Site Security Handbook

Status of this Memo

Abstract

   This handbook is a guide to developing computer security policies and
   procedures for sites that have systems on the Internet. The purpose
   of this handbook is to provide practical guidance to administrators
   trying to secure their information and services. The subjects
   covered include policy content and formation, a broad range of
   technical system and network security topics, and security incident
   response.

Table of Contents

1. Introduction

   This document provides guidance to system and network administrators
   on how to address security issues within the Internet community. It
   builds on the foundation provided in RFC 1244 and is the collective
   work of a number of contributing authors. Those authors include:
   Jules P. Aronson (aronson@nlm.nih.gov), Nevil Brownlee
   (n.brownlee@auckland.ac.nz), Frank Byrum (byrum@norfolk.infi.net),
   Joao Nuno Ferreira (ferreira@rccn.net), Barbara Fraser
   (byf@cert.org), Steve Glass (glass@ftp.com), Erik Guttman
   (erik.guttman@eng.sun.com), Tom Killalea (tomk@nwnet.net), Klaus-
   Peter Kossakowski (kossakowski@cert.dfn.de), Lorna Leone
   (lorna@staff.singnet.com.sg), Edward.P.Lewis
   (Edward.P.Lewis.1@gsfc.nasa.gov), Gary Malkin (gmalkin@xylogics.com),
   Russ Mundy (mundy@tis.com), Philip J. Nesser
   (pjnesser@martigny.ai.mit.edu), and Michael S. Ramsey
   (msr@interpath.net).

   In addition to the principle writers, a number of reviewers provided
   valuable comments. Those reviewers include: Eric Luiijf
   (luiijf@fel.tno.nl), Marijke Kaat (marijke.kaat@sec.nl), Ray Plzak
   (plzak@nic.mil) and Han Pronk (h.m.pronk@vka.nl).

   A special thank you goes to Joyce Reynolds, ISI, and Paul Holbrook,
   CICnet, for their vision, leadership, and effort in the creation of
   the first version of this handbook. It is the working group's sincere
   hope that this version will be as helpful to the community as the
   earlier one was.

1.1 Purpose of This Work

   This handbook is a guide to setting computer security policies and
   procedures for sites that have systems on the Internet (however, the
   information provided should also be useful to sites not yet connected
   to the Internet). This guide lists issues and factors that a site
   must consider when setting their own policies. It makes a number of
   recommendations and provides discussions of relevant areas.

   This guide is only a framework for setting security policies and
   procedures. In order to have an effective set of policies and
   procedures, a site will have to make many decisions, gain agreement,
   and then communicate and implement these policies.

1.2 Audience

   The audience for this document are system and network administrators,
   and decision makers (typically "middle management") at sites. For
   brevity, we will use the term "administrator" throughout this
   document to refer to system and network administrators.

   This document is not directed at programmers or those trying to
   create secure programs or systems. The focus of this document is on
   the policies and procedures that need to be in place to support the
   technical security features that a site may be implementing.

   The primary audience for this work are sites that are members of the
   Internet community. However, this document should be useful to any
   site that allows communication with other sites. As a general guide
   to security policies, this document may also be useful to sites with
   isolated systems.

1.3 Definitions

   For the purposes of this guide, a "site" is any organization that
   owns computers or network-related resources. These resources may
   include host computers that users use, routers, terminal servers, PCs
   or other devices that have access to the Internet. A site may be an
   end user of Internet services or a service provider such as a mid-
   level network. However, most of the focus of this guide is on those
   end users of Internet services. We assume that the site has the
   ability to set policies and procedures for itself with the
   concurrence and support from those who actually own the resources. It
   will be assumed that sites that are parts of larger organizations
   will know when they need to consult, collaborate, or take
   recommendations from, the larger entity.

The "Internet" is a collection of thousands of networks linked by a
common set of technical protocols which make it possible for users of
any one of the networks to communicate with, or use the services
located on, any of the other networks (FYI4, RFC 1594).

The term "administrator" is used to cover all those people who are
responsible for the day-to-day operation of system and network
resources. This may be a number of individuals or an organization.

The term "security administrator" is used to cover all those people
who are responsible for the security of information and information
technology. At some sites this function may be combined with
administrator (above); at others, this will be a separate position.

The term "decision maker" refers to those people at a site who set or
approve policy. These are often (but not always) the people who own
the resources.

1.4 Related Work

The Site Security Handbook Working Group is working on a User's Guide
to Internet Security. It will provide practical guidance to end users
to help them protect their information and the resources they use.

1.5 Basic Approach

This guide is written to provide basic guidance in developing a
security plan for your site. One generally accepted approach to
follow is suggested by Fites, et. al. [Fites 1989] and includes the
following steps:

(1) Identify what you are trying to protect.
(2) Determine what you are trying to protect it from.
(3) Determine how likely the threats are.
(4) Implement measures which will protect your assets in a cost-
     effective manner.
(5) Review the process continuously and make improvements each time
     a weakness is found.

Most of this document is focused on item 4 above, but the other steps
cannot be avoided if an effective plan is to be established at your
site. One old truism in security is that the cost of protecting
yourself against a threat should be less than the cost of recovering
if the threat were to strike you. Cost in this context should be
remembered to include losses expressed in real currency, reputation,
trustworthiness, and other less obvious measures. Without reasonable
knowledge of what you are protecting and what the likely threats are,
following this rule could be difficult.

1.6 Risk Assessment

1.6.1 General Discussion

   One of the most important reasons for creating a computer security
   policy is to ensure that efforts spent on security yield cost
   effective benefits. Although this may seem obvious, it is possible
   to be mislead about where the effort is needed. As an example, there
   is a great deal of publicity about intruders on computers systems;
   yet most surveys of computer security show that, for most
   organizations, the actual loss from "insiders" is much greater.

   Risk analysis involves determining what you need to protect, what you
   need to protect it from, and how to protect it. It is the process of
   examining all of your risks, then ranking those risks by level of
   severity. This process involves making cost-effective decisions on
   what you want to protect. As mentioned above, you should probably
   not spend more to protect something than it is actually worth.

   A full treatment of risk analysis is outside the scope of this
   document. [Fites 1989] and [Pfleeger 1989] provide introductions to
   this topic. However, there are two elements of a risk analysis that
   will be briefly covered in the next two sections:

   (1) Identifying the assets
   (2) Identifying the threats

   For each asset, the basic goals of security are availability,
   confidentiality, and integrity. Each threat should be examined with
   an eye to how the threat could affect these areas.

1.6.2 Identifying the Assets

   One step in a risk analysis is to identify all the things that need
   to be protected. Some things are obvious, like valuable proprietary
   information, intellectual property, and all the various pieces of
   hardware; but, some are overlooked, such as the people who actually
   use the systems. The essential point is to list all things that could
   be affected by a security problem.

   One list of categories is suggested by Pfleeger [Pfleeger 1989]; this
   list is adapted from that source:

   (1) Hardware: CPUs, boards, keyboards, terminals,
        workstations, personal computers, printers, disk
        drives, communication lines, terminal servers, routers.

    (2)   Software: source programs, object programs,
          utilities, diagnostic programs, operating systems,
          communication programs.

    (3)   Data: during execution, stored on-line, archived off-line,
          backups, audit logs, databases, in transit over
          communication media.

    (4)   People: users, administrators, hardware maintainers.

    (5)   Documentation: on programs, hardware, systems, local
          administrative procedures.

    (6)   Supplies: paper, forms, ribbons, magnetic media.

1.6.3 Identifying the Threats

    Once the assets requiring protection are identified, it is necessary
    to identify threats to those assets. The threats can then be
    examined to determine what potential for loss exists. It helps to
    consider from what threats you are trying to protect your assets.
    The following are classic threats that should be considered.
    Depending on your site, there will be more specific threats that
    should be identified and addressed.

    (1)   Unauthorized access to resources and/or information
    (2)   Unintented and/or unauthorized Disclosure of information
    (3)   Denial of service

2. Security Policies

    Throughout this document there will be many references to policies.
    Often these references will include recommendations for specific
    policies. Rather than repeat guidance in how to create and
    communicate such a policy, the reader should apply the advice
    presented in this chapter when developing any policy recommended
    later in this book.

2.1 What is a Security Policy and Why Have One?

    The security-related decisions you make, or fail to make, as
    administrator largely determines how secure or insecure your network
    is, how much functionality your network offers, and how easy your
    network is to use. However, you cannot make good decisions about
    security without first determining what your security goals are.
    Until you determine what your security goals are, you cannot make
    effective use of any collection of security tools because you simply
    will not know what to check for and what restrictions to impose.

   For example, your goals will probably be very different from the
   goals of a product vendor. Vendors are trying to make configuration
   and operation of their products as simple as possible, which implies
   that the default configurations will often be as open (i.e.,
   insecure) as possible. While this does make it easier to install new
   products, it also leaves access to those systems, and other systems
   through them, open to any user who wanders by.

   Your goals will be largely determined by the following key tradeoffs:

   (1)   services offered versus security provided -
         Each service offered to users carries its own security risks.
         For some services the risk outweighs the benefit of the service
         and the administrator may choose to eliminate the service rather
         than try to secure it.

   (2)   ease of use versus security -
         The easiest system to use would allow access to any user and
         require no passwords; that is, there would be no security.
         Requiring passwords makes the system a little less convenient,
         but more secure. Requiring device-generated one-time passwords
         makes the system even more difficult to use, but much more
         secure.

   (3)   cost of security versus risk of loss -
         There are many different costs to security: monetary (i.e., the
         cost of purchasing security hardware and software like firewalls
         and one-time password generators), performance (i.e., encryption
         and decryption take time), and ease of use (as mentioned above).
         There are also many levels of risk: loss of privacy (i.e., the
         reading of information by unauthorized individuals), loss of
         data (i.e., the corruption or erasure of information), and the
         loss of service (e.g., the filling of data storage space, usage
         of computational resources, and denial of network access). Each
         type of cost must be weighed against each type of loss.

   Your goals should be communicated to all users, operations staff, and
   managers through a set of security rules, called a "security policy."
   We are using this term, rather than the narrower "computer security
   policy" since the scope includes all types of information technology
   and the information stored and manipulated by the technology.

2.1.1 Definition of a Security Policy

   A security policy is a formal statement of the rules by which people
   who are given access to an organization's technology and information
   assets must abide.

2.1.2 Purposes of a Security Policy

   The main purpose of a security policy is to inform users, staff and
   managers of their obligatory requirements for protecting technology
   and information assets. The policy should specify the mechanisms
   through which these requirements can be met. Another purpose is to
   provide a baseline from which to acquire, configure and audit
   computer systems and networks for compliance with the policy.
   Therefore an attempt to use a set of security tools in the absence of
   at least an implied security policy is meaningless.

   An Appropriate Use Policy (AUP) may also be part of a security
   policy. It should spell out what users shall and shall not do on the
   various components of the system, including the type of traffic
   allowed on the networks. The AUP should be as explicit as possible
   to avoid ambiguity or misunderstanding. For example, an AUP might
   list any prohibited USENET newsgroups. (Note: Appropriate Use Policy
   is referred to as Acceptable Use Policy by some sites.)

2.1.3 Who Should be Involved When Forming Policy?

   In order for a security policy to be appropriate and effective, it
   needs to have the acceptance and support of all levels of employees
   within the organization. It is especially important that corporate
   management fully support the security policy process otherwise there
   is little chance that they will have the intended impact. The
   following is a list of individuals who should be involved in the
   creation and review of security policy documents:

   (1) site security administrator
   (2) information technology technical staff (e.g., staff from
        computing center)
   (3) administrators of large user groups within the organization
        (e.g., business divisions, computer science department within a
        university, etc.)
   (4) security incident response team
   (5) representatives of the user groups affected by the security
        policy
   (6) responsible management
   (7) legal counsel (if appropriate)

   The list above is representative of many organizations, but is not
   necessarily comprehensive. The idea is to bring in representation
   from key stakeholders, management who have budget and policy
   authority, technical staff who know what can and cannot be supported,
   and legal counsel who know the legal ramifications of various policy

choices. In some organizations, it may be appropriate to include EDP
audit personnel. Involving this group is important if resulting
policy statements are to reach the broadest possible acceptance. It
is also relevant to mention that the role of legal counsel will also
vary from country to country.

2.2 What Makes a Good Security Policy?

The characteristics of a good security policy are:

(1) It must be implementable through system administration
     procedures, publishing of acceptable use guidelines, or other
     appropriate methods.

(2) It must be enforcible with security tools, where appropriate,
     and with sanctions, where actual prevention is not technically
     feasible.

(3) It must clearly define the areas of responsibility for the
     users, administrators, and management.

The components of a good security policy include:

(1) Computer Technology Purchasing Guidelines which specify
     required, or preferred, security features. These should
     supplement existing purchasing policies and guidelines.

(2) A Privacy Policy which defines reasonable expectations of
     privacy regarding such issues as monitoring of electronic mail,
     logging of keystrokes, and access to users' files.

(3) An Access Policy which defines access rights and privileges to
     protect assets from loss or disclosure by specifying acceptable
     use guidelines for users, operations staff, and management. It
     should provide guidelines for external connections, data
     communications, connecting devices to a network, and adding new
     software to systems. It should also specify any required
     notification messages (e.g., connect messages should provide
     warnings about authorized usage and line monitoring, and not
     simply say "Welcome").

(4) An Accountability Policy which defines the responsibilities of
     users, operations staff, and management. It should specify an
     audit capability, and provide incident handling guidelines
     (i.e., what to do and who to contact if a possible intrusion is
     detected).

   (5) An Authentication Policy which establishes trust through an
        effective password policy, and by setting guidelines for remote
        location authentication and the use of authentication devices
        (e.g., one-time passwords and the devices that generate them).

   (6) An Availability statement which sets users' expectations for the
        availability of resources. It should address redundancy and
        recovery issues, as well as specify operating hours and
        maintenance down-time periods. It should also include contact
        information for reporting system and network failures.

   (7) An Information Technology System & Network Maintenance Policy
        which describes how both internal and external maintenance
        people are allowed to handle and access technology. One
        important topic to be addressed here is whether remote
        maintenance is allowed and how such access is controlled.
        Another area for consideration here is outsourcing and how it is
        managed.

   (8) A Violations Reporting Policy that indicates which types of
        violations (e.g., privacy and security, internal and external)
        must be reported and to whom the reports are made. A non-
        threatening atmosphere and the possibility of anonymous
        reporting will result in a greater probability that a violation
        will be reported if it is detected.

   (9) Supporting Information which provides users, staff, and
        management with contact information for each type of policy
        violation; guidelines on how to handle outside queries about a
        security incident, or information which may be considered
        confidential or proprietary; and cross-references to security
        procedures and related information, such as company policies and
        governmental laws and regulations.

   There may be regulatory requirements that affect some aspects of your
   security policy (e.g., line monitoring). The creators of the
   security policy should consider seeking legal assistance in the
   creation of the policy. At a minimum, the policy should be reviewed
   by legal counsel.

   Once your security policy has been established it should be clearly
   communicated to users, staff, and management. Having all personnel
   sign a statement indicating that they have read, understood, and
   agreed to abide by the policy is an important part of the process.
   Finally, your policy should be reviewed on a regular basis to see if
   it is successfully supporting your security needs.

2.3 Keeping the Policy Flexible

   In order for a security policy to be viable for the long term, it
   requires a lot of flexibility based upon an architectural security
   concept. A security policy should be (largely) independent from
   specific hardware and software situations (as specific systems tend
   to be replaced or moved overnight). The mechanisms for updating the
   policy should be clearly spelled out. This includes the process, the
   people involved, and the people who must sign-off on the changes.

   It is also important to recognize that there are exceptions to every
   rule. Whenever possible, the policy should spell out what exceptions
   to the general policy exist. For example, under what conditions is a
   system administrator allowed to go through a user's files. Also,
   there may be some cases when multiple users will have access to the
   same userid. For example, on systems with a "root" user, multiple
   system administrators may know the password and use the root account.

   Another consideration is called the "Garbage Truck Syndrome." This
   refers to what would happen to a site if a key person was suddenly
   unavailable for his/her job function (e.g., was suddenly ill or left
   the company unexpectedly). While the greatest security resides in
   the minimum dissemination of information, the risk of losing critical
   information increases when that information is not shared. It is
   important to determine what the proper balance is for your site.

3. Architecture

3.1 Objectives

3.1.1 Completely Defined Security Plans

   All sites should define a comprehensive security plan. This plan
   should be at a higher level than the specific policies discussed in
   chapter 2, and it should be crafted as a framework of broad
   guidelines into which specific policies will fit.

   It is important to have this framework in place so that individual
   policies can be consistent with the overall site security
   architecture. For example, having a strong policy with regard to
   Internet access and having weak restrictions on modem usage is
   inconsistent with an overall philosophy of strong security
   restrictions on external access.

   A security plan should define: the list of network services that will
   be provided; which areas of the organization will provide the
   services; who will have access to those services; how access will be
   provided; who will administer those services; etc.

The plan should also address how incident will be handled. Chapter 5
provides an in-depth discussion of this topic, but it is important
for each site to define classes of incidents and corresponding
responses. For example, sites with firewalls should set a threshold
on the number of attempts made to foil the firewall before triggering
a response? Escallation levels should be defined for both attacks
and responses. Sites without firewalls will have to determine if a
single attempt to connect to a host constitutes an incident? What
about a systematic scan of systems?

For sites connected to the Internet, the rampant media magnification
of Internet related security incidents can overshadow a (potentially)
more serious internal security problem. Likewise, companies who have
never been connected to the Internet may have strong, well defined,
internal policies but fail to adequately address an external
connection policy.

3.1.2 Separation of Services

There are many services which a site may wish to provide for its
users, some of which may be external. There are a variety of
security reasons to attempt to isolate services onto dedicated host
computers. There are also performance reasons in most cases, but a
detailed discussion is beyond to scope of this document.

The services which a site may provide will, in most cases, have
different levels of access needs and models of trust. Services which
are essential to the security or smooth operation of a site would be
better off being placed on a dedicated machine with very limited
access (see Section 3.1.3 "deny all" model), rather than on a machine
that provides a service (or services) which has traditionally been
less secure, or requires greater accessability by users who may
accidentally suborn security.

It is also important to distinguish between hosts which operate
within different models of trust (e.g., all the hosts inside of a
firewall and any host on an exposed network).

Some of the services which should be examined for potential
separation are outlined in section 3.2.3. It is important to remember
that security is only as strong as the weakest link in the chain.
Several of the most publicized penetrations in recent years have been
through the exploitation of vulnerabilities in electronic mail
systems. The intruders were not trying to steal electronic mail, but
they used the vulnerability in that service to gain access to other
systems.

If possible, each service should be running on a different machine
whose only duty is to provide a specific service. This helps to
isolate intruders and limit potential harm.

3.1.3 Deny all/ Allow all

There are two diametrically opposed underlying philosophies which can
be adopted when defining a security plan. Both alternatives are
legitimate models to adopt, and the choice between them will depend
on the site and its needs for security.

The first option is to turn off all services and then selectively
enable services on a case by case basis as they are needed. This can
be done at the host or network level as appropriate. This model,
which will here after be referred to as the "deny all" model, is
generally more secure than the other model described in the next
paragraph. More work is required to successfully implement a "deny
all" configuration as well as a better understanding of services.
Allowing only known services provides for a better analysis of a
particular service/protocol and the design of a security mechanism
suited to the security level of the site.

The other model, which will here after be referred to as the "allow
all" model, is much easier to implement, but is generally less secure
than the "deny all" model. Simply turn on all services, usually the
default at the host level, and allow all protocols to travel across
network boundaries, usually the default at the router level. As
security holes become apparent, they are restricted or patched at
either the host or network level.

Each of these models can be applied to different portions of the
site, depending on functionality requirements, administrative
control, site policy, etc. For example, the policy may be to use the
"allow all" model when setting up workstations for general use, but
adopt a "deny all" model when setting up information servers, like an
email hub. Likewise, an "allow all" policy may be adopted for
traffic between LAN's internal to the site, but a "deny all" policy
can be adopted between the site and the Internet.

Be careful when mixing philosophies as in the examples above. Many
sites adopt the theory of a hard "crunchy" shell and a soft "squishy"
middle. They are willing to pay the cost of security for their
external traffic and require strong security measures, but are
unwilling or unable to provide similar protections internally. This
works fine as long as the outer defenses are never breached and the
internal users can be trusted. Once the outer shell (firewall) is
breached, subverting the internal network is trivial.

3.1.4 Identify Real Needs for Services

   There is a large variety of services which may be provided, both
   internally and on the Internet at large. Managing security is, in
   many ways, managing access to services internal to the site and
   managing how internal users access information at remote sites.

   Services tend to rush like waves over the Internet. Over the years
   many sites have established anonymous FTP servers, gopher servers,
   wais servers, WWW servers, etc. as they became popular, but not
   particularly needed, at all sites. Evaluate all new services that
   are established with a skeptical attitude to determine if they are
   actually needed or just the current fad sweeping the Internet.

   Bear in mind that security complexity can grow exponentially with the
   number of services provided. Filtering routers need to be modified
   to support the new protocols. Some protocols are inherently
   difficult to filter safely (e.g., RPC and UDP services), thus
   providing more openings to the internal network. Services provided
   on the same machine can interact in catastrophic ways. For example,
   allowing anonymous FTP on the same machine as the WWW server may
   allow an intruder to place a file in the anonymous FTP area and cause
   the HTTP server to execute it.

3.2 Network and Service Configuration

3.2.1 Protecting the Infrastructure

   Many network administrators go to great lengths to protect the hosts
   on their networks. Few administrators make any effort to protect the
   networks themselves. There is some rationale to this. For example,
   it is far easier to protect a host than a network. Also, intruders
   are likely to be after data on the hosts; damaging the network would
   not serve their purposes. That said, there are still reasons to
   protect the networks. For example, an intruder might divert network
   traffic through an outside host in order to examine the data (i.e.,
   to search for passwords). Also, infrastructure includes more than
   the networks and the routers which interconnect them. Infrastructure
   also includes network management (e.g., SNMP), services (e.g., DNS,
   NFS, NTP, WWW), and security (i.e., user authentication and access
   restrictions).

   The infrastructure also needs protection against human error. When
   an administrator misconfigures a host, that host may offer degraded
   service. This only affects users who require that host and, unless

that host is a primary server, the number of affected users will
therefore be limited. However, if a router is misconfigured, all
users who require the network will be affected. Obviously, this is a
far larger number of users than those depending on any one host.

3.2.2 Protecting the Network

There are several problems to which networks are vulnerable. The
classic problem is a "denial of service" attack. In this case, the
network is brought to a state in which it can no longer carry
legitimate users' data. There are two common ways this can be done:
by attacking the routers and by flooding the network with extraneous
traffic. Please note that the term "router" in this section is used
as an example of a larger class of active network interconnection
components that also includes components like firewalls, proxy-
servers, etc.

An attack on the router is designed to cause it to stop forwarding
packets, or to forward them improperly. The former case may be due
to a misconfiguration, the injection of a spurious routing update, or
a "flood attack" (i.e., the router is bombarded with unroutable
packets, causing its performance to degrade). A flood attack on a
network is similar to a flood attack on a router, except that the
flood packets are usually broadcast. An ideal flood attack would be
the injection of a single packet which exploits some known flaw in
the network nodes and causes them to retransmit the packet, or
generate error packets, each of which is picked up and repeated by
another host. A well chosen attack packet can even generate an
exponential explosion of transmissions.

Another classic problem is "spoofing." In this case, spurious
routing updates are sent to one or more routers causing them to
misroute packets. This differs from a denial of service attack only
in the purpose behind the spurious route. In denial of service, the
object is to make the router unusable; a state which will be quickly
detected by network users. In spoofing, the spurious route will
cause packets to be routed to a host from which an intruder may
monitor the data in the packets. These packets are then re-routed to
their correct destinations. However, the intruder may or may not
have altered the contents of the packets.

The solution to most of these problems is to protect the routing
update packets sent by the routing protocols in use (e.g., RIP-2,
OSPF). There are three levels of protection: clear-text password,
cryptographic checksum, and encryption. Passwords offer only minimal
protection against intruders who do not have direct access to the
physical networks. Passwords also offer some protection against
misconfigured routers (i.e, routers which, out of the box, attempt to

route packets). The advantage of passwords is that they have a very
low overhead, in both bandwidth and CPU consumption. Checksums
protect against the injection of spurious packets, even if the
intruder has direct access to the physical network. Combined with a
sequence number, or other unique identifier, a checksum can also
protect again "replay" attacks, wherein an old (but valid at the
time) routing update is retransmitted by either an intruder or a
misbehaving router. The most security is provided by complete
encryption of sequenced, or uniquely identified, routing updates.
This prevents an intruder from determining the topology of the
network. The disadvantage to encryption is the overhead involved in
processing the updates.

RIP-2 (RFC 1723) and OSPF (RFC 1583) both support clear-text
passwords in their base design specifications. In addition, there
are extensions to each base protocol to support MD5 encryption.

Unfortunately, there is no adequate protection against a flooding
attack, or a misbehaving host or router which is flooding the
network. Fortunately, this type of attack is obvious when it occurs
and can usually be terminated relatively simply.

3.2.3 Protecting the Services

There are many types of services and each has its own security
requirements. These requirements will vary based on the intended use
of the service. For example, a service which should only be usable
within a site (e.g., NFS) may require different protection mechanisms
than a service provided for external use. It may be sufficient to
protect the internal server from external access. However, a WWW
server, which provides a home page intended for viewing by users
anywhere on the Internet, requires built-in protection. That is, the
service/protocol/server must provide whatever security may be
required to prevent unauthorized access and modification of the Web
database.

Internal services (i.e., services meant to be used only by users
within a site) and external services (i.e., services deliberately
made available to users outside a site) will, in general, have
protection requirements which differ as previously described. It is
therefore wise to isolate the internal services to one set of server
host computers and the external services to another set of server
host computers. That is, internal and external servers should not be
co-located on the same host computer. In fact, many sites go so far

as to have one set of subnets (or even different networks) which are
accessible from the outside and another set which may be accessed
only within the site. Of course, there is usually a firewall which
connects these partitions. Great care must be taken to ensure that
such a firewall is operating properly.

There is increasing interest in using intranets to connect different
parts of a organization (e.g., divisions of a company). While this
document generally differentiates between external and internal
(public and private), sites using intranets should be aware that they
will need to consider three separations and take appropriate actions
when designing and offering services. A service offered to an
intranet would be neither public, nor as completely private as a
service to a single organizational subunit. Therefore, the service
would need its own supporting system, separated from both external
and internal services and networks.

One form of external service deserves some special consideration, and
that is anonymous, or guest, access. This may be either anonymous
FTP or guest (unauthenticated) login. It is extremely important to
ensure that anonymous FTP servers and guest login userids are
carefully isolated from any hosts and file systems from which outside
users should be kept. Another area to which special attention must
be paid concerns anonymous, writable access. A site may be legally
responsible for the content of publicly available information, so
careful monitoring of the information deposited by anonymous users is
advised.

Now we shall consider some of the most popular services: name
service, password/key service, authentication/proxy service,
electronic mail, WWW, file transfer, and NFS. Since these are the
most frequently used services, they are the most obvious points of
attack. Also, a successful attack on one of these services can
produce disaster all out of proportion to the innocence of the basic
service.

3.2.3.1 Name Servers (DNS and NIS(+))

The Internet uses the Domain Name System (DNS) to perform address
resolution for host and network names. The Network Information
Service (NIS) and NIS+ are not used on the global Internet, but are
subject to the same risks as a DNS server. Name-to-address
resolution is critical to the secure operation of any network. An
attacker who can successfully control or impersonate a DNS server can
re-route traffic to subvert security protections. For example,
routine traffic can be diverted to a compromised system to be
monitored; or, users can be tricked into providing authentication
secrets. An organization should create well known, protected sites

to act as secondary name servers and protect their DNS masters from
denial of service attacks using filtering routers.

Traditionally, DNS has had no security capabilities. In particular,
the information returned from a query could not be checked for
modification or verified that it had come from the name server in
question. Work has been done to incorporate digital signatures into
the protocol which, when deployed, will allow the integrity of the
information to be cryptographically verified (see RFC 2065).

3.2.3.2 Password/Key Servers (NIS(+) and KDC)

Password and key servers generally protect their vital information
(i.e., the passwords and keys) with encryption algorithms. However,
even a one-way encrypted password can be determined by a dictionary
attack (wherein common words are encrypted to see if they match the
stored encryption). It is therefore necessary to ensure that these
servers are not accessable by hosts which do not plan to use them for
the service, and even those hosts should only be able to access the
service (i.e., general services, such as Telnet and FTP, should not
be allowed by anyone other than administrators).

3.2.3.3 Authentication/Proxy Servers (SOCKS, FWTK)

A proxy server provides a number of security enhancements. It allows
sites to concentrate services through a specific host to allow
monitoring, hiding of internal structure, etc. This funnelling of
services creates an attractive target for a potential intruder. The
type of protection required for a proxy server depends greatly on the
proxy protocol in use and the services being proxied. The general
rule of limiting access only to those hosts which need the services,
and limiting access by those hosts to only those services, is a good
starting point.

3.2.3.4 Electronic Mail

Electronic mail (email) systems have long been a source for intruder
break-ins because email protocols are among the oldest and most
widely deployed services. Also, by it's very nature, an email server
requires access to the outside world; most email servers accept input
from any source. An email server generally consists of two parts: a
receiving/sending agent and a processing agent. Since email is
delivered to all users, and is usually private, the processing agent
typically requires system (root) privileges to deliver the mail.
Most email implementations perform both portions of the service,
which means the receiving agent also has system privileges. This
opens several security holes which this document will not describe.
There are some implementations available which allow a separation of

the two agents. Such implementations are generally considered more
secure, but still require careful installation to avoid creating a
security problem.

3.2.3.5 World Wide Web (WWW)

The Web is growing in popularity exponentially because of its ease of
use and the powerful ability to concentrate information services.
Most WWW servers accept some type of direction and action from the
persons accessing their services. The most common example is taking
a request from a remote user and passing the provided information to
a program running on the server to process the request. Some of
these programs are not written with security in mind and can create
security holes. If a Web server is available to the Internet
community, it is especially important that confidential information
not be co-located on the same host as that server. In fact, it is
recommended that the server have a dedicated host which is not
"trusted" by other internal hosts.

Many sites may want to co-locate FTP service with their WWW service.
But this should only occur for anon-ftp servers that only provide
information (ftp-get). Anon-ftp puts, in combination with WWW, might
be dangerous (e.g., they could result in modifications to the
information your site is publishing to the web) and in themselves
make the security considerations for each service different.

3.2.3.6 File Transfer (FTP, TFTP)

FTP and TFTP both allow users to receive and send electronic files in
a point-to-point manner. However, FTP requires authentication while
TFTP requires none. For this reason, TFTP should be avoided as much
as possible.

Improperly configured FTP servers can allow intruders to copy,
replace and delete files at will, anywhere on a host, so it is very
important to configure this service correctly. Access to encrypted
passwords and proprietary data, and the introduction of Trojan horses
are just a few of the potential security holes that can occur when
the service is configured incorrectly. FTP servers should reside on
their own host. Some sites choose to co-locate FTP with a Web
server, since the two protocols share common security considerations
However, the the practice isn't recommended, especially when the FTP
service allows the deposit of files (see section on WWW above). As
mentioned in the opening paragraphs of section 3.2.3, services
offered internally to your site should not be co-located with
services offered externally. Each should have its own host.

TFTP does not support the same range of functions as FTP, and has no
security whatsoever. This service should only be considered for
internal use, and then it should be configured in a restricted way so
that the server only has access to a set of predetermined files
(instead of every world-readable file on the system). Probably the
most common usage of TFTP is for downloading router configuration
files to a router. TFTP should reside on its own host, and should
not be installed on hosts supporting external FTP or Web access.

3.2.3.7 NFS

The Network File Service allows hosts to share common disks. NFS is
frequently used by diskless hosts who depend on a disk server for all
of their storage needs. Unfortunately, NFS has no built-in security.
It is therefore necessary that the NFS server be accessible only by
those hosts which are using it for service. This is achieved by
specifying which hosts the file system is being exported to and in
what manner (e.g., read-only, read-write, etc.). Filesystems should
not be exported to any hosts outside the local network since this
will require that the NFS service be accessible externally. Ideally,
external access to NFS service should be stopped by a firewall.

3.2.4 Protecting the Protection

It is amazing how often a site will overlook the most obvious
weakness in its security by leaving the security server itself open
to attack. Based on considerations previously discussed, it should
be clear that: the security server should not be accessible from
off-site; should offer minimum access, except for the authentication
function, to users on-site; and should not be co-located with any
other servers. Further, all access to the node, including access to
the service itself, should be logged to provide a "paper trail" in
the event of a security breach.

3.3 Firewalls

One of the most widely deployed and publicized security measures in
use on the Internet is a "firewall." Firewalls have been given the
reputation of a general panacea for many, if not all, of the Internet
security issues. They are not. Firewalls are just another tool in
the quest for system security. They provide a certain level of
protection and are, in general, a way of implementing security policy
at the network level. The level of security that a firewall provides
can vary as much as the level of security on a particular machine.
There are the traditional trade-offs between security, ease of use,
cost, complexity, etc.

A firewall is any one of several mechanisms used to control and watch
access to and from a network for the purpose of protecting it. A
firewall acts as a gateway through which all traffic to and from the
protected network and/or systems passes. Firewalls help to place
limitations on the amount and type of communication that takes place
between the protected network and the another network (e.g., the
Internet, or another piece of the site's network).

A firewall is generally a way to build a wall between one part of a
network, a company's internal network, for example, and another part,
the global Internet, for example. The unique feature about this wall
is that there needs to be ways for some traffic with particular
characteristics to pass through carefully monitored doors
("gateways"). The difficult part is establishing the criteria by
which the packets are allowed or denied access through the doors.
Books written on firewalls use different terminology to describe the
various forms of firewalls. This can be confusing to system
administrators who are not familiar with firewalls. The thing to note
here is that there is no fixed terminology for the description of
firewalls.

Firewalls are not always, or even typically, a single machine.
Rather, firewalls are often a combination of routers, network
segments, and host computers. Therefore, for the purposes of this
discussion, the term "firewall" can consist of more than one physical
device. Firewalls are typically built using two different
components, filtering routers and proxy servers.

Filtering routers are the easiest component to conceptualize in a
firewall. A router moves data back and forth between two (or more)
different networks. A "normal" router takes a packet from network A
and "routes" it to its destination on network B. A filtering router
does the same thing but decides not only how to route the packet, but
whether it should route the packet. This is done by installing a
series of filters by which the router decides what to do with any
given packet of data.

A discussion concerning capabilities of a particular brand of router,
running a particular software version is outside the scope of this
document. However, when evaluating a router to be used for filtering
packets, the following criteria can be important when implementing a
filtering policy: source and destination IP address, source and
destination TCP port numbers, state of the TCP "ack" bit, UDP source
and destination port numbers, and direction of packet flow (i.e.. A-
>B or B->A). Other information necessary to construct a secure
filtering scheme are whether the router reorders filter instructions
(designed to optimize filters, this can sometimes change the meaning
and cause unintended access), and whether it is possible to apply

filters for inbound and outbound packets on each interface (if the
router filters only outbound packets then the router is "outside" of
its filters and may be more vulnerable to attack). In addition to
the router being vulnerable, this distinction between applying
filters on inbound or outbound packets is especially relevant for
routers with more than 2 interfaces. Other important issues are the
ability to create filters based on IP header options and the fragment
state of a packet. Building a good filter can be very difficult and
requires a good understanding of the type of services (protocols)
that will be filtered.

For better security, the filters usually restrict access between the
two connected nets to just one host, the bastion host. It is only
possible to access the other network via this bastion host. As only
this host, rather than a few hundred hosts, can get attacked, it is
easier to maintain a certain level of security because only this host
has to be protected very carefully. To make resources available to
legitimate users across this firewall, services have to be forwarded
by the bastion host. Some servers have forwarding built in (like
DNS-servers or SMTP-servers), for other services (e.g., Telnet, FTP,
etc.), proxy servers can be used to allow access to the resources
across the firewall in a secure way.

A proxy server is way to concentrate application services through a
single machine. There is typically a single machine (the bastion
host) that acts as a proxy server for a variety of protocols (Telnet,
SMTP, FTP, HTTP, etc.) but there can be individual host computers for
each service. Instead of connecting directly to an external server,
the client connects to the proxy server which in turn initiates a
connection to the requested external server. Depending on the type
of proxy server used, it is possible to configure internal clients to
perform this redirection automatically, without knowledge to the
user, others might require that the user connect directly to the
proxy server and then initiate the connection through a specified
format.

There are significant security benefits which can be derived from
using proxy servers. It is possible to add access control lists to
protocols, requiring users or systems to provide some level of
authentication before access is granted. Smarter proxy servers,
sometimes called Application Layer Gateways (ALGs), can be written
which understand specific protocols and can be configured to block
only subsections of the protocol. For example, an ALG for FTP can
tell the difference between the "put" command and the "get" command;
an organization may wish to allow users to "get" files from the
Internet, but not be able to "put" internal files on a remote server.
By contrast, a filtering router could either block all FTP access, or
none, but not a subset.

Proxy servers can also be configured to encrypt data streams based on
a variety of parameters. An organization might use this feature to
allow encrypted connections between two locations whose sole access
points are on the Internet.

Firewalls are typically thought of as a way to keep intruders out,
but they are also often used as a way to let legitimate users into a
site. There are many examples where a valid user might need to
regularly access the "home" site while on travel to trade shows and
conferences, etc. Access to the Internet is often available but may
be through an untrusted machine or network. A correctly configured
proxy server can allow the correct users into the site while still
denying access to other users.

The current best effort in firewall techniques is found using a
combination of a pair of screening routers with one or more proxy
servers on a network between the two routers. This setup allows the
external router to block off any attempts to use the underlying IP
layer to break security (IP spoofing, source routing, packet
fragments), while allowing the proxy server to handle potential
security holes in the higher layer protocols. The internal router's
purpose is to block all traffic except to the proxy server. If this
setup is rigidly implemented, a high level of security can be
achieved.

Most firewalls provide logging which can be tuned to make security
administration of the network more convenient. Logging may be
centralized and the system may be configured to send out alerts for
abnormal conditions. It is important to regularly monitor these logs
for any signs of intrusions or break-in attempts. Since some
intruders will attempt to cover their tracks by editing logs, it is
desirable to protect these logs. A variety of methods is available,
including: write once, read many (WORM) drives; papers logs; and
centralized logging via the "syslog" utility. Another technique is
to use a "fake" serial printer, but have the serial port connected to
an isolated machine or PC which keeps the logs.

Firewalls are available in a wide range of quality and strengths.
Commercial packages start at approximately $10,000US and go up to
over $250,000US. "Home grown" firewalls can be built for smaller
amounts of capital. It should be remembered that the correct setup
of a firewall (commercial or homegrown) requires a significant amount
of skill and knowledge of TCP/IP. Both types require regular
maintenance, installation of software patches and updates, and
regular monitoring. When budgeting for a firewall, these additional
costs should be considered in addition to the cost of the physical
elements of the firewall.

As an aside, building a "home grown" firewall requires a significant
amount of skill and knowledge of TCP/IP. It should not be trivially
attempted because a perceived sense of security is worse in the long
run than knowing that there is no security. As with all security
measures, it is important to decide on the threat, the value of the
assets to be protected, and the costs to implement security.

A final note about firewalls. They can be a great aid when
implementing security for a site and they protect against a large
variety of attacks. But it is important to keep in mind that they
are only one part of the solution. They cannot protect your site
against all types of attack.

4. Security Services and Procedures

This chapter guides the reader through a number of topics that should
be addressed when securing a site. Each section touches on a
security service or capability that may be required to protect the
information and systems at a site. The topics are presented at a
fairly high-level to introduce the reader to the concepts.

Throughout the chapter, you will find significant mention of
cryptography. It is outside the scope of this document to delve into
details concerning cryptography, but the interested reader can obtain
more information from books and articles listed in the reference
section of this document.

4.1 Authentication

For many years, the prescribed method for authenticating users has
been through the use of standard, reusable passwords. Originally,
these passwords were used by users at terminals to authenticate
themselves to a central computer. At the time, there were no
networks (internally or externally), so the risk of disclosure of the
clear text password was minimal. Today, systems are connected
together through local networks, and these local networks are further
connected together and to the Internet. Users are logging in from
all over the globe; their reusable passwords are often transmitted
across those same networks in clear text, ripe for anyone in-between
to capture. And indeed, the CERT* Coordination Center and other
response teams are seeing a tremendous number of incidents involving
packet sniffers which are capturing the clear text passwords.

With the advent of newer technologies like one-time passwords (e.g.,
S/Key), PGP, and token-based authentication devices, people are using
password-like strings as secret tokens and pins. If these secret
tokens and pins are not properly selected and protected, the
authentication will be easily subverted.

4.1.1 One-Time passwords

   As mentioned above, given today's networked environments, it is
   recommended that sites concerned about the security and integrity of
   their systems and networks consider moving away from standard,
   reusable passwords. There have been many incidents involving Trojan
   network programs (e.g., telnet and rlogin) and network packet
   sniffing programs. These programs capture clear text
   hostname/account name/password triplets. Intruders can use the
   captured information for subsequent access to those hosts and
   accounts. This is possible because 1) the password is used over and
   over (hence the term "reusable"), and 2) the password passes across
   the network in clear text.

   Several authentication techniques have been developed that address
   this problem. Among these techniques are challenge-response
   technologies that provide passwords that are only used once (commonly
   called one-time passwords). There are a number of products available
   that sites should consider using. The decision to use a product is
   the responsibility of each organization, and each organization should
   perform its own evaluation and selection.

4.1.2 Kerberos

   Kerberos is a distributed network security system which provides for
   authentication across unsecured networks. If requested by the
   application, integrity and encryption can also be provided. Kerberos
   was originally developed at the Massachusetts Institute of Technology
   (MIT) in the mid 1980s. There are two major releases of Kerberos,
   version 4 and 5, which are for practical purposes, incompatible.

   Kerberos relies on a symmetric key database using a key distribution
   center (KDC) which is known as the Kerberos server. A user or
   service (known as "principals") are granted electronic "tickets"
   after properly communicating with the KDC. These tickets are used
   for authentication between principals. All tickets include a time
   stamp which limits the time period for which the ticket is valid.
   Therefore, Kerberos clients and server must have a secure time
   source, and be able to keep time accurately.

   The practical side of Kerberos is its integration with the
   application level. Typical applications like FTP, telnet, POP, and
   NFS have been integrated with the Kerberos system. There are a
   variety of implementations which have varying levels of integration.
   Please see the Kerberos FAQ available at http://www.ov.com/misc/krb-
   faq.html for the latest information.

4.1.3 Choosing and Protecting Secret Tokens and PINs

   When selecting secret tokens, take care to choose them carefully.
   Like the selection of passwords, they should be robust against brute
   force efforts to guess them. That is, they should not be single
   words in any language, any common, industry, or cultural acronyms,
   etc. Ideally, they will be longer rather than shorter and consist of
   pass phrases that combine upper and lower case character, digits, and
   other characters.

   Once chosen, the protection of these secret tokens is very important.
   Some are used as pins to hardware devices (like token cards) and
   these should not be written down or placed in the same location as
   the device with which they are associated. Others, such as a secret
   Pretty Good Privacy (PGP) key, should be protected from unauthorized
   access.

   One final word on this subject. When using cryptography products,
   like PGP, take care to determine the proper key length and ensure
   that your users are trained to do likewise. As technology advances,
   the minimum safe key length continues to grow. Make sure your site
   keeps up with the latest knowledge on the technology so that you can
   ensure that any cryptography in use is providing the protection you
   believe it is.

4.1.4 Password Assurance

   While the need to eliminate the use of standard, reusable passwords
   cannot be overstated, it is recognized that some organizations may
   still be using them. While it's recommended that these organizations
   transition to the use of better technology, in the mean time, we have
   the following advice to help with the selection and maintenance of
   traditional passwords. But remember, none of these measures provides
   protection against disclosure due to sniffer programs.

   (1) The importance of robust passwords - In many (if not most) cases
        of system penetration, the intruder needs to gain access to an
        account on the system. One way that goal is typically
        accomplished is through guessing the password of a legitimate
        user. This is often accomplished by running an automated
        password cracking program, which utilizes a very large
        dictionary, against the system's password file. The only way to
        guard against passwords being disclosed in this manner is
        through the careful selection of passwords which cannot be
        easily guessed (i.e., combinations of numbers, letters, and
        punctuation characters). Passwords should also be as long as
        the system supports and users can tolerate.

(2) Changing default passwords - Many operating systems and
     application programs are installed with default accounts and
     passwords. These must be changed immediately to something that
     cannot be guessed or cracked.

(3) Restricting access to the password file - In particular, a site
     wants to protect the encrypted password portion of the file so
     that would-be intruders don't have them available for cracking.
     One effective technique is to use shadow passwords where the
     password field of the standard file contains a dummy or false
     password. The file containing the legitimate passwords are
     protected elsewhere on the system.

(4) Password aging - When and how to expire passwords is still a
     subject of controversy among the security community. It is
     generally accepted that a password should not be maintained once
     an account is no longer in use, but it is hotly debated whether
     a user should be forced to change a good password that's in
     active use. The arguments for changing passwords relate to the
     prevention of the continued use of penetrated accounts.
     However, the opposition claims that frequent password changes
     lead to users writing down their passwords in visible areas
     (such as pasting them to a terminal), or to users selecting very
     simple passwords that are easy to guess. It should also be
     stated that an intruder will probably use a captured or guessed
     password sooner rather than later, in which case password aging
     provides little if any protection.

     While there is no definitive answer to this dilemma, a password
     policy should directly address the issue and provide guidelines
     for how often a user should change the password. Certainly, an
     annual change in their password is usually not difficult for
     most users, and you should consider requiring it. It is
     recommended that passwords be changed at least whenever a
     privileged account is compromised, there is a critical change in
     personnel (especially if it is an administrator!), or when an
     account has been compromised. In addition, if a privileged
     account password is compromised, all passwords on the system
     should be changed.

(5) Password/account blocking - Some sites find it useful to disable
     accounts after a predefined number of failed attempts to
     authenticate. If your site decides to employ this mechanism, it
     is recommended that the mechanism not "advertise" itself. After

disabling, even if the correct password is presented, the
message displayed should remain that of a failed login attempt.
Implementing this mechanism will require that legitimate users
contact their system administrator to request that their account
be reactivated.

(6) A word about the finger daemon - By default, the finger daemon
     displays considerable system and user information. For example,
     it can display a list of all users currently using a system, or
     all the contents of a specific user's .plan file. This
     information can be used by would-be intruders to identify
     usernames and guess their passwords. It is recommended that
     sites consider modifying finger to restrict the information
     displayed.

4.2 Confidentiality

   There will be information assets that your site will want to protect
   from disclosure to unauthorized entities. Operating systems often
   have built-in file protection mechanisms that allow an administrator
   to control who on the system can access, or "see," the contents of a
   given file. A stronger way to provide confidentiality is through
   encryption. Encryption is accomplished by scrambling data so that it
   is very difficult and time consuming for anyone other than the
   authorized recipients or owners to obtain the plain text. Authorized
   recipients and the owner of the information will possess the
   corresponding decryption keys that allow them to easily unscramble
   the text to a readable (clear text) form. We recommend that sites
   use encryption to provide confidentiality and protect valuable
   information.

   The use of encryption is sometimes controlled by governmental and
   site regulations, so we encourage administrators to become informed
   of laws or policies that regulate its use before employing it. It is
   outside the scope of this document to discuss the various algorithms
   and programs available for this purpose, but we do caution against
   the casual use of the UNIX crypt program as it has been found to be
   easily broken. We also encourage everyone to take time to understand
   the strength of the encryption in any given algorithm/product before
   using it. Most well-known products are well-documented in the
   literature, so this should be a fairly easy task.

4.3 Integrity

   As an administrator, you will want to make sure that information
   (e.g., operating system files, company data, etc.) has not been
   altered in an unauthorized fashion. This means you will want to
   provide some assurance as to the integrity of the information on your

systems. One way to provide this is to produce a checksum of the
unaltered file, store that checksum offline, and periodically (or
when desired) check to make sure the checksum of the online file
hasn't changed (which would indicate the data has been modified).

Some operating systems come with checksumming programs, such as the
UNIX sum program. However, these may not provide the protection you
actually need. Files can be modified in such a way as to preserve
the result of the UNIX sum program! Therefore, we suggest that you
use a cryptographically strong program, such as the message digesting
program MD5 [ref], to produce the checksums you will be using to
assure integrity.

There are other applications where integrity will need to be assured,
such as when transmitting an email message between two parties. There
are products available that can provide this capability. Once you
identify that this is a capability you need, you can go about
identifying technologies that will provide it.

4.4 Authorization

   Authorization refers to the process of granting privileges to
   processes and, ultimately, users. This differs from authentication
   in that authentication is the process used to identify a user. Once
   identified (reliably), the privileges, rights, property, and
   permissible actions of the user are determined by authorization.

   Explicitly listing the authorized activities of each user (and user
   process) with respect to all resources (objects) is impossible in a
   reasonable system. In a real system certain techniques are used to
   simplify the process of granting and checking authorization(s).

   One approach, popularized in UNIX systems, is to assign to each
   object three classes of user: owner, group and world. The owner is
   either the creator of the object or the user assigned as owner by the
   super-user. The owner permissions (read, write and execute) apply
   only to the owner. A group is a collection of users which share
   access rights to an object. The group permissions (read, write and
   execute) apply to all users in the group (except the owner). The
   world refers to everybody else with access to the system. The world
   permissions (read, write and execute) apply to all users (except the
   owner and members of the group).

   Another approach is to attach to an object a list which explicitly
   contains the identity of all permitted users (or groups). This is an
   Access Control List (ACL). The advantage of ACLs are that they are

easily maintained (one central list per object) and it's very easy to
visually check who has access to what. The disadvantages are the
extra resources required to store such lists, as well as the vast
number of such lists required for large systems.

4.5 Access

4.5.1 Physical Access

   Restrict physical access to hosts, allowing access only to those
   people who are supposed to use the hosts. Hosts include "trusted"
   terminals (i.e., terminals which allow unauthenticated use such as
   system consoles, operator terminals and terminals dedicated to
   special tasks), and individual microcomputers and workstations,
   especially those connected to your network. Make sure people's work
   areas mesh well with access restrictions; otherwise they will find
   ways to circumvent your physical security (e.g., jamming doors open).

   Keep original and backup copies of data and programs safe. Apart
   from keeping them in good condition for backup purposes, they must be
   protected from theft. It is important to keep backups in a separate
   location from the originals, not only for damage considerations, but
   also to guard against thefts.

   Portable hosts are a particular risk. Make sure it won't cause
   problems if one of your staff's portable computer is stolen.
   Consider developing guidelines for the kinds of data that should be
   allowed to reside on the disks of portable computers as well as how
   the data should be protected (e.g., encryption) when it is on a
   portable computer.

   Other areas where physical access should be restricted is the wiring
   closets and important network elements like file servers, name server
   hosts, and routers.

4.5.2 Walk-up Network Connections

   By "walk-up" connections, we mean network connection points located
   to provide a convenient way for users to connect a portable host to
   your network.

   Consider whether you need to provide this service, bearing in mind
   that it allows any user to attach an unauthorized host to your
   network. This increases the risk of attacks via techniques such as

IP address spoofing, packet sniffing, etc. Users and site management
must appreciate the risks involved. If you decide to provide walk-up
connections, plan the service carefully and define precisely where
you will provide it so that you can ensure the necessary physical
access security.

A walk-up host should be authenticated before its user is permitted
to access resources on your network. As an alternative, it may be
possible to control physical access. For example, if the service is
to be used by students, you might only provide walk-up connection
sockets in student laboratories.

If you are providing walk-up access for visitors to connect back to
their home networks (e.g., to read e-mail, etc.) in your facility,
consider using a separate subnet that has no connectivity to the
internal network.

Keep an eye on any area that contains unmonitored access to the
network, such as vacant offices. It may be sensible to disconnect
such areas at the wiring closet, and consider using secure hubs and
monitoring attempts to connect unauthorized hosts.

4.5.3 Other Network Technologies

Technologies considered here include X.25, ISDN, SMDS, DDS and Frame
Relay. All are provided via physical links which go through
telephone exchanges, providing the potential for them to be diverted.
Crackers are certainly interested in telephone switches as well as in
data networks!

With switched technologies, use Permanent Virtual Circuits or Closed
User Groups whenever this is possible. Technologies which provide
authentication and/or encryption (such as IPv6) are evolving rapidly;
consider using them on links where security is important.

4.5.4 Modems

4.5.4.1 Modem Lines Must Be Managed

Although they provide convenient access to a site for its users, they
can also provide an effective detour around the site's firewalls.
For this reason it is essential to maintain proper control of modems.

Don't allow users to install a modem line without proper
authorization. This includes temporary installations (e.g., plugging
a modem into a facsimile or telephone line overnight).

Maintain a register of all your modem lines and keep your register up
to date. Conduct regular (ideally automated) site checks for
unauthorized modems.

4.5.4.2 Dial-in Users Must Be Authenticated

A username and password check should be completed before a user can
access anything on your network. Normal password security
considerations are particularly important (see section 4.1.1).

Remember that telephone lines can be tapped, and that it is quite
easy to intercept messages to cellular phones. Modern high-speed
modems use more sophisticated modulation techniques, which makes them
somewhat more difficult to monitor, but it is prudent to assume that
hackers know how to eavesdrop on your lines. For this reason, you
should use one-time passwords if at all possible.

It is helpful to have a single dial-in point (e.g., a single large
modem pool) so that all users are authenticated in the same way.

Users will occasionally mis-type a password. Set a short delay - say
two seconds - after the first and second failed logins, and force a
disconnect after the third. This will slow down automated password
attacks. Don't tell the user whether the username, the password, or
both, were incorrect.

4.5.4.3 Call-back Capability

Some dial-in servers offer call-back facilities (i.e., the user dials
in and is authenticated, then the system disconnects the call and
calls back on a specified number). Call-back is useful since if
someone were to guess a username and password, they are disconnected,
and the system then calls back the actual user whose password was
cracked; random calls from a server are suspicious, at best. This
does mean users may only log in from one location (where the server
is configured to dial them back), and of course there may be phone
charges associated with there call-back location.

This feature should be used with caution; it can easily be bypassed.
At a minimum, make sure that the return call is never made from the
same modem as the incoming one. Overall, although call-back can
improve modem security, you should not depend on it alone.

4.5.4.4 All Logins Should Be Logged

All logins, whether successful or unsuccessful should be logged.
However, do not keep correct passwords in the log. Rather, log them
simply as a successful login attempt. Since most bad passwords are

   mistyped by authorized users, they only vary by a single character
   from the actual password. Therefore if you can't keep such a log
   secure, don't log it at all.

   If Calling Line Identification is available, take advantage of it by
   recording the calling number for each login attempt. Be sensitive to
   the privacy issues raised by Calling Line Identification. Also be
   aware that Calling Line Identification is not to be trusted (since
   intruders have been known to break into phone switches and forward
   phone numbers or make other changes); use the data for informational
   purposes only, not for authentication.

4.5.4.5 Choose Your Opening Banner Carefully

   Many sites use a system default contained in a message of the day
   file for their opening banner. Unfortunately, this often includes the
   type of host hardware or operating system present on the host. This
   can provide valuable information to a would-be intruder. Instead,
   each site should create its own specific login banner, taking care to
   only include necessary information.

   Display a short banner, but don't offer an "inviting" name (e.g.,
   University of XYZ, Student Records System). Instead, give your site
   name, a short warning that sessions may be monitored, and a
   username/password prompt. Verify possible legal issues related to
   the text you put into the banner.

   For high-security applications, consider using a "blind" password
   (i.e., give no response to an incoming call until the user has typed
   in a password). This effectively simulates a dead modem.

4.5.4.6 Dial-out Authentication

   Dial-out users should also be authenticated, particularly since your
   site will have to pay their telephone charges.

   Never allow dial-out from an unauthenticated dial-in call, and
   consider whether you will allow it from an authenticated one. The
   goal here is to prevent callers using your modem pool as part of a
   chain of logins. This can be hard to detect, particularly if a
   hacker sets up a path through several hosts on your site.

   At a minimum, don't allow the same modems and phone lines to be used
   for both dial-in and dial-out. This can be implemented easily if you
   run separate dial-in and dial-out modem pools.

4.5.4.7 Make Your Modem Programming as "Bullet-proof" as Possible

   Be sure modems can't be reprogrammed while they're in service. At a
   minimum, make sure that three plus signs won't put your dial-in
   modems into command mode!

   Program your modems to reset to your standard configuration at the
   start of each new call. Failing this, make them reset at the end of
   each call. This precaution will protect you against accidental
   reprogramming of your modems. Resetting at both the end and the
   beginning of each call will assure an even higher level of confidence
   that a new caller will not inherit a previous caller's session.

   Check that your modems terminate calls cleanly. When a user logs out
   from an access server, verify that the server hangs up the phone line
   properly. It is equally important that the server forces logouts
   from whatever sessions were active if the user hangs up unexpectedly.

4.6 Auditing

   This section covers the procedures for collecting data generated by
   network activity, which may be useful in analyzing the security of a
   network and responding to security incidents.

4.6.1 What to Collect

   Audit data should include any attempt to achieve a different security
   level by any person, process, or other entity in the network. This
   includes login and logout, super user access (or the non-UNIX
   equivalent), ticket generation (for Kerberos, for example), and any
   other change of access or status. It is especially important to note
   "anonymous" or "guest" access to public servers.

   The actual data to collect will differ for different sites and for
   different types of access changes within a site. In general, the
   information you want to collect includes: username and hostname, for
   login and logout; previous and new access rights, for a change of
   access rights; and a timestamp. Of course, there is much more
   information which might be gathered, depending on what the system
   makes available and how much space is available to store that
   information.

   One very important note: do not gather passwords. This creates an
   enormous potential security breach if the audit records should be
   improperly accessed. Do not gather incorrect passwords either, as
   they often differ from valid passwords by only a single character or
   transposition.

4.6.2 Collection Process

   The collection process should be enacted by the host or resource
   being accessed. Depending on the importance of the data and the need
   to have it local in instances in which services are being denied,
   data could be kept local to the resource until needed or be
   transmitted to storage after each event.

   There are basically three ways to store audit records: in a
   read/write file on a host, on a write-once/read-many device (e.g., a
   CD-ROM or a specially configured tape drive), or on a write-only
   device (e.g., a line printer). Each method has advantages and
   disadvantages.

   File system logging is the least resource intensive of the three
   methods and the easiest to configure. It allows instant access to
   the records for analysis, which may be important if an attack is in
   progress. File system logging is also the least reliable method. If
   the logging host has been compromised, the file system is usually the
   first thing to go; an intruder could easily cover up traces of the
   intrusion.

   Collecting audit data on a write-once device is slightly more effort
   to configure than a simple file, but it has the significant advantage
   of greatly increased security because an intruder could not alter the
   data showing that an intrusion has occurred. The disadvantage of
   this method is the need to maintain a supply of storage media and the
   cost of that media. Also, the data may not be instantly available.

   Line printer logging is useful in system where permanent and
   immediate logs are required. A real time system is an example of
   this, where the exact point of a failure or attack must be recorded.
   A laser printer, or other device which buffers data (e.g., a print
   server), may suffer from lost data if buffers contain the needed data
   at a critical instant. The disadvantage of, literally, "paper
   trails" is the need to keep the printer fed and the need to scan
   records by hand. There is also the issue of where to store the,
   potentially, enormous volume of paper which may be generated.

   For each of the logging methods described, there is also the issue of
   securing the path between the device generating the log and actual
   logging device (i.e., the file server, tape/CD-ROM drive, printer).
   If that path is compromised, logging can be stopped or spoofed or
   both. In an ideal world, the logging device would be directly

attached by a single, simple, point-to-point cable. Since that is
usually impractical, the path should pass through the minimum number
of networks and routers. Even if logs can be blocked, spoofing can
be prevented with cryptographic checksums (it probably isn't
necessary to encrypt the logs because they should not contain
sensitive information in the first place).

4.6.3 Collection Load

Collecting audit data may result in a rapid accumulation of bytes so
storage availability for this information must be considered in
advance. There are a few ways to reduce the required storage space.
First, data can be compressed, using one of many methods. Or, the
required space can be minimized by keeping data for a shorter period
of time with only summaries of that data kept in long-term archives.
One major drawback to the latter method involves incident response.
Often, an incident has been ongoing for some period of time when a
site notices it and begins to investigate. At that point in time,
it's very helpful to have detailed audit logs available. If these are
just summaries, there may not be sufficient detail to fully handle
the incident.

4.6.4 Handling and Preserving Audit Data

Audit data should be some of the most carefully secured data at the
site and in the backups. If an intruder were to gain access to audit
logs, the systems themselves, in addition to the data, would be at
risk.

Audit data may also become key to the investigation, apprehension,
and prosecution of the perpetrator of an incident. For this reason,
it is advisable to seek the advice of legal council when deciding how
audit data should be treated. This should happen before an incident
occurs.

If a data handling plan is not adequately defined prior to an
incident, it may mean that there is no recourse in the aftermath of
an event, and it may create liability resulting from improper
treatment of the data.

4.6.5 Legal Considerations

Due to the content of audit data, there are a number of legal
questions that arise which might need to be addressed by your legal
counsel. If you collect and save audit data, you need to be prepared
for consequences resulting both from its existence and its content.

One area concerns the privacy of individuals. In certain instances,
audit data may contain personal information. Searching through the
data, even for a routine check of the system's security, could
represent an invasion of privacy.

A second area of concern involves knowledge of intrusive behavior
originating from your site. If an organization keeps audit data, is
it responsible for examining it to search for incidents? If a host
in one organization is used as a launching point for an attack
against another organization, can the second organization use the
audit data of the first organization to prove negligence on the part
of that organization?

The above examples are meant to be comprehensive, but should motivate
your organization to consider the legal issues involved with audit
data.

4.7 Securing Backups

The procedure of creating backups is a classic part of operating a
computer system. Within the context of this document, backups are
addressed as part of the overall security plan of a site. There are
several aspects to backups that are important within this context:

(1) Make sure your site is creating backups
(2) Make sure your site is using offsite storage for backups. The
     storage site should be carefully selected for both its security
     and its availability.
(3) Consider encrypting your backups to provide additional protection
     of the information once it is off-site. However, be aware that
     you will need a good key management scheme so that you'll be
     able to recover data at any point in the future. Also, make
     sure you will have access to the necessary decryption programs
     at such time in the future as you need to perform the
     decryption.
(4) Don't always assume that your backups are good. There have been
     many instances of computer security incidents that have gone on
     for long periods of time before a site has noticed the incident.
     In such cases, backups of the affected systems are also tainted.
(5) Periodically verify the correctness and completeness of your
     backups.

5. Security Incident Handling

This chapter of the document will supply guidance to be used before,
during, and after a computer security incident occurs on a host,
network, site, or multi-site environment. The operative philosophy
in the event of a breach of computer security is to react according

to a plan. This is true whether the breach is the result of an
external intruder attack, unintentional damage, a student testing
some new program to exploit a software vulnerability, or a
disgruntled employee. Each of the possible types of events, such as
those just listed, should be addressed in advance by adequate
contingency plans.

Traditional computer security, while quite important in the overall
site security plan, usually pays little attention to how to actually
handle an attack once one occurs. The result is that when an attack
is in progress, many decisions are made in haste and can be damaging
to tracking down the source of the incident, collecting evidence to
be used in prosecution efforts, preparing for the recovery of the
system, and protecting the valuable data contained on the system.

One of the most important, but often overlooked, benefits for
efficient incident handling is an economic one. Having both
technical and managerial personnel respond to an incident requires
considerable resources. If trained to handle incidents efficiently,
less staff time is required when one occurs.

Due to the world-wide network most incidents are not restricted to a
single site. Operating systems vulnerabilities apply (in some cases)
to several millions of systems, and many vulnerabilities are
exploited within the network itself. Therefore, it is vital that all
sites with involved parties be informed as soon as possible.

Another benefit is related to public relations. News about computer
security incidents tends to be damaging to an organization's stature
among current or potential clients. Efficient incident handling
minimizes the potential for negative exposure.

A final benefit of efficient incident handling is related to legal
issues. It is possible that in the near future organizations may be
held responsible because one of their nodes was used to launch a
network attack. In a similar vein, people who develop patches or
workarounds may be sued if the patches or workarounds are
ineffective, resulting in compromise of the systems, or, if the
patches or workarounds themselves damage systems. Knowing about
operating system vulnerabilities and patterns of attacks, and then
taking appropriate measures to counter these potential threats, is
critical to circumventing possible legal problems.

The sections in this chapter provide an outline and starting point
for creating your site's policy for handling security incidents. The
sections are:

(1) Preparing and planning (what are the goals and objectives in
     handling an incident).
(2) Notification (who should be contacted in the case of an
     incident).
        - Local managers and personnel
        - Law enforcement and investigative agencies
        - Computer security incidents handling teams
        - Affected and involved sites
        - Internal communications
        - Public relations and press releases
(3) Identifying an incident (is it an incident and how serious is
     it).
(4) Handling (what should be done when an incident occurs).
        - Notification (who should be notified about the incident)
        - Protecting evidence and activity logs (what records should be
          kept from before, during, and after the incident)
        - Containment (how can the damage be limited)
        - Eradication (how to eliminate the reasons for the incident)
        - Recovery (how to reestablish service and systems)
        - Follow Up (what actions should be taken after the incident)
(5) Aftermath (what are the implications of past incidents).
(6) Administrative response to incidents.

The remainder of this chapter will detail the issues involved in each
of the important topics listed above, and provide some guidance as to
what should be included in a site policy for handling incidents.

5.1 Preparing and Planning for Incident Handling

Part of handling an incident is being prepared to respond to an
incident before the incident occurs in the first place. This
includes establishing a suitable level of protections as explained in
the preceding chapters. Doing this should help your site prevent
incidents as well as limit potential damage resulting from them when
they do occur. Protection also includes preparing incident handling
guidelines as part of a contingency plan for your organization or
site. Having written plans eliminates much of the ambiguity which
occurs during an incident, and will lead to a more appropriate and
thorough set of responses. It is vitally important to test the
proposed plan before an incident occurs through "dry runs". A team
might even consider hiring a tiger team to act in parallel with the
dry run. (Note: a tiger team is a team of specialists that try to
penetrate the security of a system.)

Learning to respond efficiently to an incident is important for a
number of reasons:

(1)   Protecting the assets which could be compromised
(2)   Protecting resources which could be utilized more
      profitably if an incident did not require their services
(3)   Complying with (government or other) regulations
(4)   Preventing the use of your systems in attacks against other
      systems (which could cause you to incur legal liability)
(5)   Minimizing the potential for negative exposure

As in any set of pre-planned procedures, attention must be paid to a
set of goals for handling an incident. These goals will be
prioritized differently depending on the site. A specific set of
objectives can be identified for dealing with incidents:

(1)   Figure out how it happened.
(2)   Find out how to avoid further exploitation of the same
        vulnerability.
(3)   Avoid escalation and further incidents.
(4)   Assess the impact and damage of the incident.
(5)   Recover from the incident.
(6)   Update policies and procedures as needed.
(7)   Find out who did it (if appropriate and possible).

Due to the nature of the incident, there might be a conflict between
analyzing the original source of a problem and restoring systems and
services. Overall goals (like assuring the integrity of critical
systems) might be the reason for not analyzing an incident. Of
course, this is an important management decision; but all involved
parties must be aware that without analysis the same incident may
happen again.

It is also important to prioritize the actions to be taken during an
incident well in advance of the time an incident occurs. Sometimes
an incident may be so complex that it is impossible to do everything
at once to respond to it; priorities are essential. Although
priorities will vary from institution to institution, the following
suggested priorities may serve as a starting point for defining your
organization's response:

(1)   Priority one -- protect human life and people's
      safety; human life always has precedence over all
      other considerations.

(2)   Priority two -- protect classified and/or sensitive
      data. Prevent exploitation of classified and/or
      sensitive systems, networks or sites. Inform affected

          classified and/or sensitive systems, networks or sites
          about already occurred penetrations.
          (Be aware of regulations by your site or by government)

     (3) Priority three -- protect other data, including
          proprietary, scientific, managerial and other data,
          because loss of data is costly in terms of resources.
          Prevent exploitations of other systems, networks or
          sites and inform already affected systems, networks or
          sites about successful penetrations.

     (4) Priority four -- prevent damage to systems (e.g., loss
          or alteration of system files, damage to disk drives,
          etc.). Damage to systems can result in costly down
          time and recovery.

     (5) Priority five -- minimize disruption of computing
          resources (including processes). It is better in many
          cases to shut a system down or disconnect from a network
          than to risk damage to data or systems. Sites will have
          to evaluate the trade-offs between shutting down and
          disconnecting, and staying up. There may be service
          agreements in place that may require keeping systems
          up even in light of further damage occurring. However,
          the damage and scope of an incident may be so extensive
          that service agreements may have to be over-ridden.

An important implication for defining priorities is that once human
life and national security considerations have been addressed, it is
generally more important to save data than system software and
hardware. Although it is undesirable to have any damage or loss
during an incident, systems can be replaced. However, the loss or
compromise of data (especially classified or proprietary data) is
usually not an acceptable outcome under any circumstances.

Another important concern is the effect on others, beyond the systems
and networks where the incident occurs. Within the limits imposed by
government regulations it is always important to inform affected
parties as soon as possible. Due to the legal implications of this
topic, it should be included in the planned procedures to avoid
further delays and uncertainties for the administrators.

Any plan for responding to security incidents should be guided by
local policies and regulations. Government and private sites that
deal with classified material have specific rules that they must
follow.

The policies chosen by your site on how it reacts to incidents will
shape your response. For example, it may make little sense to create
mechanisms to monitor and trace intruders if your site does not plan
to take action against the intruders if they are caught. Other
organizations may have policies that affect your plans. Telephone
companies often release information about telephone traces only to
law enforcement agencies.

Handling incidents can be tedious and require any number of routine
tasks that could be handled by support personnel. To free the
technical staff it may be helpful to identify support staff who will
help with tasks like: photocopying, fax'ing, etc.

5.2 Notification and Points of Contact

   It is important to establish contacts with various personnel before a
   real incident occurs. Many times, incidents are not real
   emergencies. Indeed, often you will be able to handle the activities
   internally. However, there will also be many times when others
   outside your immediate department will need to be included in the
   incident handling. These additional contacts include local managers
   and system administrators, administrative contacts for other sites on
   the Internet, and various investigative organizations. Getting to
   know these contacts before incidents occurs will help to make your
   incident handling process more efficient.

   For each type of communication contact, specific "Points of Contact"
   (POC) should be defined. These may be technical or administrative in
   nature and may include legal or investigative agencies as well as
   service providers and vendors. When establishing these contact, it
   is important to decide how much information will be shared with each
   class of contact. It is especially important to define, ahead of
   time, what information will be shared with the users at a site, with
   the public (including the press), and with other sites.

   Settling these issues are especially important for the local person
   responsible for handling the incident, since that is the person
   responsible for the actual notification of others. A list of
   contacts in each of these categories is an important time saver for
   this person during an incident. It can be quite difficult to find an
   appropriate person during an incident when many urgent events are
   ongoing. It is strongly recommended that all relevant telephone
   numbers (also electronic mail addresses and fax numbers) be included
   in the site security policy. The names and contact information of
   all individuals who will be directly involved in the handling of an
   incident should be placed at the top of this list.

5.2.1 Local Managers and Personnel

   When an incident is under way, a major issue is deciding who is in
   charge of coordinating the activity of the multitude of players. A
   major mistake that can be made is to have a number of people who are
   each working independently, but are not working together. This will
   only add to the confusion of the event and will probably lead to
   wasted or ineffective effort.

   The single POC may or may not be the person responsible for handling
   the incident. There are two distinct roles to fill when deciding who
   shall be the POC and who will be the person in charge of the
   incident. The person in charge of the incident will make decisions
   as to the interpretation of policy applied to the event. In
   contrast, the POC must coordinate the effort of all the parties
   involved with handling the event.

   The POC must be a person with the technical expertise to successfully
   coordinate the efforts of the system managers and users involved in
   monitoring and reacting to the attack. Care should be taken when
   identifying who this person will be. It should not necessarily be
   the same person who has administrative responsibility for the
   compromised systems since often such administrators have knowledge
   only sufficient for the day to day use of the computers, and lack in
   depth technical expertise.

   Another important function of the POC is to maintain contact with law
   enforcement and other external agencies to assure that multi-agency
   involvement occurs. The level of involvement will be determined by
   management decisions as well as legal constraints.

   A single POC should also be the single person in charge of collecting
   evidence, since as a rule of thumb, the more people that touch a
   potential piece of evidence, the greater the possibility that it will
   be inadmissible in court. To ensure that evidence will be acceptable
   to the legal community, collecting evidence should be done following
   predefined procedures in accordance with local laws and legal
   regulations.

   One of the most critical tasks for the POC is the coordination of all
   relevant processes. Responsibilities may be distributed over the
   whole site, involving multiple independent departments or groups.
   This will require a well coordinated effort in order to achieve
   overall success. The situation becomes even more complex if multiple
   sites are involved. When this happens, rarely will a single POC at
   one site be able to adequately coordinate the handling of the entire
   incident. Instead, appropriate incident response teams should be
   involved.

The incident handling process should provide some escalation
mechanisms. In order to define such a mechanism, sites will need to
create an internal classification scheme for incidents. Associated
with each level of incident will be the appropriate POC and
procedures. As an incident is escalated, there may be a change in
the POC which will need to be communicated to all others involved in
handling the incident. When a change in the POC occurs, old POC
should brief the new POC in all background information.

Lastly, users must know how to report suspected incidents. Sites
should establish reporting procedures that will work both during and
outside normal working hours. Help desks are often used to receive
these reports during normal working hours, while beepers and
telephones can be used for out of hours reporting.

5.2.2 Law Enforcement and Investigative Agencies

In the event of an incident that has legal consequences, it is
important to establish contact with investigative agencies (e.g, the
FBI and Secret Service in the U.S.) as soon as possible. Local law
enforcement, local security offices, and campus police departments
should also be informed as appropriate.   This section describes many
of the issues that will be confronted, but it is acknowledged that
each organization will have its own local and governmental laws and
regulations that will impact how they interact with law enforcement
and investigative agencies. The most important point to make is that
each site needs to work through these issues.

A primary reason for determining these point of contact well in
advance of an incident is that once a major attack is in progress,
there is little time to call these agencies to determine exactly who
the correct point of contact is. Another reason is that it is
important to cooperate with these agencies in a manner that will
foster a good working relationship, and that will be in accordance
with the working procedures of these agencies. Knowing the working
procedures in advance, and the expectations of your point of contact
is a big step in this direction. For example, it is important to
gather evidence that will be admissible in any subsequent legal
proceedings, and this will require prior knowledge of how to gather
such evidence. A final reason for establishing contacts as soon as
possible is that it is impossible to know the particular agency that
will assume jurisdiction in any given incident. Making contacts and
finding the proper channels early on will make responding to an
incident go considerably more smoothly.

If your organization or site has a legal counsel, you need to notify
this office soon after you learn that an incident is in progress. At
a minimum, your legal counsel needs to be involved to protect the
legal and financial interests of your site or organization. There
are many legal and practical issues, a few of which are:

(1)   Whether your site or organization is willing to risk negative
      publicity or exposure to cooperate with legal prosecution
      efforts.

(2)   Downstream liability--if you leave a compromised system as is so
      it can be monitored and another computer is damaged because the
      attack originated from your system, your site or organization
      may be liable for damages incurred.

(3)   Distribution of information--if your site or organization
      distributes information about an attack in which another site or
      organization may be involved or the vulnerability in a product
      that may affect ability to market that product, your site or
      organization may again be liable for any damages (including
      damage of reputation).

(4)   Liabilities due to monitoring--your site or organization may be
      sued if users at your site or elsewhere discover that your site
      is monitoring account activity without informing users.

Unfortunately, there are no clear precedents yet on the liabilities
or responsibilities of organizations involved in a security incident
or who might be involved in supporting an investigative effort.
Investigators will often encourage organizations to help trace and
monitor intruders. Indeed, most investigators cannot pursue computer
intrusions without extensive support from the organizations involved.
However, investigators cannot provide protection from liability
claims, and these kinds of efforts may drag out for months and may
take a lot of effort.

On the other hand, an organization's legal council may advise extreme
caution and suggest that tracing activities be halted and an intruder
shut out of the system. This, in itself, may not provide protection
from liability, and may prevent investigators from identifying the
perpetrator.

The balance between supporting investigative activity and limiting
liability is tricky. You'll need to consider the advice of your legal
counsel and the damage the intruder is causing (if any) when making
your decision about what to do during any particular incident.

Your legal counsel should also be involved in any decision to contact
investigative agencies when an incident occurs at your site. The
decision to coordinate efforts with investigative agencies is most
properly that of your site or organization. Involving your legal
counsel will also foster the multi-level coordination between your
site and the particular investigative agency involved, which in turn
results in an efficient division of labor. Another result is that
you are likely to obtain guidance that will help you avoid future
legal mistakes.

Finally, your legal counsel should evaluate your site's written
procedures for responding to incidents. It is essential to obtain a
"clean bill of health" from a legal perspective before you actually
carry out these procedures.

It is vital, when dealing with investigative agencies, to verify that
the person who calls asking for information is a legitimate
representative from the agency in question. Unfortunately, many well
intentioned people have unknowingly leaked sensitive details about
incidents, allowed unauthorized people into their systems, etc.,
because a caller has masqueraded as a representative of a government
agency. (Note: this word of caution actually applies to all external
contacts.)

A similar consideration is using a secure means of communication.
Because many network attackers can easily re-route electronic mail,
avoid using electronic mail to communicate with other agencies (as
well as others dealing with the incident at hand). Non-secured phone
lines (the phones normally used in the business world) are also
frequent targets for tapping by network intruders, so be careful!

There is no one established set of rules for responding to an
incident when the local government becomes involved. Normally (in
the U.S.), except by legal order, no agency can force you to monitor,
to disconnect from the network, to avoid telephone contact with the
suspected attackers, etc. Each organization will have a set of local
and national laws and regulations that must be adhered to when
handling incidents. It is recommended that each site be familiar with
those laws and regulations, and identify and get know the contacts
for agencies with jurisdiction well in advance of handling an
incident.

5.2.3 Computer Security Incident Handling Teams

There are currently a number of of Computer Security Incident
Response teams (CSIRTs) such as the CERT Coordination Center, the
German DFN-CERT, and other teams around the globe. Teams exist for
many major government agencies and large corporations. If such a

team is available, notifying it should be of primary consideration
during the early stages of an incident. These teams are responsible
for coordinating computer security incidents over a range of sites
and larger entities. Even if the incident is believed to be
contained within a single site, it is possible that the information
available through a response team could help in fully resolving the
incident.

If it is determined that the breach occurred due to a flaw in the
system's hardware or software, the vendor (or supplier) and a
Computer Security Incident Handling team should be notified as soon
as possible. This is especially important because many other systems
are vulnerable, and these vendor and response team organizations can
help disseminate help to other affected sites.

In setting up a site policy for incident handling, it may be
desirable to create a subgroup, much like those teams that already
exist, that will be responsible for handling computer security
incidents for the site (or organization). If such a team is created,
it is essential that communication lines be opened between this team
and other teams. Once an incident is under way, it is difficult to
open a trusted dialogue between other teams if none has existed
before.

5.2.4 Affected and Involved Sites

If an incident has an impact on other sites, it is good practice to
inform them. It may be obvious from the beginning that the incident
is not limited to the local site, or it may emerge only after further
analysis.

Each site may choose to contact other sites directly or they can pass
the information to an appropriate incident response team. It is often
very difficult to find the responsible POC at remote sites and the
incident response team will be able to facilitate contact by making
use of already established channels.

The legal and liability issues arising from a security incident will
differ from site to site. It is important to define a policy for the
sharing and logging of information about other sites before an
incident occurs.

Information about specific people is especially sensitive, and may be
subject to privacy laws. To avoid problems in this area, irrelevant
information should be deleted and a statement of how to handle the
remaining information should be included. A clear statement of how
this information is to be used is essential. No one who informs a
site of a security incident wants to read about it in the public

press. Incident response teams are valuable in this respect. When
they pass information to responsible POCs, they are able to protect
the anonymity of the original source. But, be aware that, in many
cases, the analysis of logs and information at other sites will
reveal addresses of your site.

All the problems discussed above should be not taken as reasons not
to involve other sites. In fact, the experiences of existing teams
reveal that most sites informed about security problems are not even
aware that their site had been compromised. Without timely
information, other sites are often unable to take action against
intruders.

5.2.5 Internal Communications

It is crucial during a major incident to communicate why certain
actions are being taken, and how the users (or departments) are
expected to behave. In particular, it should be made very clear to
users what they are allowed to say (and not say) to the outside world
(including other departments). For example, it wouldn't be good for
an organization if users replied to customers with something like,
"I'm sorry the systems are down, we've had an intruder and we are
trying to clean things up." It would be much better if they were
instructed to respond with a prepared statement like, "I'm sorry our
systems are unavailable, they are being maintained for better service
in the future."

Communications with customers and contract partners should be handled
in a sensible, but sensitive way. One can prepare for the main issues
by preparing a checklist. When an incident occurs, the checklist can
be used with the addition of a sentence or two for the specific
circumstances of the incident.

Public relations departments can be very helpful during incidents.
They should be involved in all planning and can provide well
constructed responses for use when contact with outside departments
and organizations is necessary.

5.2.6 Public Relations - Press Releases

There has been a tremendous growth in the amount of media coverage
dedicated to computer security incidents in the United States. Such
press coverage is bound to extend to other countries as the Internet
continues to grow and expand internationally. Readers from countries
where such media attention has not yet occurred, can learn from the
experiences in the U.S. and should be forwarned and prepared.

One of the most important issues to consider is when, who, and how
much to release to the general public through the press. There are
many issues to consider when deciding this particular issue. First
and foremost, if a public relations office exists for the site, it is
important to use this office as liaison to the press. The public
relations office is trained in the type and wording of information
released, and will help to assure that the image of the site is
protected during and after the incident (if possible). A public
relations office has the advantage that you can communicate candidly
with them, and provide a buffer between the constant press attention
and the need of the POC to maintain control over the incident.

If a public relations office is not available, the information
released to the press must be carefully considered. If the
information is sensitive, it may be advantageous to provide only
minimal or overview information to the press. It is quite possible
that any information provided to the press will be quickly reviewed
by the perpetrator of the incident. Also note that misleading the
press can often backfire and cause more damage than releasing
sensitive information.

While it is difficult to determine in advance what level of detail to
provide to the press, some guidelines to keep in mind are:

(1) Keep the technical level of detail low. Detailed
     information about the incident may provide enough
     information for others to launch similar attacks on
     other sites, or even damage the site's ability to
     prosecute the guilty party once the event is over.

(2) Keep the speculation out of press statements.
     Speculation of who is causing the incident or the
     motives are very likely to be in error and may cause
     an inflamed view of the incident.

(3) Work with law enforcement professionals to assure that
     evidence is protected. If prosecution is involved,
     assure that the evidence collected is not divulged to
     the press.

(4) Try not to be forced into a press interview before you are
     prepared. The popular press is famous for the "2 am"
     interview, where the hope is to catch the interviewee off
     guard and obtain information otherwise not available.

(5) Do not allow the press attention to detract from the
     handling of the event. Always remember that the successful
     closure of an incident is of primary importance.

5.3 Identifying an Incident

5.3.1 Is It Real?

   This stage involves determining if a problem really exists. Of
   course many if not most signs often associated with virus infection,
   system intrusions, malicious users, etc., are simply anomalies such
   as hardware failures or suspicious system/user behavior. To assist
   in identifying whether there really is an incident, it is usually
   helpful to obtain and use any detection software which may be
   available. Audit information is also extremely useful, especially in
   determining whether there is a network attack. It is extremely
   important to obtain a system snapshot as soon as one suspects that
   something is wrong. Many incidents cause a dynamic chain of events
   to occur, and an initial system snapshot may be the most valuable
   tool for identifying the problem and any source of attack. Finally,
   it is important to start a log book. Recording system events,
   telephone conversations, time stamps, etc., can lead to a more rapid
   and systematic identification of the problem, and is the basis for
   subsequent stages of incident handling.

   There are certain indications or "symptoms" of an incident that
   deserve special attention:

   (1)   System crashes.
   (2)   New user accounts (the account RUMPLESTILTSKIN has been
          unexpectedly created), or high activity on a previously
          low usage account.
   (3)   New files (usually with novel or strange file names,
          such as data.xx or k or .xx).
   (4)   Accounting discrepancies (in a UNIX system you might
          notice the shrinking of an accounting file called
          /usr/admin/lastlog, something that should make you very
          suspicious that there may be an intruder).
   (5)   Changes in file lengths or dates (a user should be
          suspicious if .EXE files in an MS DOS computer have
          unexplainedly grown by over 1800 bytes).
   (6)   Attempts to write to system (a system manager notices
          that a privileged user in a VMS system is attempting to
          alter RIGHTSLIST.DAT).
   (7)   Data modification or deletion (files start to disappear).
   (8)   Denial of service (a system manager and all other users
          become locked out of a UNIX system, now in single user mode).
   (9)   Unexplained, poor system performance
   (10)   Anomalies ("GOTCHA" is displayed on the console or there
          are frequent unexplained "beeps").
   (11)   Suspicious probes (there are numerous unsuccessful login
          attempts from another node).

   (12) Suspicious browsing (someone becomes a root user on a UNIX
         system and accesses file after file on many user accounts.)
   (13) Inability of a user to log in due to modifications of his/her
         account.

   By no means is this list comprehensive; we have just listed a number
   of common indicators. It is best to collaborate with other technical
   and computer security personnel to make a decision as a group about
   whether an incident is occurring.

5.3.2 Types and Scope of Incidents

   Along with the identification of the incident is the evaluation of
   the scope and impact of the problem. It is important to correctly
   identify the boundaries of the incident in order to effectively deal
   with it and prioritize responses.

   In order to identify the scope and impact a set of criteria should be
   defined which is appropriate to the site and to the type of
   connections available. Some of the issues include:

   (1) Is this a multi-site incident?
   (2) Are many computers at your site affected by this incident?
   (3) Is sensitive information involved?
   (4) What is the entry point of the incident (network,
        phone line, local terminal, etc.)?
   (5) Is the press involved?
   (6) What is the potential damage of the incident?
   (7) What is the estimated time to close out the incident?
   (8) What resources could be required to handle the incident?
   (9) Is law enforcement involved?

5.3.3 Assessing the Damage and Extent

   The analysis of the damage and extent of the incident can be quite
   time consuming, but should lead to some insight into the nature of
   the incident, and aid investigation and prosecution. As soon as the
   breach has occurred, the entire system and all of its components
   should be considered suspect. System software is the most probable
   target. Preparation is key to be able to detect all changes for a
   possibly tainted system. This includes checksumming all media from
   the vendor using a algorithm which is resistant to tampering. (See
   sections 4.3)

   Assuming original vendor distribution media are available, an
   analysis of all system files should commence, and any irregularities
   should be noted and referred to all parties involved in handling the
   incident. It can be very difficult, in some cases, to decide which

backup media are showing a correct system status. Consider, for
example, that the incident may have continued for months or years
before discovery, and the suspect may be an employee of the site, or
otherwise have intimate knowledge or access to the systems. In all
cases, the pre-incident preparation will determine what recovery is
possible.

If the system supports centralized logging (most do), go back over
the logs and look for abnormalities. If process accounting and
connect time accounting is enabled, look for patterns of system
usage. To a lesser extent, disk usage may shed light on the
incident. Accounting can provide much helpful information in an
analysis of an incident and subsequent prosecution. Your ability to
address all aspects of a specific incident strongly depends on the
success of this analysis.

5.4 Handling an Incident

Certain steps are necessary to take during the handling of an
incident. In all security related activities, the most important
point to be made is that all sites should have policies in place.
Without defined policies and goals, activities undertaken will remain
without focus. The goals should be defined by management and legal
counsel in advance.

One of the most fundamental objectives is to restore control of the
affected systems and to limit the impact and damage. In the worst
case scenario, shutting down the system, or disconnecting the system
from the network, may the only practical solution.

As the activities involved are complex, try to get as much help as
necessary. While trying to solve the problem alone, real damage
might occur due to delays or missing information. Most
administrators take the discovery of an intruder as a personal
challenge. By proceeding this way, other objectives as outlined in
the local policies may not always be considered. Trying to catch
intruders may be a very low priority, compared to system integrity,
for example. Monitoring a hacker's activity is useful, but it might
not be considered worth the risk to allow the continued access.

5.4.1 Types of Notification and Exchange of Information

When you have confirmed that an incident is occurring, the
appropriate personnel must be notified. How this notification is
achieved is very important to keeping the event under control both
from a technical and emotional standpoint. The circumstances should
be described in as much detail as possible, in order to aid prompt
acknowledgment and understanding of the problem. Great care should

be taken when determining to which groups detailed technical
information is given during the notification. For example, it is
helpful to pass this kind of information to an incident handling team
as they can assist you by providing helpful hints for eradicating the
vulnerabilities involved in an incident. On the other hand, putting
the critical knowledge into the public domain (e.g., via USENET
newsgroups or mailing lists) may potentially put a large number of
systems at risk of intrusion. It is invalid to assume that all
administrators reading a particular newsgroup have access to
operating system source code, or can even understand an advisory well
enough to take adequate steps.

First of all, any notification to either local or off-site personnel
must be explicit. This requires that any statement (be it an
electronic mail message, phone call, fax, beeper, or semaphore)
providing information about the incident be clear, concise, and fully
qualified. When you are notifying others that will help you handle
an event, a "smoke screen" will only divide the effort and create
confusion. If a division of labor is suggested, it is helpful to
provide information to each participant about what is being
accomplished in other efforts. This will not only reduce duplication
of effort, but allow people working on parts of the problem to know
where to obtain information relevant to their part of the incident.

Another important consideration when communicating about the incident
is to be factual. Attempting to hide aspects of the incident by
providing false or incomplete information may not only prevent a
successful resolution to the incident, but may even worsen the
situation.

The choice of language used when notifying people about the incident
can have a profound effect on the way that information is received.
When you use emotional or inflammatory terms, you raise the potential
for damage and negative outcomes of the incident. It is important to
remain calm both in written and spoken communications.

Another consideration is that not all people speak the same language.
Due to this fact, misunderstandings and delay may arise, especially
if it is a multi-national incident. Other international concerns
include differing legal implications of a security incident and
cultural differences. However, cultural differences do not only
exist between countries. They even exist within countries, between
different social or user groups. For example, an administrator of a
university system might be very relaxed about attempts to connect to
the system via telnet, but the administrator of a military system is
likely to consider the same action as a possible attack.

Another issue associated with the choice of language is the
notification of non-technical or off-site personnel. It is important
to accurately describe the incident without generating undue alarm or
confusion. While it is more difficult to describe the incident to a
non-technical audience, it is often more important. A non-technical
description may be required for upper-level management, the press, or
law enforcement liaisons. The importance of these communications
cannot be underestimated and may make the difference between
resolving the incident properly and escalating to some higher level
of damage.

If an incident response team becomes involved, it might be necessary
to fill out a template for the information exchange. Although this
may seem to be an additional burden and adds a certain delay, it
helps the team to act on this minimum set of information. The
response team may be able to respond to aspects of the incident of
which the local administrator is unaware. If information is given out
to someone else, the following minimum information should be
provided:

(1) timezone of logs, ... in GMT or local time
(2) information about the remote system, including host names,
     IP addresses and (perhaps) user IDs
(3) all log entries relevant for the remote site
(4) type of incident (what happened, why should you care)

If local information (i.e., local user IDs) is included in the log
entries, it will be necessary to sanitize the entries beforehand to
avoid privacy issues. In general, all information which might assist
a remote site in resolving an incident should be given out, unless
local policies prohibit this.

5.4.2 Protecting Evidence and Activity Logs

When you respond to an incident, document all details related to the
incident. This will provide valuable information to yourself and
others as you try to unravel the course of events. Documenting all
details will ultimately save you time. If you don't document every
relevant phone call, for example, you are likely to forget a
significant portion of information you obtain, requiring you to
contact the source of information again. At the same time, recording
details will provide evidence for prosecution efforts, providing the
case moves in that direction. Documenting an incident will also help
you perform a final assessment of damage (something your management,
as well as law enforcement officers, will want to know), and will
provide the basis for later phases of the handling process:
eradication, recovery, and follow-up "lessons learned."

During the initial stages of an incident, it is often infeasible to
determine whether prosecution is viable, so you should document as if
you are gathering evidence for a court case. At a minimum, you
should record:

(1) all system events (audit records)
(2) all actions you take (time tagged)
(3) all external conversations (including the person with whom
     you talked, the date and time, and the content of the
     conversation)

The most straightforward way to maintain documentation is keeping a
log book. This allows you to go to a centralized, chronological
source of information when you need it, instead of requiring you to
page through individual sheets of paper. Much of this information is
potential evidence in a court of law. Thus, when a legal follow-up
is a possibility, one should follow the prepared procedures and avoid
jeopardizing the legal follow-up by improper handling of possible
evidence. If appropriate, the following steps may be taken.

(1) Regularly (e.g., every day) turn in photocopied, signed
     copies of your logbook (as well as media you use to record
     system events) to a document custodian.
(2) The custodian should store these copied pages in a secure
     place (e.g., a safe).
(3) When you submit information for storage, you should
     receive a signed, dated receipt from the document
     custodian.

Failure to observe these procedures can result in invalidation of any
evidence you obtain in a court of law.

5.4.3 Containment

The purpose of containment is to limit the extent of an attack. An
essential part of containment is decision making (e.g., determining
whether to shut a system down, disconnect from a network, monitor
system or network activity, set traps, disable functions such as
remote file transfer, etc.).

Sometimes this decision is trivial; shut the system down if the
information is classified, sensitive, or proprietary. Bear in mind
that removing all access while an incident is in progress obviously
notifies all users, including the alleged problem users, that the
administrators are aware of a problem; this may have a deleterious

effect on an investigation. In some cases, it is prudent to remove
all access or functionality as soon as possible, then restore normal
operation in limited stages. In other cases, it is worthwhile to
risk some damage to the system if keeping the system up might enable
you to identify an intruder.

This stage should involve carrying out predetermined procedures.
Your organization or site should, for example, define acceptable
risks in dealing with an incident, and should prescribe specific
actions and strategies accordingly. This is especially important
when a quick decision is necessary and it is not possible to first
contact all involved parties to discuss the decision. In the absence
of predefined procedures, the person in charge of the incident will
often not have the power to make difficult management decisions (like
to lose the results of a costly experiment by shutting down a
system). A final activity that should occur during this stage of
incident handling is the notification of appropriate authorities.

5.4.4 Eradication

Once the incident has been contained, it is time to eradicate the
cause. But before eradicating the cause, great care should be taken
to collect all necessary information about the compromised system(s)
and the cause of the incident as they will likely be lost when
cleaning up the system.

Software may be available to help you in the eradication process,
such as anti-virus software. If any bogus files have been created,
archive them before deleting them. In the case of virus infections,
it is important to clean and reformat any media containing infected
files. Finally, ensure that all backups are clean. Many systems
infected with viruses become periodically re-infected simply because
people do not systematically eradicate the virus from backups. After
eradication, a new backup should be taken.

Removing all vulnerabilities once an incident has occurred is
difficult. The key to removing vulnerabilities is knowledge and
understanding of the breach.

It may be necessary to go back to the original distribution media and
re-customize the system. To facilitate this worst case scenario, a
record of the original system setup and each customization change
should be maintained. In the case of a network-based attack, it is
important to install patches for each operating system vulnerability
which was exploited.

As discussed in section 5.4.2, a security log can be most valuable
during this phase of removing vulnerabilities. The logs showing how
the incident was discovered and contained can be used later to help
determine how extensive the damage was from a given incident. The
steps taken can be used in the future to make sure the problem does
not resurface. Ideally, one should automate and regularly apply the
same test as was used to detect the security incident.

If a particular vulnerability is isolated as having been exploited,
the next step is to find a mechanism to protect your system. The
security mailing lists and bulletins would be a good place to search
for this information, and you can get advice from incident response
teams.

5.4.5 Recovery

Once the cause of an incident has been eradicated, the recovery phase
defines the next stage of action. The goal of recovery is to return
the system to normal. In general, bringing up services in the order
of demand to allow a minimum of user inconvenience is the best
practice. Understand that the proper recovery procedures for the
system are extremely important and should be specific to the site.

5.4.6 Follow-Up

Once you believe that a system has been restored to a "safe" state,
it is still possible that holes, and even traps, could be lurking in
the system. One of the most important stages of responding to
incidents is also the most often omitted, the follow-up stage. In
the follow-up stage, the system should be monitored for items that
may have been missed during the cleanup stage. It would be prudent
to utilize some of the tools mentioned in chapter 7 as a start.
Remember, these tools don't replace continual system monitoring and
good systems administration practices.

The most important element of the follow-up stage is performing a
postmortem analysis. Exactly what happened, and at what times? How
well did the staff involved with the incident perform? What kind of
information did the staff need quickly, and how could they have
gotten that information as soon as possible? What would the staff do
differently next time?

After an incident, it is prudent to write a report describing the
exact sequence of events: the method of discovery, correction
procedure, monitoring procedure, and a summary of lesson learned.
This will aid in the clear understanding of the problem. Creating a
formal chronology of events (including time stamps) is also important
for legal reasons.

A follow-up report is valuable for many reasons. It provides a
reference to be used in case of other similar incidents. It is also
important to, as quickly as possible obtain a monetary estimate of
the amount of damage the incident caused. This estimate should
include costs associated with any loss of software and files
(especially the value of proprietary data that may have been
disclosed), hardware damage, and manpower costs to restore altered
files, reconfigure affected systems, and so forth. This estimate may
become the basis for subsequent prosecution activity. The report can
also help justify an organization's computer security effort to
management.

5.5 Aftermath of an Incident

   In the wake of an incident, several actions should take place. These
   actions can be summarized as follows:

   (1) An inventory should be taken of the systems' assets,
        (i.e., a careful examination should determine how the
        system was affected by the incident).

   (2) The lessons learned as a result of the incident
        should be included in revised security plan to
        prevent the incident from re-occurring.

   (3) A new risk analysis should be developed in light of the
        incident.

   (4) An investigation and prosecution of the individuals
        who caused the incident should commence, if it is
        deemed desirable.

   If an incident is based on poor policy, and unless the policy is
   changed, then one is doomed to repeat the past. Once a site has
   recovered from and incident, site policy and procedures should be
   reviewed to encompass changes to prevent similar incidents. Even
   without an incident, it would be prudent to review policies and
   procedures on a regular basis. Reviews are imperative due to today's
   changing computing environments.

   The whole purpose of this post mortem process is to improve all
   security measures to protect the site against future attacks. As a
   result of an incident, a site or organization should gain practical
   knowledge from the experience. A concrete goal of the post mortem is
   to develop new proactive methods. Another important facet of the
   aftermath may be end user and administrator education to prevent a
   reoccurrence of the security problem.

5.6 Responsibilities

5.6.1 Not Crossing the Line

   It is one thing to protect one's own network, but quite another to
   assume that one should protect other networks. During the handling
   of an incident, certain system vulnerabilities of one's own systems
   and the systems of others become apparent. It is quite easy and may
   even be tempting to pursue the intruders in order to track them.
   Keep in mind that at a certain point it is possible to "cross the
   line," and, with the best of intentions, become no better than the
   intruder.

   The best rule when it comes to propriety is to not use any facility
   of remote sites which is not public. This clearly excludes any entry
   onto a system (such as a remote shell or login session) which is not
   expressly permitted. This may be very tempting; after a breach of
   security is detected, a system administrator may have the means to
   "follow it up," to ascertain what damage is being done to the remote
   site. Don't do it! Instead, attempt to reach the appropriate point
   of contact for the affected site.

5.6.2 Good Internet Citizenship

   During a security incident there are two choices one can make.
   First, a site can choose to watch the intruder in the hopes of
   catching him; or, the site can go about cleaning up after the
   incident and shut the intruder out of the systems. This is a
   decision that must be made very thoughtfully, as there may be legal
   liabilities if you choose to leave your site open, knowing that an
   intruder is using your site as a launching pad to reach out to other
   sites. Being a good Internet citizen means that you should try to
   alert other sites that may have been impacted by the intruder. These
   affected sites may be readily apparent after a thorough review of
   your log files.

5.6.3 Administrative Response to Incidents

   When a security incident involves a user, the site's security policy
   should describe what action is to be taken. The transgression should
   be taken seriously, but it is very important to be sure of the role
   the user played. Was the user naive? Could there be a mistake in
   attributing the security breach to the user? Applying administrative
   action that assumes the user intentionally caused the incident may

not be appropriate for a user who simply made a mistake. It may be
appropriate to include sanctions more suitable for such a situation
in your policies (e.g., education or reprimand of a user) in addition
to more stern measures for intentional acts of intrusion and system
misuse.

6. Ongoing Activities

   At this point in time, your site has hopefully developed a complete
   security policy and has developed procedures to assist in the
   configuration and management of your technology in support of those
   policies. How nice it would be if you could sit back and relax at
   this point and know that you were finished with the job of security.
   Unfortunately, that isn't possible. Your systems and networks are
   not a static environment, so you will need to review policies and
   procedures on a regular basis. There are a number of steps you can
   take to help you keep up with the changes around you so that you can
   initiate corresponding actions to address those changes. The
   following is a starter set and you may add others as appropriate for
   your site.

   (1) Subscribe to advisories that are issued by various security incident
        response teams, like those of the CERT Coordination Center, and
        update your systems against those threats that apply to your site's
        technology.

   (2) Monitor security patches that are produced by the vendors of your
        equipment, and obtain and install all that apply.

   (3) Actively watch the configurations of your systems to identify any
        changes that may have occurred, and investigate all anomalies.

   (4) Review all security policies and procedures annually (at a minimum).

   (5) Read relevant mailing lists and USENET newsgroups to keep up to
        date with the latest information being shared by fellow
        administrators.

   (6) Regularly check for compliance with policies and procedures. This
        audit should be performed by someone other than the people who
        define or implement the policies and procedures.

7. Tools and Locations

   This chapter provides a brief list of publicly available security
   technology which can be downloaded from the Internet. Many of the
   items described below will undoubtedly be surpassed or made obsolete
   before this document is published.

   Some of the tools listed are applications such as end user programs
   (clients) and their supporting system infrastructure (servers).
   Others are tools that a general user will never see or need to use,
   but may be used by applications, or by administrators to troubleshoot
   security problems or to guard against intruders.

   A sad fact is that there are very few security conscious applications
   currently available. Primarily, this is caused by the need for a
   security infrastructure which must first be put into place for most
   applications to operate securely. There is considerable effort
   currently taking place to build this infrastructure so that
   applications can take advantage of secure communications.

   Most of the tools and applications described below can be found in
   one of the following archive sites:

   (1)   CERT Coordination Center
         ftp://info.cert.org:/pub/tools
   (2)   DFN-CERT
         ftp://ftp.cert.dfn.de/pub/tools/
   (3)   Computer Operations, Audit, and Security Tools (COAST)
         coast.cs.purdue.edu:/pub/tools

   It is important to note that many sites, including CERT and COAST are
   mirrored throughout the Internet. Be careful to use a "well known"
   mirror site to retrieve software, and to use verification tools (md5
   checksums, etc.) to validate that software. A clever cracker might
   advertise security software that has intentionally been designed to
   provide access to data or systems.

Tools

   COPS
   DES
   Drawbridge
   identd (not really a security tool)
   ISS
   Kerberos
   logdaemon
   lsof
   MD5
   PEM
   PGP
   rpcbind/portmapper replacement
   SATAN
   sfingerd
   S/KEY
   smrsh

```
ssh
swatch
TCP-Wrapper
tiger
Tripwire*
TROJAN.PL
```

8. Mailing Lists and Other Resources

    It would be impossible to list all of the mail-lists and other
    resources dealing with site security. However, these are some "jump-
    points" from which the reader can begin. All of these references are
    for the "INTERNET" constituency. More specific (vendor and
    geographical) resources can be found through these references.

    Mailing Lists

    (1)   CERT(TM) Advisory
           Send mail to: cert-advisory-request@cert.org
           Message Body: subscribe cert <FIRST NAME> <LAST NAME>

           A CERT advisory provides information on how to obtain a patch or
           details of a workaround for a known computer security problem.
           The CERT Coordination Center works with vendors to produce a
           workaround or a patch for a problem, and does not publish
           vulnerability information until a workaround or a patch is
           available. A CERT advisory may also be a warning to our
           constituency about ongoing attacks (e.g.,
           "CA-91:18.Active.Internet.tftp.Attacks").

           CERT advisories are also published on the USENET newsgroup:
                       comp.security.announce

           CERT advisory archives are available via anonymous FTP from
           info.cert.org in the /pub/cert_advisories directory.

    (2)   VIRUS-L List
           Send mail to: listserv%lehiibm1.bitnet@mitvma.mit.edu
           Message Body: subscribe virus-L FIRSTNAME LASTNAME

           VIRUS-L is a moderated mailing list with a focus
           on computer virus issues. For more information,
           including a copy of the posting guidelines, see
           the file "virus-l.README", available by anonymous
           FTP from cs.ucr.edu.

   (3) Internet Firewalls
        Send mail to: majordomo@greatcircle.com
        Message Body: subscribe firewalls user@host

        The Firewalls mailing list is a discussion forum for
        firewall administrators and implementors.

   USENET newsgroups

   (1) comp.security.announce
        The comp.security.announce newsgroup is moderated
        and is used solely for the distribution of CERT
        advisories.

   (2) comp.security.misc
        The comp.security.misc is a forum for the
        discussion of computer security, especially as it
        relates to the UNIX(r) Operating System.

   (3) alt.security
        The alt.security newsgroup is also a forum for the
        discussion of computer security, as well as other
        issues such as car locks and alarm systems.

   (4) comp.virus
        The comp.virus newsgroup is a moderated newsgroup
        with a focus on computer virus issues. For more
        information, including a copy of the posting
        guidelines, see the file "virus-l.README",
        available via anonymous FTP on info.cert.org
        in the /pub/virus-l directory.

   (5) comp.risks
        The comp.risks newsgroup is a moderated forum on
        the risks to the public in computers and related
        systems.

   World-Wide Web Pages

   (1) http://www.first.org/

        Computer Security Resource Clearinghouse. The main focus is on
        crisis response information; information on computer
        security-related threats, vulnerabilities, and solutions. At the
        same time, the Clearinghouse strives to be a general index to
        computer security information on a broad variety of subjects,
        including general risks, privacy, legal issues, viruses,
        assurance, policy, and training.

(2) http://www.telstra.com.au/info/security.html

This Reference Index contains a list of links to information
sources on Network and Computer Security. There is no implied
fitness to the Tools, Techniques and Documents contained within this
archive. Many if not all of these items work well, but we do
not guarantee that this will be so. This information is for the
education and legitimate use of computer security techniques only.

(3) http://www.alw.nih.gov/Security/security.html

This page features general information about computer security.
Information is organized by source and each section is organized
by topic. Recent modifications are noted in What's New page.

(4) http://csrc.ncsl.nist.gov

This archive at the National Institute of Standards and Technology's
Computer Security Resource Clearinghouse page contains a number of
announcements, programs, and documents related to computer security.

* CERT and Tripwire are registered in the U.S. Patent and Trademark Offic
e

9. References

The following references may not be available in all countries.

[Appelman, et. al., 1995] Appelman, Heller, Ehrman, White, and
McAuliffe, "The Law and The Internet", USENIX 1995 Technical
Conference on UNIX and Advanced Computing, New Orleans, LA, January
16-20, 1995.

[ABA, 1989] American Bar Association, Section of Science and
Technology, "Guide to the Prosecution of Telecommunication Fraud by
the Use of Computer Crime Statutes", American Bar Association, 1989.

[Aucoin, 1989] R. Aucoin, "Computer Viruses: Checklist for Recovery",
Computers in Libraries, Vol. 9, No. 2, Pg. 4, February 1989.

[Barrett, 1996] D. Barrett, "Bandits on the Information
Superhighway", O'Reilly & Associates, Sebastopol, CA, 1996.

[Bates, 1992] R. Bates, "Disaster Recovery Planning: Networks,
Telecommunications and Data Communications", McGraw-Hill, 1992.

[Bellovin, 1989] S. Bellovin, "Security Problems in the TCP/IP
Protocol Suite", Computer Communication Review, Vol 19, 2, pp. 32-48,
April 1989.

[Bellovin, 1990] S. Bellovin, and M. Merritt, "Limitations of the
Kerberos Authentication System", Computer Communications Review,
October 1990.

[Bellovin, 1992] S. Bellovin, "There Be Dragon", USENIX: Proceedings
of the Third Usenix Security Symposium, Baltimore, MD. September,
1992.

[Bender, 1894] D. Bender, "Computer Law: Evidence and Procedure", M.
Bender, New York, NY, 1978-present.

[Bloombecker, 1990] B. Bloombecker, "Spectacular Computer Crimes",
Dow Jones- Irwin, Homewood, IL. 1990.

[Brand, 1990] R. Brand, "Coping with the Threat of Computer Security
Incidents: A Primer from Prevention through Recovery", R. Brand, 8
June 1990.

[Brock, 1989] J. Brock, "November 1988 Internet Computer Virus and
the Vulnerability of National Telecommunications Networks to Computer
Viruses", GAO/T-IMTEC-89-10, Washington, DC, 20 July 1989.

[BS 7799] British Standard, BS Tech Cttee BSFD/12, Info. Sec. Mgmt,
"BS 7799 : 1995 Code of Practice for Information Security
Management", British Standards Institution, London, 54, Effective 15
February 1995.

[Caelli, 1988] W. Caelli, Editor, "Computer Security in the Age of
Information", Proceedings of the Fifth IFIP International Conference
on Computer Security, IFIP/Sec '88.

[Carroll, 1987] J. Carroll, "Computer Security", 2nd Edition,
Butterworth Publishers, Stoneham, MA, 1987.

[Cavazos and Morin, 1995] E. Cavazos and G. Morin, "Cyber-Space and
The Law", MIT Press, Cambridge, MA, 1995.

[CCH, 1989] Commerce Clearing House, "Guide to Computer Law",
(Topical Law Reports), Chicago, IL., 1989.

[Chapman, 1992] B. Chapman, "Network(In) Security Through IP Packet
Filtering", USENIX: Proceedings of the Third UNIX Security Symposium,
Baltimore, MD, September 1992.

[Chapman and Zwicky, 1995] B. Chapman and E. Zwicky, "Building
Internet Firewalls", O'Reilly and Associates, Sebastopol, CA, 1995.

[Cheswick, 1990] B. Cheswick, "The Design of a Secure Internet Gateway", Proceedings of the Summer Usenix Conference, Anaheim, CA, June 1990.

[Cheswick1] W. Cheswick, "An Evening with Berferd In Which a Cracker is Lured, Endured, and Studied", AT&T Bell Laboratories.

[Cheswick and Bellovin, 1994] W. Cheswick and S. Bellovin, "Firewalls and Internet Security: Repelling the Wily Hacker", Addison-Wesley, Reading, MA, 1994.

[Conly, 1989] C. Conly, "Organizing for Computer Crime Investigation and Prosecution", U.S. Dept. of Justice, Office of Justice Programs, Under Contract Number OJP-86-C-002, National Institute of Justice, Washington, DC, July 1989.

[Cooper, 1989] J. Cooper, "Computer and Communications Security: Strategies for the 1990s", McGraw-Hill, 1989.

[CPSR, 1989] Computer Professionals for Social Responsibility, "CPSR Statement on the Computer Virus", CPSR, Communications of the ACM, Vol. 32, No. 6, Pg. 699, June 1989.

[CSC-STD-002-85, 1985] Department of Defense, "Password Management Guideline", CSC-STD-002-85, 12 April 1985, 31 pages.

[Curry, 1990] D. Curry, "Improving the Security of Your UNIX System", SRI International Report ITSTD-721-FR-90-21, April 1990.

[Curry, 1992] D. Curry, "UNIX System Security: A Guide for Users and Systems Administrators", Addision-Wesley, Reading, MA, 1992.

[DDN88] Defense Data Network, "BSD 4.2 and 4.3 Software Problem Resolution", DDN MGT Bulletin #43, DDN Network Information Center, 3 November 1988.

[DDN89] DCA DDN Defense Communications System, "DDN Security Bulletin 03", DDN Security Coordination Center, 17 October 1989.

[Denning, 1990] P. Denning, Editor, "Computers Under Attack: Intruders, Worms, and Viruses", ACM Press, 1990.

[Eichin and Rochlis, 1989] M. Eichin, and J. Rochlis, "With Microscope and Tweezers: An Analysis of the Internet Virus of November 1988", Massachusetts Institute of Technology, February 1989.

[Eisenberg, et. al., 89] T. Eisenberg, D. Gries, J. Hartmanis, D. Holcomb, M. Lynn, and T. Santoro, "The Computer Worm", Cornell University, 6 February 1989.

[Ermann, Willians, and Gutierrez, 1990] D. Ermann, M. Williams, and C. Gutierrez, Editors, "Computers, Ethics, and Society", Oxford University Press, NY, 1990. (376 pages, includes bibliographical references).

[Farmer and Spafford, 1990] D. Farmer and E. Spafford, "The COPS Security Checker System", Proceedings of the Summer 1990 USENIX Conference, Anaheim, CA, Pgs. 165-170, June 1990.

[Farrow, 1991] Rik Farrow, "UNIX Systems Security", Addison-Wesley, Reading, MA, 1991.

[Fenwick, 1985] W. Fenwick, Chair, "Computer Litigation, 1985: Trial Tactics and Techniques", Litigation Course Handbook Series No. 280, Prepared for distribution at the Computer Litigation, 1985: Trial Tactics and Techniques Program, February-March 1985.

[Fites 1989] M. Fites, P. Kratz, and A. Brebner, "Control and Security of Computer Information Systems", Computer Science Press, 1989.

[Fites, Johnson, and Kratz, 1992] Fites, Johnson, and Kratz, "The Computer Virus Crisis", Van Hostrand Reinhold, 2nd edition, 1992.

[Forester and Morrison, 1990] T. Forester, and P. Morrison, "Computer Ethics: Tales and Ethical Dilemmas in Computing", MIT Press, Cambridge, MA, 1990.

[Foster and Morrision, 1990] T. Forester, and P. Morrison, "Computer Ethics: Tales and Ethical Dilemmas in Computing", MIT Press, Cambridge, MA, 1990. (192 pages including index.)

[GAO/IMTEX-89-57, 1989] U.S. General Accounting Office, "Computer Security - Virus Highlights Need for Improved Internet Management", United States General Accounting Office, Washington, DC, 1989.

[Garfinkel and Spafford, 1991] S. Garfinkel, and E. Spafford, "Practical Unix Security", O'Reilly & Associates, ISBN 0-937175-72-2, May 1991.

[Garfinkel, 1995] S. Garfinkel, "PGP:Pretty Good Privacy", O'Reilly & Associates, Sebastopol, CA, 1996.

[Garfinkel and Spafford, 1996] S. Garfinkel and E. Spafford,
"Practical UNIX and Internet Security", O'Reilly & Associates,
Sebastopol, CA, 1996.

[Gemignani, 1989] M. Gemignani, "Viruses and Criminal Law",
Communications of the ACM, Vol. 32, No. 6, Pgs. 669-671, June 1989.

[Goodell, 1996] J. Goodell, "The Cyberthief and the Samurai: The True
Story of Kevin Mitnick-And The Man Who Hunted Him Down", Dell
Publishing, 1996.

[Gould, 1989] C. Gould, Editor, "The Information Web: Ethical and
Social Implications of Computer Networking", Westview Press, Boulder,
CO, 1989.

[Greenia, 1989] M. Greenia, "Computer Security Information
Sourcebook", Lexikon Services, Sacramento, CA, 1989.

[Hafner and Markoff, 1991] K. Hafner and J. Markoff, "Cyberpunk:
Outlaws and Hackers on the Computer Frontier", Touchstone, Simon &
Schuster, 1991.

[Hess, Safford, and Pooch] D. Hess, D. Safford, and U. Pooch, "A Unix
Network Protocol Security Study: Network Information Service", Texas
A&M University.

[Hoffman, 1990] L. Hoffman, "Rogue Programs: Viruses, Worms, and
Trojan Horses", Van Nostrand Reinhold, NY, 1990. (384 pages,
includes bibliographical references and index.)

[Howard, 1995] G. Howard, "Introduction to Internet Security: From
Basics to Beyond", Prima Publishing, Rocklin, CA, 1995.

[Huband and Shelton, 1986] F. Huband, and R. Shelton, Editors,
"Protection of Computer Systems and Software: New Approaches for
Combating Theft of Software and Unauthorized Intrusion", Papers
presented at a workshop sponsored by the National Science Foundation,
1986.

[Hughes, 1995] L. Hughes Jr., "Actually Useful Internet Security
Techniques", New Riders Publishing, Indianapolis, IN, 1995.

[IAB-RFC1087, 1989] Internet Activities Board, "Ethics and the
Internet", RFC 1087, IAB, January 1989. Also appears in the
Communications of the ACM, Vol. 32, No. 6, Pg. 710, June 1989.

[Icove, Seger, and VonStorch, 1995] D. Icove, K. Seger, and W. VonStorch, "Computer Crime: A Crimefighter's Handbook", O'Reilly & Associates, Sebastopol, CA, 1995.

[IVPC, 1996] IVPC, "International Virus Prevention Conference '96 Proceedings", NCSA, 1996.

[Johnson and Podesta] D. Johnson, and J. Podesta, "Formulating A Company Policy on Access to and Use and Disclosure of Electronic Mail on Company Computer Systems".

[Kane, 1994] P. Kane, "PC Security and Virus Protection Handbook: The Ongoing War Against Information Sabotage", M&T Books, 1994.

[Kaufman, Perlman, and Speciner, 1995] C. Kaufman, R. Perlman, and M. Speciner, "Network Security: PRIVATE Communication in a PUBLIC World", Prentice Hall, Englewood Cliffs, NJ, 1995.

[Kent, 1990] S. Kent, "E-Mail Privacy for the Internet: New Software and Strict Registration Procedures will be Implemented this Year", Business Communications Review, Vol. 20, No. 1, Pg. 55, 1 January 1990.

[Levy, 1984] S. Levy, "Hacker: Heroes of the Computer Revolution", Delta, 1984.

[Lewis, 1996] S. Lewis, "Disaster Recovery Yellow Pages", The Systems Audit Group, 1996.

[Littleman, 1996] J. Littleman, "The Fugitive Game: Online with Kevin Mitnick", Little, Brown, Boston, MA., 1996.

[Lu and Sundareshan, 1989] W. Lu and M. Sundareshan, "Secure Communication in Internet Environments: A Hierarchical Key Management Scheme for End-to-End Encryption", IEEE Transactions on Communications, Vol. 37, No. 10, Pg. 1014, 1 October 1989.

[Lu and Sundareshan, 1990] W. Lu and M. Sundareshan, "A Model for Multilevel Security in Computer Networks", IEEE Transactions on Software Engineering, Vol. 16, No. 6, Page 647, 1 June 1990.

[Martin and Schinzinger, 1989] M. Martin, and R. Schinzinger, "Ethics in Engineering", McGraw Hill, 2nd Edition, 1989.

[Merkle] R. Merkle, "A Fast Software One Way Hash Function", Journal of Cryptology, Vol. 3, No. 1.

[McEwen, 1989] J. McEwen, "Dedicated Computer Crime Units", Report
Contributors: D. Fester and H. Nugent, Prepared for the National
Institute of Justice, U.S. Department of Justice, by Institute for
Law and Justice, Inc., under contract number OJP-85-C-006,
Washington, DC, 1989.

[MIT, 1989] Massachusetts Institute of Technology, "Teaching Students
About Responsible Use of Computers", MIT, 1985-1986. Also reprinted
in the Communications of the ACM, Vol. 32, No. 6, Pg. 704, Athena
Project, MIT, June 1989.

[Mogel, 1989] Mogul, J., "Simple and Flexible Datagram Access
Controls for UNIX-based Gateways", Digital Western Research
Laboratory Research Report 89/4, March 1989.

[Muffett, 1992] A. Muffett, "Crack Version 4.1: A Sensible Password
Checker for Unix"

[NCSA1, 1995] NCSA, "NCSA Firewall Policy Guide", 1995.

[NCSA2, 1995] NCSA, "NCSA's Corporate Computer Virus Prevention
Policy Model", NCSA, 1995.

[NCSA, 1996] NCSA, "Firewalls & Internet Security Conference '96
Proceedings", 1996.

[NCSC-89-660-P, 1990] National Computer Security Center, "Guidelines
for Formal Verification Systems", Shipping list no.: 89-660-P, The
Center, Fort George G. Meade, MD, 1 April 1990.

[NCSC-89-254-P, 1988] National Computer Security Center, "Glossary of
Computer Security Terms", Shipping list no.: 89-254-P, The Center,
Fort George G. Meade, MD, 21 October 1988.

[NCSC-C1-001-89, 1989] Tinto, M., "Computer Viruses: Prevention,
Detection, and Treatment", National Computer Security Center C1
Technical Report C1-001-89, June 1989.

[NCSC Conference, 1989] National Computer Security Conference, "12th
National Computer Security Conference: Baltimore Convention Center,
Baltimore, MD, 10-13 October, 1989: Information Systems Security,
Solutions for Today - Concepts for Tomorrow", National Institute of
Standards and National Computer Security Center, 1989.

[NCSC-CSC-STD-003-85, 1985] National Computer Security Center,
"Guidance for Applying the Department of Defense Trusted Computer
System Evaluation Criteria in Specific Environments", CSC-STD-003-85,
NCSC, 25 June 1985.

[NCSC-STD-004-85, 1985] National Computer Security Center, "Technical
Rationale Behind CSC-STD-003-85: Computer Security Requirements",
CSC-STD-004-85, NCSC, 25 June 1985.

[NCSC-STD-005-85, 1985] National Computer Security Center, "Magnetic
Remanence Security Guideline", CSC-STD-005-85, NCSC, 15 November
1985.

[NCSC-TCSEC, 1985] National Computer Security Center, "Trusted
Computer System Evaluation Criteria", DoD 5200.28-STD, CSC-STD-001-
83, NCSC, December 1985.

[NCSC-TG-003, 1987] NCSC, "A Guide to Understanding DISCRETIONARY
ACCESS CONTROL in Trusted Systems", NCSC-TG-003, Version-1, 30
September 1987, 29 pages.

[NCSC-TG-001, 1988] NCSC, "A Guide to Understanding AUDIT in Trusted
Systems", NCSC-TG-001, Version-2, 1 June 1988, 25 pages.

[NCSC-TG-004, 1988] National Computer Security Center, "Glossary of
Computer Security Terms", NCSC-TG-004, NCSC, 21 October 1988.

[NCSC-TG-005, 1987] National Computer Security Center, "Trusted
Network Interpretation", NCSC-TG-005, NCSC, 31 July 1987.

[NCSC-TG-006, 1988] NCSC, "A Guide to Understanding CONFIGURATION
MANAGEMENT in Trusted Systems", NCSC-TG-006, Version-1, 28 March
1988, 31 pages.

[NCSC-TRUSIX, 1990] National Computer Security Center, "Trusted UNIX
Working Group (TRUSIX) rationale for selecting access control list
features for the UNIX system", Shipping list no.: 90-076-P, The
Center, Fort George G. Meade, MD, 1990.

[NRC, 1991] National Research Council, "Computers at Risk: Safe
Computing in the Information Age", National Academy Press, 1991.

[Nemeth, et. al, 1995] E. Nemeth, G. Snyder, S. Seebass, and T. Hein,
"UNIX Systems Administration Handbook", Prentice Hall PTR, Englewood
Cliffs, NJ, 2nd ed. 1995.

[NIST, 1989] National Institute of Standards and Technology,
"Computer Viruses and Related Threats: A Management Guide", NIST
Special Publication 500-166, August 1989.

[NSA] National Security Agency, "Information Systems Security
Products and Services Catalog", NSA, Quarterly Publication.

[NSF, 1988] National Science Foundation, "NSF Poses Code of
Networking Ethics", Communications of the ACM, Vol. 32, No. 6, Pg.
688, June 1989. Also appears in the minutes of the regular meeting
of the Division Advisory Panel for Networking and Communications
Research and Infrastructure, Dave Farber, Chair, November 29-30,
1988.

[NTISSAM, 1987] NTISS, "Advisory Memorandum on Office Automation
Security Guideline", NTISSAM COMPUSEC/1-87, 16 January 1987, 58
pages.

[OTA-CIT-310, 1987] United States Congress, Office of Technology
Assessment, "Defending Secrets, Sharing Data: New Locks and Keys for
Electronic Information", OTA-CIT-310, October 1987.

[OTA-TCT-606] Congress of the United States, Office of Technology
Assessment, "Information Security and Privacy in Network
Environments", OTA-TCT-606, September 1994.

[Palmer and Potter, 1989] I. Palmer, and G. Potter, "Computer
Security Risk Management", Van Nostrand Reinhold, NY, 1989.

[Parker, 1989] D. Parker, "Computer Crime: Criminal Justice Resource
Manual", U.S. Dept. of Justice, National Institute of Justice, Office
of Justice Programs, Under Contract Number OJP-86-C-002, Washington,
D.C., August 1989.

[Parker, Swope, and Baker, 1990] D. Parker, S. Swope, and B. Baker,
"Ethical Conflicts: Information and Computer Science, Technology and
Business", QED Information Sciences, Inc., Wellesley, MA. (245
pages).

[Pfleeger, 1989] C. Pfleeger, "Security in Computing", Prentice-Hall,
Englewood Cliffs, NJ, 1989.

[Quarterman, 1990] J. Quarterman, J., "The Matrix: Computer Networks
and Conferencing Systems Worldwide", Digital Press, Bedford, MA,
1990.

[Ranum1, 1992] M. Ranum, "An Internet Firewall", Proceedings of World
Conference on Systems Management and Security, 1992.

[Ranum2, 1992] M. Ranum, "A Network Firewall", Digital Equipment
Corporation Washington Open Systems Resource Center, June 12, 1992.

[Ranum, 1993] M. Ranum, "Thinking About Firewalls", 1993.

[Ranum and Avolio, 1994] M. Ranum and F. Avolio, "A Toolkit and Methods for Internet Firewalls", Trustest Information Systems, 1994.

[Reinhardt, 1992] R. Reinhardt, "An Architectural Overview of UNIX Network Security"

[Reinhardt, 1993] R. Reinhardt, "An Architectural Overview of UNIX Network Security", ARINC Research Corporation, February 18, 1993.

[Reynolds-RFC1135, 1989] The Helminthiasis of the Internet, RFC 1135, USC/Information Sciences Institute, Marina del Rey, CA, December 1989.

[Russell and Gangemi, 1991] D. Russell and G. Gangemi, "Computer Security Basics" O'Reilly & Associates, Sebastopol, CA, 1991.

[Schneier 1996] B. Schneier, "Applied Cryptography: Protocols, Algorithms, and Source Code in C", John Wiley & Sons, New York, second edition, 1996.

[Seeley, 1989] D. Seeley, "A Tour of the Worm", Proceedings of 1989 Winter USENIX Conference, Usenix Association, San Diego, CA, February 1989.

[Shaw, 1986] E. Shaw Jr., "Computer Fraud and Abuse Act of 1986", Congressional Record (3 June 1986), Washington, D.C., 3 June 1986.

[Shimomura, 1996] T. Shimomura with J. Markoff, "Takedown:The Pursuit and Capture of Kevin Mitnick, America's Most Wanted Computer Outlaw- by the Man Who Did It", Hyperion, 1996.

[Shirey, 1990] R. Shirey, "Defense Data Network Security Architecture", Computer Communication Review, Vol. 20, No. 2, Page 66, 1 April 1990.

[Slatalla and Quittner, 1995] M. Slatalla and J. Quittner, "Masters of Deception: The Gang that Ruled Cyberspace", Harper Collins Publishers, 1995.

[Smith, 1989] M. Smith, "Commonsense Computer Security: Your Practical Guide to Preventing Accidental and Deliberate Electronic Data Loss", McGraw-Hill, New York, NY, 1989.

[Smith, 1995] D. Smith, "Forming an Incident Response Team", Sixth Annual Computer Security Incident Handling Workshop, Boston, MA, July 25-29, 1995.

[Spafford, 1988] E. Spafford, "The Internet Worm Program: An
Analysis", Computer Communication Review, Vol. 19, No. 1, ACM SIGCOM,
January 1989. Also issued as Purdue CS Technical Report CSD-TR-823,
28 November 1988.

[Spafford, 1989] G. Spafford, "An Analysis of the Internet Worm",
Proceedings of the European Software Engineering Conference 1989,
Warwick England, September 1989. Proceedings published by Springer-
Verlag as: Lecture Notes in Computer Science #387. Also issued as
Purdue Technical Report #CSD-TR-933.

[Spafford, Keaphy, and Ferbrache, 1989] E. Spafford, K. Heaphy, and
D. Ferbrache, "Computer Viruses: Dealing with Electronic Vandalism
and Programmed Threats", ADAPSO, 1989. (109 pages.)

[Stallings1, 1995] W. Stallings, "Internet Security Handbook", IDG
Books, Foster City CA, 1995.

[Stallings2, 1995] W. Stallings, "Network and InterNetwork Security",
Prentice Hall, , 1995.

[Stallings3, 1995] W. Stallings, "Protect Your Privacy: A Guide for
PGP Users" PTR Prentice Hall, 1995.

[Stoll, 1988] C. Stoll, "Stalking the Wily Hacker", Communications of
the ACM, Vol. 31, No. 5, Pgs. 484-497, ACM, New York, NY, May 1988.

[Stoll, 1989] C. Stoll, "The Cuckoo's Egg", ISBN 00385-24946-2,
Doubleday, 1989.

[Treese and Wolman, 1993] G. Treese and A. Wolman, "X Through the
Firewall, and Other Applications Relays", Digital Equipment
Corporation, Cambridge Research Laboratory, CRL 93/10, May 3, 1993.

[Trible, 1986] P. Trible, "The Computer Fraud and Abuse Act of 1986",
U.S. Senate Committee on the Judiciary, 1986.

[Venema] W. Venema, "TCP WRAPPER: Network monitoring, access control,
and booby traps", Mathematics and Computing Science, Eindhoven
University of Technology, The Netherlands.

[USENIX, 1988] USENIX, "USENIX Proceedings: UNIX Security Workshop",
Portland, OR, August 29-30, 1988.

[USENIX, 1990] USENIX, "USENIX Proceedings: UNIX Security II
Workshop", Portland, OR, August 27-28, 1990.

[USENIX, 1992] USENIX, "USENIX Symposium Proceedings: UNIX Security III", Baltimore, MD, September 14-16, 1992.

[USENIX, 1993] USENIX, "USENIX Symposium Proceedings: UNIX Security IV", Santa Clara, CA, October 4-6, 1993.

[USENIX, 1995] USENIX, "The Fifth USENIX UNIX Security Symposium", Salt Lake City, UT, June 5-7, 1995.

[Wood, et.al., 1987] C. Wood, W. Banks, S. Guarro, A. Garcia, V. Hampel, and H. Sartorio, "Computer Security: A Comprehensive Controls Checklist", John Wiley and Sons, Interscience Publication, 1987.

[Wrobel, 1993] L. Wrobel, "Writing Disaster Recovery Plans for Telecommunications Networks and LANS", Artech House, 1993.

[Vallabhaneni, 1989] S. Vallabhaneni, "Auditing Computer Security: A Manual with Case Studies", Wiley, New York, NY, 1989.

Security Considerations

   This entire document discusses security issues.

Editor Information

   Barbara Y. Fraser
   Software Engineering Institute
   Carnegie Mellon University
   5000 Forbes Avenue
   Pittsburgh, PA 15213

   Phone: (412) 268-5010
   Fax: (412) 268-6989
   EMail: byf@cert.org