below and then record them.

Note: We assume that you are deploying workstations and servers in an existing infrastructure, which includes an existing network. The security issues related to the network architecture, including where you place servers and workstations on the network, are outside the scope of this practice.

➤ *Identify the purpose of each computer.*

Document how the computer will be used. Consider the following:

- What categories of information will be stored on the computer?
- What categories of information will be processed on the computer (but retrieved from and stored on another computer)?
- What are the security requirements for that information?
- What network service(s) will be provided by the computer?
- What are the security requirements for those services?

➤ *Identify the computer hardware and operating system requirements.*

Document the needed hardware, including the processor architecture, memory requirements, secondary storage requirements (such as hard disk drives and removable-medium drives), networking requirements (such as modems or Ethernet cards), display monitor, and other audiovisual requirements (such as video cameras or microphones).

Document the hardware configuration. This will help you select and securely configure the software.

Document the operating system features needed, even if you are confined to using one vendor's operating system. Include the requirements for both general security features (such as capabilities for user authentication and file access controls) and for special security features (such as an encrypting file system, or a built-in feature to erase memory and disk blocks before reallocating them).

Do not purchase or deploy systems that fail to meet your security requirements.

➤ *Identify the network services that will be provided on the computer.*

The network services you list in your deployment plan may include electronic mail, access to corporate databases, and access to the World Wide Web. For each service, document whether the computer will be configured as a client, a server, or both. Include also peer-to-peer network services, such as file sharing.

Clients: Workstations are normally configured as clients for several network services. You should document the planned behavior of those clients: the levels of access required, the type of access (read, write, etc.), and other aspects of the configurations required for client software.

Servers: As a general rule, a network server should be dedicated to a single service. This usually simplifies configuration, which reduces the likelihood of configuration errors. It also can eliminate unexpected and unsafe interactions among the services that present opportunities for intruders.

In some cases, it may be appropriate to offer more than one service on a single host computer. For example, the server software from many vendors combines the file transfer protocol (FTP) and the hypertext transfer protocol (HTTP) services in a single package. For some organizations, it is appropriate to provide access to public information via both protocols from the same server host.

➤ *Identify the network service software, both client and server, to be installed on the computer.*

Many operating systems include network service software for both clients and servers. You may simply choose to use those packages. For major services, however, there are several choices that vendors may provide. When making your choice, pay special attention to the ability of candidate packages to meet your security requirements, and document your selection.

➤ *Identify other application or utility software that will be installed on the computer.*

List in your deployment plan not only user-oriented application software, but also system-related software and security-related software. The module *Preparing to Detect Signs of Intrusion* [Kochmar 98] provides details on selecting some kinds of security-related software.

➤ *Identify the users or categories of users of the computer.*

For workstations, you will sometimes be able to identify an individual who will be the primary user; but more often, you will have to define categories of users. The categories are based on user roles that reflect their authorized activity. The roles are often based on similar work assignments and similar needs for access to particular information resources—system administrators, software developers, data entry personnel, etc. If appropriate, include categories of remote users and temporary or guest users.

For network servers, document the categories of users that will be allowed access to the provided services. For public servers connected to the Internet, the category of users is probably everyone. For internal servers, you may need to categorize users by their organizational department, physical location, or job responsibilities. You also need a category of administrative users who will need access to administer the network server and possibly one for backup operators.

In general, you should prevent the use of a network server as a workstation. This will ensure that its users are only those who are accessing the provided service (almost always from another computer on the network) and those responsible for server administration.

➤ *Determine the privileges that each category of user will have on the computer.*

Documenting privileges typically requires creating a matrix of users (defined in the previous step) and privileges. The privileges are customarily in groups that define what system resources or services a user can read, write, change, execute, create, delete, install, remove, turn on, or turn off.

➤ *Decide how users will be authenticated and how authentication data will be protected.*

For workstations, it is common to authenticate users via the authentication capability provided with the operating system.

For network servers, there are usually two kinds of authentication: (1) that provided with the operating system, commonly used for authenticating administrative users and (2) authentication provided by the network service software, commonly used for authenticating users of the service. A particular software implementation of a network service may use the provided authentication capability, and thus it may be necessary for users of that service to have a local identity (usually a local account) on the server.

Authentication mechanisms can be both procedural and technological. The most common approach is the use of passwords; but other mechanisms can be used, such as keys, tokens, and biometric devices (devices that recognize a person based on biological characteristics such as fingerprints or patterns of the retinal blood vessels).

Because authentication mechanisms like passwords require information to be accessible to the authentication software, carefully document how that information will be protected. Authentication data is critical security information that requires a high level of protection.

➤ *Determine how appropriate access to information resources will be enforced.*

For many resources, such as program and data files, the access controls provided by the operating system are the most obvious means to enforce access privileges. Also, consider using encryption technologies to protect the confidentiality of information. In some cases, protection mechanisms will need to be augmented by policies that guide user's behavior related to their workstations.

- *Develop intrusion detection strategies for the computer.*

Many of the common intrusion detection methods depend on the existence of various logs that your systems produce and on the availability of auditing tools that analyze those logs. In your deployment plan, describe the kinds of information that will be collected on each computer in support of security. This will help you install the appropriate software tools and configure these tools and the operating system to collect the necessary information.

This topic is elaborated in the security improvement module *Preparing to Detect Signs of Intrusion* [Kochmar 98].

- *Write a plan for backup and recovery of information resources stored on the computer.*

Possessing recent, secure backup copies of information resources makes it possible for you to quickly restore the integrity and availability of information resources. Successful restoration depends on configuring the operating system, installing appropriate tools, and following defined operating procedures. You need to document backup processes; roles, responsibilities, and how the physical media that store the backup data are handled, stored, and managed. Consider using encryption technologies to protect the backups.

For some network servers, such as those providing public services like the World Wide Web, it is common to develop the information content of those services on a different host machine. The authoritative version of this content is maintained (and backed up) on this second computer, and then transferred to the public server. This method makes it unnecessary to perform file backups of the server itself. If the information is ever compromised, you can restore it by transferring a copy from the authoritative version. For more information, refer to the practice “Configure computers for file backups.”

- *Determine how network services will be maintained or restored after various kinds of faults.*

To maintain the availability of services essential to your business, you generally need some level of redundancy. For example, you may want to specify when to use hot, warm, and cold backups. Hot backups provide the capability to immediately switch configurations because the backup system is run in parallel with the primary system. Warm backups require some degree of reconfiguration before you use them since they are not run in full parallel with operational systems. You must start cold backups from a shutdown state and bring them up to date before using them.

Write a plan to ensure that no single failure (power supply, hardware, software, etc.) will make an essential service unavailable for a period of time you consider unacceptable.

- *Develop and follow a documented procedure for installing an operating system.*

In your procedure, include steps to implement all the decisions you made in the steps above and describe all the parameters that are set during installation.

In many cases, the parameters are recorded in scripts or configuration files that are executed or read during various phases of the installation. Make all your parameter choices explicit, even if they match the vendor’s current default settings. (This may seem to be unnecessary, but it can prevent security problems if you subsequently reuse your scripts or configuration files to configure future workstations and servers.) Your explicit choices will still be used even if the vendor’s defaults have changed with new releases. Your installation procedure should also specify the vendor’s security-related updates or patches that are to be applied to the operating system.

➤ *Determine how the computer will be connected to your network.*

There are concerns relating to network connections that can affect the configuration and use of any one computer.

LANs: Many organizations use a fast networking technology such as Ethernet for their local area networks. In these cases, information traversing a network segment can be seen by any computer on that segment. This suggests that you should only place “trusted” computers on the same network segment, or else encrypt information before transmitting it.¹

Modems: Modems permit direct connectivity between one of your computers (and thus, potentially, your internal network) and the external networks reachable by the public telephone network. Many organizations forbid users to attach a modem to a workstation.

It is also important to document the use of modems on a network server. As a general rule, do not attach modems to any servers other than those whose purpose is to provide dial-in access.

➤ *Identify the security concerns related to day-to-day administration of the computer.*

If your organization is small, it may be feasible to administer both workstations and network servers individually from their consoles. We recommend this method because it is the most secure.

In most cases, however, workstations and servers are some distance from the offices of the system administrators. As a result, a significant amount of day-to-day administration is done from the administrator’s workstation via the network. Permitting remote administration in a secure manner typically requires configuring the operating system and installing various administration software tools, possibly including tools to encrypt administration commands and data between the target computer and the administrator’s workstation. Thus, you need to document your administration processes to configure the computer appropriately.

Note that a detailed administration plan is not required here. Such a plan is necessary for a well-run organization, but it addresses a broader range of issues than just security and is therefore outside the scope of this practice.

➤ *Identify actions to protect information contained on hardware that is no longer in use.*

Determine what steps you need to take to ensure that the information contained on hardware being updated, replaced, removed from service, or disposed of is eliminated to the extent possible. For example, erase and reformat disks, rewrite tapes, and clear firmware passwords. The extent of your actions is dependent upon the sensitivity of the information. You may need to physically destroy hardware containing highly sensitive information to ensure that the hardware cannot be used and that the information cannot be accessed.

1. However, note that network protocols based on TCP/IP require at least part of the information in a packet (source, destination, port) to be unencrypted, which exposes the network to traffic analysis.

Policy considerations

Your organization's security policy for networked systems should

- require that a detailed computer deployment plan be developed and followed whenever computers are being deployed (or redeployed)
- require that access to your deployment plan be permitted to only those who require the information to perform their jobs

Where to find updates

The latest version of this practice, plus implementation details for selected technologies, is available on the Web at URL

<http://www.cert.org/security-improvement/practices/p025.html>

2 ***Keep operating systems and applications software up to date.***

You need to stay informed of vendors' security-related updates to their products, which may be called updates, upgrades, patches, service packs, or hot fixes. Whenever an update is released, you need to evaluate it, determine if it is applicable to your organization's computers, and, if so, install it.

Why this is important

Because software systems are so complex, it is common for security-related problems to be discovered only after the software has been in widespread use. Although most vendors try to address known security flaws in a timely manner, there is normally a gap from the time the problem is publicly known, the time the vendor requires to prepare the correction, and the time you install the update. This gap gives potential intruders an opportunity to take advantage of this flaw and mount an attack on your computers and networks. To keep this time interval as short as possible, you need to stay aware of

- announcements of security-related problems that may apply to your systems
- immediate steps you can take to reduce your exposure to the vulnerability, such as disabling the affected software
- permanent fixes from vendors

Installing applicable vendors' updates as soon as they are available can reduce your vulnerability to attack.

How to do it

- *Develop a list of sources of information about security problems and software updates for your system and application software.*

The most common sources of current information include Web sites of vendors and computer- and network-security organizations¹. There are also mailing lists, some of which are sponsored by vendors, and USENET news groups.

See *Preparing to Detect Signs of Intrusion* [Kochmar 98], specifically the implementation "Maintaining currency by periodically reviewing public and vendor information sources." This implementation is available at <http://www.cert.org/security-improvement/implementations/i040.01.html>.

- *Establish a procedure for monitoring those information sources.*

In the case of mailing lists, you usually receive announcements about security problems and software updates soon after they are available. Web sites vary considerably in the

1. For example, the CERT/CC site at URL <http://www.cert.org>.

timeliness of their announcements, so you need to decide how often to look for information there. Some of the news-oriented Web sites are updated one or more times a day, so daily monitoring is a good idea.

➤ *Evaluate updates for applicability to your systems.*

Not all updates are applicable to the configuration of the computers and networks in your organization and to your organization's security requirements.

Evaluate all the updates to determine their applicability, and weigh the cost of deploying an update against the benefits.

➤ *Plan the installation of applicable updates.*

The installation of an update can itself cause security problems:

- During the update process, the computer may temporarily be placed in a more vulnerable state.
- If the update is scheduled inappropriately, it might make a computer or information resources unavailable when needed.
- If an update must be performed on a large number of computers, there can be a period of time when some computers on the network are using different and potentially incompatible versions of software, which might cause information loss or corruption.
- The update may introduce new vulnerabilities.

Updates can also cause a number of problems in other installed software. You may want to consider running a previously developed regression test suite to compare current performance with past performance. Another approach is to install the update in an isolated test environment and run a series of user trials before releasing the update on your operational systems.

A number of software packages exist that will tell you the differences to the system just prior to and after installing the update. We recommend that you use one of these to fully understand and analyze the effects of the update on your systems.

In addition, you should always backup your system prior to applying any updates.

Update approaches that depend on an administrator physically visiting each computer are labor intensive but will work for networks with a small number of computers. You will need to employ automated tools to roll-out updates to a large number of computers. Some of these tools are provided by vendors for their specific products. You may need to develop tools that are tailored to your environment if vendor tools are insufficient. When using automated tools to roll-out updates, the affected computers and the network are likely to be vulnerable to attack during the update process. To lessen this vulnerability, you should use only an isolated network segment when propagating the updates.

➤ *Install the updates using a documented plan.*

Follow the plan developed in the previous step. This helps ensure that you deploy computers consistently throughout your organization.

➤ *Deploy new computers with up-to-date software.*

When new workstations and network servers are being deployed, it is common to install the operating system and other software from the original distribution media supplied by vendors. However, those software versions may not include recent security-related

updates. Maintain an archive of updates that you have evaluated and chosen to install on existing computers, so that you can install them on new computers before deployment.

Also acquire and install the most up-to-date driver software (often available from vendors' Web sites) for all components and peripheral devices. Those drivers typically address performance and security issues that have been discovered since the components were packaged and shipped from the factory. Be sure to read all the release documentation associated with the updated drivers before using them. Also, whenever possible, verify the integrity and authenticity of the new driver software, using methods such as cryptographic checksums supplied by the vendor.

- *After making any changes in a computer's configuration or its information content, create new cryptographic checksums or other integrity-checking baseline information for that computer.*

Refer to the modules *Detecting Signs of Intrusion* [Firth 97a] and *Preparing to Detect Signs of Intrusion* [Kochmar 98] for additional information on the role of checking the integrity of baseline information to support intrusion detection.

Policy considerations

Your organization's security policy for networked systems should require that systems administrators install necessary security-related software updates in a timely manner.

Other information

We understand that you may not have sufficient information to decide whether or not to apply an update and that you may not have a comprehensive test environment within which to evaluate the effects of an update. We recommend that you implement the steps in this practice to the extent possible and practical.

Where to find updates

The latest version of this practice, plus implementation details for selected technologies, is available on the Web at URL

<http://www.cert.org/security-improvement/practices/p027.html>

3

Offer only essential network services and operating system services on the server host machine.

Ideally, each network service should be on a dedicated, single-purpose host. Many computers are configured by default to provide a wider set of services and applications than required to provide a particular network service, so you may need to configure the server to eliminate or disable them.

Why this is important

Offering only the essential network services on a particular host can enhance your network security in several ways:

- Other services cannot be used to attack the host and impair or remove desired network services.
- Different services may be administered by different individuals. By isolating services so each host and service has a single administrator, you will minimize the possibility of conflicts between the administrators.
- The host can be configured to better suit the requirements of the particular service. Different services might require different hardware and software configurations, which could lead to needless vulnerabilities or service restrictions.
- By reducing services, the number of logs and log entries is reduced so detecting unexpected behavior becomes easier.

How to do it

We recommend that you use the configuration principle “deny first, then allow.” That is, turn off as many services and applications as possible and then selectively turn on those that are essential.

- *Determine the functions that you intend to support with your network server.*

The services you enable on a selected host depend on the functions you want the host to provide. Functions could support the selected network service, other services hosted on this computer, or development and maintenance of the operating system and applications.

Determine the configuration of the host for

- file systems (e.g., whether any file services will be used by this host)
- system maintenance (e.g., with multiuser systems, whether all maintenance will be done only via the console or remotely)
- server maintenance (e.g., if all the maintenance will be done on another host and only the resulting files downloaded to this host, there is no need to provide any compilers or editors on this host)
- network configuration (DNS vs. NIS)
- protocols offered (IP, IPX, Appletalk, DecNet, etc.)

- *If there are alternative ways of providing the same function, select the more secure way.*

For example, on UNIX systems, remote system maintenance (i.e., not from the console) can be provided using *remote shell (rsh)* or *secure shell (ssh)* capabilities; of the two, **ssh** is more secure.

- *Once you determine the minimal set of services and applications, ensure that only those are available on the host.*

Either do not install unnecessary services or turn the services off and remove the corresponding files (and any other unnecessary files) from the host. Be careful with network service programs. Some provide multiple services (such as combined HTTP and FTP services) and you will have to reconfigure them or disable unneeded services.

When considering services to enable or disable, administrators typically think of those services that run as processes. This includes, for example, telnet, FTP, electronic mail, and Web services. However, most of today's systems also provide services directly from the kernel. An example would be a netmask request. That request is typically broadcast onto the local area network, and all systems that see that request answer it, if not otherwise instructed. The kernel of those answering systems is providing the netmask service, more than likely unbeknownst to the administrator of that computer.

You need to determine what services are provided by the kernel and what controls the operating system provides to configure those services. These services are frequently not documented and are often not controllable. There is no tool that we know of to test for the presence of such services in a manner similar to the way the strobe tool for UNIX systems tests for services running as processes. The best source of information is the system vendor.

We recommend that you configure computers to offer only the services that your deployment plan specifies they should provide. Providing multiple services or combining the role of workstation and server on the same machine makes it harder to maintain security.

- *After you make all configuration choices, create and record cryptographic checksums or other integrity-checking baseline information for your critical system software.*

Refer to the practice "Generate information required to verify the integrity of your systems and data." in the module *Preparing to Detect Signs of Intrusion* [Kochmar 98].

Policy considerations

Your organization's security policy for networked systems should require that network servers be configured to offer only essential services.

Other information

You may need to configure the server differently according to the other features provided by the selected host operating system or environment. For example, certain operating systems provide extensive access control mechanisms that minimize or prevent the possibility of unauthorized access at relatively fine levels of granularity.

Where to find updates

The latest version of this practice, plus implementation details for selected technologies, is available on the Web at URL

<http://www.cert.org/security-improvement/practices/p012.html>

4 *Configure computers for user authentication.*

An organization's security policy for networked systems should specify that only authorized users may access the computers. To enforce this, you need to configure the computer to authenticate a prospective user, who must prove that they are authorized for such access.

Configuring the computer for authentication usually involves configuring parts of the operating system, firmware, and applications such as the software that implements a network service. If your organization has authentication servers, configuring a new workstation or network server for user authentication may require you to make configuration changes on another computer. In special cases, you may also use authentication hardware such as tokens, one-time password devices, or biometric devices (devices that can recognize a person based on biological characteristics, such as fingerprints or patterns in retinal blood vessels).

Why this is important

Unauthorized users can jeopardize the security of information stored on or accessible from a computer. To prevent this, you must configure the computer to authenticate all users who attempt access.

How to do it

This practice is most effective if you include it as part of the initial installation and configuration of the operating system.

Your deployment plan documents the users or user categories and the approach to authenticating those users. The following steps describe how to implement that part of the plan.

- *Configure the system to use available authentication capabilities.*

If the computer's firmware offers the feature of requiring a password when the system is turned on, enable that feature and set the password. This feature is sometimes known as a BIOS or EEPROM password.

Enabling this feature will require your intervention if the system crashes because you can't configure the computer to restart automatically. This is usually acceptable for workstations because if the user is not present, it is not necessary to restart the computer immediately. However, enabling this feature can present problems for network servers, which normally operate 24 hours a day. When the system crashes, an administrator may not be available to restart the system.

➤ *Remove unneeded default accounts and groups.*

The default configuration of the operating system often includes guest accounts, administrator accounts, and accounts associated with local and network services. The names and passwords for those accounts are well known. Remove or disable unnecessary accounts to eliminate their use by intruders.

➤ *Change default passwords.*

For default accounts that you want to keep on the system, change the passwords to make it harder for intruders to compromise the accounts. Also disable passwords for accounts that need to exist but do not require an interactive login.

➤ *Create the user groups for the particular computer.*

Assign users to the appropriate groups. Then assign rights to the groups, as documented in your deployment plan. This approach is preferable to assigning rights to individual users.

➤ *Create the user accounts for the particular computer.*

Your deployment plan identifies who will be authorized to use each computer and its services. Create only the necessary accounts. Check your password policy, and set account passwords appropriately.

A password policy should address

- length: a minimum length for passwords. It is common to specify a minimum length of eight characters.
- complexity: the mix of characters required. It is common to require passwords to contain both uppercase and lowercase letters and at least one nonalphanumeric character.
- aging: how long a password may remain unchanged. It is common to require users to change their passwords at least once a month. The policy should permit users to do so only through approved authentication mechanisms.
- reuse: whether a password may be reused. Some users try to defeat a password aging requirement by changing the password to one they have used before.
- authority: who is allowed to change passwords

Finally, if you have retained any of the default administrator accounts, consider changing their names.

➤ *Ensure users follow your password policy.*

Document your password policy, communicate it to users, and train them to always follow the policy.

Configure the password-setting software to reject passwords that don't conform to your policy, if the operating system provides this feature.

➤ *Configure computers to require reauthentication after idle periods.*

This step is most useful for workstations, but consider it for network servers as well, especially if the server will be administered from the console.

Most operating systems include software to display a changing image (screensaver) on a monitor or software to power down monitors and disks (energy saver) after a short period of inactivity.

This inactivity may indicate that the workstation is unattended though a user is still logged in. Requiring reauthentication when the user returns prevents an unauthorized person from using an active session while the authorized user is away.

If possible, configure the operating system to terminate a session (log out) after a specified idle period (typically 15 minutes). Alternatively, install a third-party “locking screen saver” to do the same thing.

Consider requiring users to shut down or lock workstations when they leave the machine unattended. This prevents a period of vulnerability between the time the user leaves and the time the locking screensaver is activated.

- *Configure computers to deny login after a small number of failed attempts.*

It is relatively easy for an unauthorized user to gain access to a computer by using automated software tools that attempt all passwords. If your operating system provides the capability, configure it to deny login after three failed attempts. Typically, the account is “locked out” for a period of time (such as 30 minutes) or until a user with appropriate authority reactivates it.

This is another situation that requires you to make a decision that balances security and convenience. Implementing this recommendation can help prevent some kinds of attacks, but it can also allow a malicious intruder to make failed login attempts to eliminate user access - a denial of service condition. You may not consider this configuration acceptable for network servers because it makes the server unavailable to the authorized user whose account was compromised.

In some cases, you need to distinguish between failed login attempts at the console and those coming in through the network. Failed network login attempts should not prevent an authorized user or administrator from logging in at the console.

Note that all failed login attempts should be logged. (Refer to the security improvement modules *Detecting Signs of Intrusion* [Firth 97a] and *Preparing to Detect Signs of Intrusion* [Kochmar 98].)

- *Install and configure other authentication mechanisms as required by your organization’s security plan and policies.*

Consider using other authentication mechanisms such as tokens, one-time password systems, and biometric hardware and software. They can be expensive, but they may be justified in some circumstances.

- *For network servers, configure the authentication capability of the network service software, if any.*

The authentication capabilities of network service software packages vary, and we plan to address them in detail in future modules on specific network services. However, note that some packages provide their own mechanisms for authenticating users, while others depend on the underlying operating system. Be sure that both are configured appropriately.

Policy considerations

Your organization’s policy for networked systems should

- describe under what conditions an account is created and deleted. This should include what account actions are taken (disabled, deleted, transferred) and how files are

handled when an employee, contractor, or vendor who has an account on your systems no longer works for your organization.

- require appropriate authentication of all users on all computers that can access information resources; this includes authenticating users of network services hosted by your servers
- include an appropriate password policy
- prohibit users from recording and storing passwords in places that could be discovered by intruders

Your organization's acceptable use policy for workstations should require that users shut down or lock their unattended workstations.

When writing a password policy, remember that requiring users to have complex passwords may have the undesired result of the user's writing the passwords on paper that they keep near the computer (often stuck to the machine) or with personal papers (in a wallet, purse, or briefcase). If that paper is observed, lost, or stolen, it creates a potential vulnerability.

If a password policy is especially difficult to follow, it creates in users a desire to find ways around it. This attitude can negatively influence users' compliance with other aspects of security policies.

Where to find updates

The latest version of this practice, plus implementation details for selected technologies, is available on the Web at URL

<http://www.cert.org/security-improvement/practices/p028.html>

5

Configure computer operating systems with appropriate object, device, and file access controls.

Many operating systems provide the capability to specify access privileges individually for files, devices, and other data or code objects. We recommend that you configure the settings on files and other objects to take advantage of this capability and protect information stored on the computer.

Why this is important

By carefully setting access controls, you can reduce both intentional and unintentional security breaches. For example, denying read access helps to protect confidentiality of information, and denying unnecessary write access can help maintain the integrity of information. Limiting the execution privilege of most system-related tools to authorized system administrators can prevent most users from making configuration changes that could reduce security. It also can restrict the ability of intruders to use those tools to attack the system or other systems on the network.

How to do it

Note that access controls should be implemented during initial installation and configuration of the operating system and carefully maintained thereafter.

► *Identify the protection needed for files, devices, and objects on the computer.*

One method that you can use to identify needed protection is to construct a matrix with categories of files and objects on one axis and groups of users (defined by roles and access authority) on the other. Then record in the matrix the kinds of access privileges allowed for that class of objects and that class of users. The privileges are based on the security requirements (such as confidentiality, integrity, and availability) of the various classes of resources.

For example, you may have file categories that include administrative information (user names, passwords, privileges, etc.), applications, development tools, operating system files, and user data files. The latter may be further subdivided into categories such as customer accounts, inventory records, research data, and management reports. You may have user groups that include system administrators, network service daemons, and users from various departments.

As you begin to identify privileges, you may need to split some rows and columns. This happens, for example, when you discover that a single group of users is really two groups because their need to access a particular resource is not uniform.

You may also want to distinguish local access privileges from network access privileges for a class of files.

Application programs may request and be granted increased access privileges for some of their operations—a change that is not obvious to the users of that application. You may not want all those users to have increased privileges. Therefore, it is important to take great care in assigning privileges to users and groups.

➤ *Create the needed user groups.*

When you take the previous step, you may identify categories of users not documented in the computer deployment plan with enough detail. Configure the operating system to recognize the needed user groups, and then assign individual users (including network service daemons) to the appropriate groups.

Carefully consider whether to retain a guest account or group and if you do, consider greatly limiting its access.

➤ *Configure the access controls.*

Configure the access controls for all protected files, devices, and other objects, using the matrix created in the first step above as a guide.

Pay attention to access control inheritance when defining categories of files and users. Ensure that you configure the operating system so that newly created files and directories inherit appropriate access controls, and that access controls are propagated down the directory hierarchies when you assign them.

➤ *Install and configure file-encryption capabilities for sensitive data.*

Some operating systems provide optional file encryption; there are also third-party file-encryption packages available. These may be useful if the operating system's access controls are insufficient for maintaining the confidentiality of file contents. This can be the case if the operating system provides few or no access control features, or when the relationships among categories of files and categories of users are so complex that it would be difficult to use only access controls to administer the security policy.

Note that this recommendation pertains only to encryption of files stored on the computer itself. Encryption of information for transmission over a network is a separate issue.

Policy considerations

Your organization's security policy for networked systems should

- specify access limitations for the information that will be stored on computers
- how access to files that have been encrypted with a user key is performed. This is particularly important when a user no longer works for your organization.

Other information

Some operating systems provide more than one file system with different access-control capabilities. It is important to choose the file system that best meets your needs for file access control. Your decision may affect the low-level formatting of storage devices and thus should be made early in the process of configuring the operating system.

Where to find updates

The latest version of this practice, plus implementation details for selected technologies, is available on the Web at URL

<http://www.cert.org/security-improvement/practices/p029.html>

6 ***Identify and enable system and network logging mechanisms.***

Collecting data generated by system, network, application, and user activities is essential for analyzing the security of these assets and detecting signs of intrusion. Log files contain information about past activities. Different systems provide various types of logging information; some systems do not collect adequate information in their default condition. You should identify the types of logs and logging mechanisms available for each system asset (file access logs, process logs, network logs, application-specific logs, etc.), identify the data recorded within each log, and then enable the collection of the desired data.

Why this is important

Log files are often the only record of suspicious behavior. Failure to enable the mechanisms to record this information and use them to initiate alert mechanisms will greatly weaken or eliminate your ability to detect intrusion attempts and to determine whether or not they succeeded. Similar problems can result from not having the necessary procedures and tools in place to process and analyze your log files.

You may need your logs to

- alert you that an intrusion is occurring
- help you recover your systems
- help you to conduct an investigation
- provide information required for legal proceedings
- provide information required for insurance claims

How to do it

- *Identify the information to be logged.*

Identify

- types of information you can log
- mechanisms used for logging
- locations where the logging is performed
- locations where the log files are stored

A table of log categories and types of log information within each category are listed below. You may want to use this list as a guide to the types of information to log (although not all systems are able to log every type on the list). Tailor logging selections to meet your site's needs.

Log Category	Types of information to log
Users	<ul style="list-style-type: none"> • Login/logout information: location and time of failed attempts, attempted logins to privileged accounts • Changes in authentication status, such as enabling privileges
Processes	<ul style="list-style-type: none"> • Real and effective user executing the process • Process start-up time, arguments • Process exit status, time, duration, resources consumed
Systems	<ul style="list-style-type: none"> • Actions requiring special privileges • Status/errors reported by hardware and software subsystems • Changes in system status, including shutdowns and restarts
Networks	<ul style="list-style-type: none"> • Service initiation requests • The name of the user/host requesting the service • Network traffic • New connections • Connection duration
File Systems	<ul style="list-style-type: none"> • Changes to access control lists and file protections • File accesses (opening, creating, executing, deleting)
Applications	<ul style="list-style-type: none"> • Applications- and services-specific information, e.g., mail logs, FTP logs, Web server logs, modem logs, firewall logs

If possible, do not log passwords, even incorrect ones. Logging correct passwords creates an enormous potential vulnerability if a non-authorized user or intruder accesses audit records. Recording incorrect passwords is also risky as they often differ from valid passwords by only a single character or transposition. Turning off password logging may require resetting a system default. If you cannot turn off password logging, you need to exercise special care in protecting access to log files that contain this information, as described in the last two steps in this practice.

You may want to log data about password use such as the number of failed attempts, accesses to specific accounts, etc.

- *Determine if the logging mechanisms provided with your systems and networks sufficiently capture the required information.*

Determine the logging mechanisms available for the platforms at your site, how the log files are named, and where they are located. The names of these log files can differ even among versions of the same operating system delivered by a single vendor, so it is important that you verify this each time you upgrade your systems.

Identify what types of information each logging mechanism can capture. The combination of mechanisms should capture the information identified in the table categories noted above. There may be differences in the log file contents provided by different vendors, even for similar types of systems.

If the logging mechanisms provided by your vendor are insufficient to capture the data you need, use other tools to capture the additional information. You may need to develop these.

➤ *Enable logging.*

Using the logging mechanisms provided by the vendor and any supplemental tools, enable all logging that you have selected from the previous step. For help, refer to the administration documentation for your systems to learn how to enable each of the logging mechanisms and refer to documentation for other tools used during setup. This documentation will specify whether these mechanisms need to be enabled only once, each time the system is rebooted, or at regular intervals during the system's normal operation. Some logging mechanisms let you select different levels of detail.

Pay attention to the location of the log data: some tools allow you to specify a file or directory where the data is logged while others write their data to a predefined default location. Make sure that you have sufficient space for the data that is generated. Ensure that the logged data is protected correctly, based on previously determined ACLs (access control lists).

Be aware that multiple logging mechanisms may contribute log records to a single log file, such as syslog in UNIX systems.

➤ *Protect logs to ensure they are reliable.*

To protect sensitive information, ensure that log files are protected from being accessed or modified by unauthorized users. Confirm that only authorized users can access utilities that reconfigure logging mechanisms, turn the utilities on and off, and write to, modify, and read log data.

It is important to collect and archive log files so that they cannot be accessed by an intruder to remove or alter signs of an intrusion or add erroneous information. Use the following methods to ensure log files are not modified:

- Log data to a file on a separate host, preferably one in a physically secure location that is not easily accessible from the network. For example, capturing log data using a computer via a dedicated serial line provides a way of storing the log files more securely than if they were written on the logging host's disks.
- Log selected data to a write-once/read-many device (such as CD-ROM or a specially configured tape drive) or to a write-only device (such as a printer) to eliminate the possibility of the data being modified once it is written.
- If supported by your systems, set selected log file attributes that enable only new information to be appended to the log files (i.e., new records can be added, those already recorded cannot be modified).
- Encrypt log files, particularly those that contain sensitive data or those being transmitted across a network.

Logging directly to disk on the local host is easiest to configure and allows instant access to file records for analysis, but it is also the least reliable. Collecting log files on a write-once device requires slightly more effort to configure but is more secure. However, data is not as easily accessed and you need to maintain a supply of storage media.

Printing the logging results is useful when you require permanent and immediate log files, but printed logs can be difficult to search, require manual analysis, and require a potentially large storage space.

When the host generating the logging data is different from the host recording it, you must secure the path between them. For environments where short distances separate the generating host from the recording host, you can attach them with single point-to-point cable(s). For environments where this approach is not practical, minimize the number of

networks and routers used to make the connection or encrypt sensitive log data as it is generated.

You need to prepare systems that perform logging to ensure that they do not stop functioning in the event of a logging denial of service attack. A UNIX example would be an intruder launching an attack that fills up the syslog files so that when the logging partition is full, logging ceases. Two means of preparation are creating separate file partitions for different log information and filtering network messages to decrease the likelihood of such attacks.

► *Document your management plan for handling log files.*

Handle the total volume of logged information. We recommend that you log as much as possible for your systems and networks. While log files can consume a great deal of storage very quickly, it is difficult to anticipate which logs will be critical in the event of an intrusion. Based on your log collection and storage approach, you may want to compress log files to allow them to remain accessible online for easier review and to conserve space.

Determine what logging data is most useful to collect. However, you need to balance the importance of recording system, network, and user activities with the resources available to store, process, review, and secure them. Questions that help you determine the usefulness of logging data include

- What is the host's sole or primary purpose? For example, if a host is acting as a Web server, you want to capture Web logs.
- How many users are assigned to the host or system and how important is it for you to know who is logged on? This helps you decide how much login/logout information to capture.
- How important is it to be able to use your logs to recover a compromised system? This helps you set the priority for capturing information such as data and file transaction logs.
- What are the range of services that can be performed on this host or system? Process accounting information is useful to detect unauthorized services.
- What is your organization's ability and capacity to process and analyze all collected logs to obtain useful information when it is needed?

Rotate log files. This activity consists of

- making a copy of the active (online) log files at regular intervals (ranging from daily to weekly)
- renaming the files so information contained in the file is not further augmented
- resetting file contents
- verifying that logging still works

Rotating log files allows you to limit the volume of log data you have to examine at any given time. It also allows you to keep log files open for a limited duration so that damage is bounded if an active log file is compromised. In this way, you create a collection of log files that contain well-defined time intervals of recorded data. You can then consolidate logs from different systems by matching time intervals. This will help you gain a network-wide perspective on the activities. To perform this consolidation, you will likely need to merge log files from different systems into a central log file and adjust the timestamps used in each to match a master clock.

Back up and archive log files. Move your log files to permanent storage or capture them as part of your regular backup procedure if you want to retrieve them later if the need arises. Document the method you use to access archived log files. Create backups before you execute any automated tools that truncate and reset the log files so that minimal logging data is lost.

Encrypt log files. We recommend encrypting log files that contain sensitive data as the log data is being recorded. Protect the encryption software and place a copy of your encryption keys on a floppy disk or write-only CD-ROM in a secure location such as a safe or safety deposit box. If the keys are lost, the log files cannot be used. If possible, use public key encryption. The logs can be encrypted using the public key (which can be safely stored online) and the corresponding private key (stored off-line) can then be used to decrypt the logs.

Ensure that you have the system and personnel resources necessary to analyze logs on a regular basis and on demand (i.e., in some cases, daily, and as alert events occur).

Dispose of log files. Ensure that all media containing log file data are disposed of in a secure manner (e.g., shredding hardcopy output, sanitizing disks, destroying CD-ROMs).

Policy considerations

Your organization's security policy for networked systems should require that you document a management plan for handling log files. This plan should include what to log, when and why to log, where to log, and who is responsible for all aspects of the plan.

Other information

See the security improvement modules *Detecting Signs of Intrusion* [Firth 97a], *Preparing to Detect Signs of Intrusion* [Kochmar 98], and *Responding to Intrusions* [Kossakowski 99].

Where to find updates

The latest version of this practice, plus implementation details for selected technologies, is available on the Web at URL

<http://www.cert.org/security-improvement/practices/p058.html>

7

Configure computers for file backups.

Before deploying a computer, you need to develop a file backup and restoration plan and configure the computer to implement that plan.

Why this is important

File backups allow you to restore the availability and integrity of information resources following security breaches and accidents. Without a backup, you may be unable to restore a computer's data after system failures and security breaches.

How to do it

- *Develop a file backup and restoration plan.*

Develop a plan that is broad enough to cover all the workstations and servers you plan to deploy.

First, determine what categories of files will be backed up. For example, you may choose to back up only user data files because damaged system files should be reloaded from the original distribution media. In general, you need to make trade-offs among speed of the backup process, the amount of storage needed for the backed-up files, and the effort required to restore one or all files from the backed-up versions.

There are two common technological approaches to file backups for workstations. With the first, files are backed up locally at each workstation, often by the user(s) of that workstation. With the second, backups are centrally administered, with data copied from workstations via networks. The first approach has the advantage of not requiring that protected information traverse the network, which reduces the chances of it being monitored, intercepted, or corrupted. On the other hand, it has the disadvantages of requiring additional storage devices on each workstation, increased efforts to keep the many backup devices and media secure, and training users to perform the backups.

For network servers, a third technological approach is often used. With this approach, the authoritative version of the information content of the server is created and maintained on a secure machine that is backed up. The information is periodically transferred to the server for access by clients. If the server is compromised and its content damaged, it can be reloaded from the secure system maintaining the authoritative version. This approach is typically used for public servers, such as Web servers¹, because the content changes at more predictable intervals than, for example, a network file server that supports user workstations.

1. This approach is elaborated in the security improvement module *Security for a Public Web Site* [Firth 97b].

Determine the appropriate medium to contain your backup files based on your requirements for speed (for both reading and writing), reliability, and storage duration. Media you should consider include magnetic tape, optical disk, and CD-ROM.

The plan should be designed to ensure that backups are performed in a secure manner and that the contents of the backups remain secure.

We recommend that the plan specify that

- the source data is encrypted before being transmitted to the storage medium
- the data remains encrypted on the backup storage media
- the storage media are kept in a physically secure facility that is protected from man-made and natural disasters

➤ *Install file backup tools.*

Select file backup tools to allow you to implement your backup plan. You may need to use third-party software, although the backup capabilities of some operating systems are likely to be sufficient. You may also need to install storage devices, either centrally or on each workstation and server, to store the backup copies.

If you choose central administration and storage of backed-up files, be sure that the chosen tools adequately protect the confidentiality and integrity of information as it travels the network to the backup device. We recommend that you use encryption technologies.

Note that the tools used to recover backed-up files are normally kept offline, rather than on individual workstations and servers. If a computer has been compromised and you need to recover a file, you cannot trust the integrity of any of the tools on that computer.

➤ *Configure the backup tools and initiate the scheduled backups.*

Tool configurations need to reflect your backup and restoration plan. Configure the tools to save access control settings along with file contents, if that feature is available.

Do the first full backup just before deploying the computer, and then confirm that you can perform a full restoration from that backup (Refer to the step *Test the ability to recover from backups* below).

➤ *Confirm that the scheduled backups are being performed successfully.*

In many organizations, file backups are completely automated, so system administrators tend to forget that they are happening. Therefore, confirm that the backup procedures for a newly deployed workstation are actually working.

➤ *Test the ability to recover from backups.*

For many system administrators, recovering a file from a backup is an uncommon activity. This step assures that if you need to recover a file, the tools and processes will work.

Performing this test periodically will help you to discover problems with the backup procedures so you can correct them before losing data.

Some backup restoration software does not accurately recover the correct file protection and file ownership controls. Check these attributes of restored files to ensure they are being set correctly.

Periodically test to ensure that you can perform a full system recovery from your backups.

Policy considerations

Your organization's security policy for networked systems should

- require the creation of a file backup and recovery plan
- inform users of their responsibilities (if any) for file backup and recovery

Other information

Be aware that file backups taken from compromised machines may contain damaged files, services, or other information left behind by an intruder (back doors, Trojan horses). Exercise caution when you use these backups to restore your computers.

Refer to the practices "Eliminate all means of intruder access." and "Return systems to normal operation." found in the module *Responding to Intrusions* [Kossakowski] for a discussion of approaches to consider when you are choosing backup methods.

Where to find updates

The latest version of this practice, plus implementation details for selected technologies, is available on the Web at URL

<http://www.cert.org/security-improvement/practices/p032.html>

8

Protect computers from viruses and similar programmed threats.

There are several kinds of software that can surreptitiously breach computer security, such as a¹

- virus: a code fragment (not an independent program) that reproduces by attaching to another program. It may damage data directly, or it may degrade system performance by taking over system resources, which are then not available to authorized users.
- worm: an independent program that reproduces by copying itself from one system to another, usually over a network. Like a virus, a worm may damage data directly, or it may degrade system performance by consuming system resources and even shutting down a network.
- Trojan horse: an independent program that appears to perform a useful function but that hides another unauthorized program inside it. When an authorized user performs the apparent function, the Trojan horse performs the unauthorized function as well (often usurping the privileges of the user).

You should configure computers to take countermeasures against these threats. In addition, you should establish policies and train users to help prevent these programmed threats from being installed on their workstations.

Why this is important

Programmed threats can cause significant security breaches; confidential information can be captured and transmitted, critical information can be modified, and the software configuration of a computer can be changed to permit subsequent intrusions.

Recovering from programmed threats can be expensive. Installing preventative measures and instituting user training can significantly reduce your exposure to these threats at a fraction of the cost it would take to recover from them.

How to do it

- *Develop a plan for protecting computers from viruses and similar programmed threats.*

The plan should specify how much responsibility and authority users and system administrators should have to take specific actions to protect their computers against viruses and similar programmed threats.

In the plan, describe how users should use the available virus-detection tools for workstations, and describe any limitations on the authority of users to download and/or install new software.

1. Definitions are adapted from Deborah Russell and G. T. Gangemi, Sr., *Computer Security Basics*, O'Reilly & Associates, Inc., 1991.

➤ *Install appropriate virus-prevention tools.*

Note that copies of virus-detection and eradication tools are usually kept offline; otherwise it is possible that the virus could modify the detection tools to prevent its own detection. You should actively check for viruses online, but periodically you should also use the offline, trusted copies of the tools to scan your systems.

This step is most effective if you perform it as part of the initial installation and configuration of the operating system.

➤ *Train users in virus prevention and recognition techniques.²*

Train users to understand how viruses and other programmed threats propagate and what they can do to help prevent further propagation. This includes training them to use virus scanning tools on software obtained from public sources (such as shareware) prior to loading and executing it.

Many viruses manifest themselves in predictable ways. Train users to recognize virus symptoms, report them, and run appropriate virus eradication tools (if your plan permits them to use these tools).

Keep users apprised of new programmed threats and related intrusion scenarios.

➤ *Update the tools as needed, especially when new viruses are discovered.*

Many virus-protection tools use a database of known virus characteristics. Vendors frequently release updated versions of those databases on a weekly or monthly basis. Ensure that your computers have the most recent versions. Updating your virus-protection tools using vendor updates as they become available is one of the primary methods to prevent virus infections.

Policy considerations

Your organization's workstation acceptable use policy or security policy for networked systems should

- define users' authority (or lack thereof) to download and/or install software on the computer
- specify who has the responsibility to scan for viruses and eradicate them — users or system administrators
- prohibit users from running executable files that they have received as email attachments or downloaded from untrusted sites

Where to find updates

The latest version of this practice, plus implementation details for selected technologies, is available on the Web at URL

<http://www.cert.org/security-improvement/practices/p033.html>

2. Note that this step is primarily applicable to workstations rather than network servers. For servers, the administrator is responsible for virus prevention and recognition.

9

Configure computers for secure remote administration.

Administration of a workstation or network server includes updating user account information, examining the logs, installing new or updated software, and maintaining an appropriate configuration. These tasks usually can be performed locally from the workstation or server console or remotely from a separate host via a network connection. Although the former approach is more secure and we recommended it whenever feasible, the latter is more common.

Why this is important

Remote administration of computers is increasingly common because of the significant cost benefits—many tasks can be automated and the administrator does not have to physically visit each computer. However, remote administration tools must be configured to operate securely.

Although the normal operational state of your computer may be secure, during the performance of administrative tasks, your computer may be in a transient vulnerable state. This is especially true for remote administration of public servers that have been placed outside your firewall, because this requires that you open a network connection through the firewall. Such a connection may be vulnerable to some forms of attack, and it may open the door to anyone on the Internet being able to “administer” your server. The result could be the loss of confidentiality or integrity of information resources on the server, an intruder gaining access to resources on your internal network, or an intruder being able to use your server or workstation as an intermediate host for attacks on other internal or external hosts.

How to do it

- *Ensure that the computer accepts administration commands only from an authenticated administrator.*

Configure the computer to use a strong method to authenticate the identity of the user who is initiating the administrative processes. In particular, avoid authentication methods that require the transmission of a password in clear text, unless it is a one-time password.

- *Ensure that the computer will allow administration from only one particular host.*

Authenticate the host in a manner that does not depend on network-resolved information such as IP addresses or DNS names, because intruders can falsify such information.

- *Ensure that all administration tasks operate at the minimum necessary privilege level.*

Administration tasks sometimes require increased privilege levels. Take care to raise privilege levels only as needed.

- *Ensure that confidential information, including system configuration information, cannot be intercepted and read by intruders.*

Methods such as encryption help to ensure that network packets travelling between the administrator's host machine and the computer being administered would not, if intercepted, provide an intruder with information that would allow subsequent access to either the computer or your organization's internal network.

- *Use a movable storage medium to transfer information from the authoritative copy to public servers outside your firewall. [This step applies only to Securing Network Servers.]*

For some network servers, particularly those providing public services like the World Wide Web, it is common to develop the information content of those services on a different host machine. The authoritative version of that content is maintained (and backed up) on that other machine, and then transferred to the public server at appropriate intervals. The transfer can be performed most securely by using a movable storage medium. This could include a writable CD-ROM, diskette, hard disk cartridge, or tape. Since this procedure does not require a network connection through your firewall, it is more secure.

During the transfer, you may need to stop or disable your server. Some servers can be configured to continue operating and to send a "Service temporarily unavailable" message in response to all requests.

Do not use a transfer method that mounts a file system from a host inside the firewall on the Web server host using NFS. There are inherent problems in the NFS protocol that could make that internal host vulnerable to attack.

- *If you choose to inspect the computer log files from a host other than the computer, use a secure method of transferring the logs to that host.*

Movable storage media and file encryption are two suitable methods for transferring logs.

- *After making any changes in a computer's configuration or in its information content, create new cryptographic checksums or other integrity-checking baseline information for your server.*

See the modules *Detecting Signs of Intrusion* [Firth 97a] and *Preparing to Detect Signs of Intrusion* [Kochmar 98] for additional information on the role of checking the integrity of baseline information in support of intrusion detection.

Policy considerations

Your organization's security policy for networked systems should require the use of secure procedures for administration of network servers and workstations.

Where to find updates

The latest version of this practice, plus implementation details for selected technologies, is available on the Web at URL

<http://www.cert.org/security-improvement/practices/p062.html>

10

Allow only appropriate physical access to computers.

In addition to the steps you take to prevent inappropriate electronic access to a computer, you should also strive to allow only appropriate physical access. What this means can vary depending on the locations of computers—whether they are in locked offices or in open-plan space, for example.

Physical access also includes activities such as installing or removing hardware.

Why this is important

If unauthorized persons can physically access a computer, the integrity of that system is at considerable risk. If a system is connected to internal networks, then intruders can access resources in a way that bypasses all of your network perimeter defenses.

To preserve the confidentiality and availability of data, you must prevent the computer and its storage media from being removed from the facility by unauthorized persons.

If new hardware can be installed, such as a modem, it may make new electronic access paths to the computer and your network available to intruders.

How to do it

- *Prevent installation of unauthorized hardware and modification of authorized hardware.*

Installation of new hardware can lead to security problems in several ways:

- Installing a modem allows a direct connection from the computer to the public telephone network, which may then permit electronic access into your network from anywhere in the world, bypassing your perimeter defenses.
- Installing a removable-media storage device or printer makes it easy to copy information and carry it away from your site.
- Installing a boot device that precedes the authorized device in the boot sequence allows the computer to be restarted in a configuration that bypasses your security precautions.

You should lock the computer case, if possible. This may require third-party locking devices.

You may also want to remove or disable the external connectors on the computer.

- *Deploy the computer in a secure facility.*

Deploying the computer in a secure facility helps to prevent unauthorized access to the computer, theft, and destruction.

As a general rule, do not deploy network servers in an individual's office.

- *Secure the network wiring and other network connection components.*

For security purposes, ensure that the network cabling is not placed in a physical location where it can be easily accessed. Note that this requires you to trade-off the convenience of access for network maintenance with security.

Policy considerations

Your organization's security policy for networked systems should

- specify who is or is not allowed to install new hardware or modify existing hardware in a computer
- specify the circumstances under which users may or may not use storage devices with removable media
- specify the circumstances under which users may take storage media or printed information away from your site
- require that network servers be deployed in physically secure locations

Where to find updates

The latest version of this practice, plus implementation details for selected technologies, is available on the Web at URL

<http://www.cert.org/security-improvement/practices/p037.html>