                        Site Security Handbook

Status of this Memo

   This handbook is the product of the Site Security Policy Handbook
   Working Group (SSPHWG), a combined effort of the Security Area and
   User Services Area of the Internet Engineering Task Force (IETF).
   This FYI RFC provides information for the Internet community. It
   does not specify an Internet standard. Distribution of this memo is
   unlimited.

Contributing Authors

   The following are the authors of the Site Security Handbook. Without
   their dedication, this handbook would not have been possible.

   Dave Curry (Purdue University), Sean Kirkpatrick (Unisys), Tom
   Longstaff (LLNL), Greg Hollingsworth (Johns Hopkins University),
   Jeffrey Carpenter (University of Pittsburgh), Barbara Fraser (CERT),
   Fred Ostapik (SRI NISC), Allen Sturtevant (LLNL), Dan Long (BBN), Jim
   Duncan (Pennsylvania State University), and Frank Byrum (DEC).

Editors' Note

   This FYI RFC is a first attempt at providing Internet users guidance
   on how to deal with security issues in the Internet. As such, this
   document is necessarily incomplete. There are some clear shortfalls;
   for example, this document focuses mostly on resources available in
   the United States. In the spirit of the Internet's "Request for
   Comments" series of notes, we encourage feedback from users of this
   handbook. In particular, those who utilize this document to craft
   their own policies and procedures.

   This handbook is meant to be a starting place for further research
   and should be viewed as a useful resource, but not the final
   authority. Different organizations and jurisdictions will have
   different resources and rules. Talk to your local organizations,
   consult an informed lawyer, or consult with local and national law
   enforcement. These groups can help fill in the gaps that this
   document cannot hope to cover.

Finally, we intend for this FYI RFC to grow and evolve. Please send
comments and suggestions to: ssphwg@cert.sei.cmu.edu.

Table of Contents

1. Introduction

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. This guide
   lists issues and factors that a site must consider when setting their
   own policies. It makes some recommendations and gives 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 the policies.

1.2 Audience

   The audience for this work are system administrators and decision
   makers (who are more traditionally called "administrators" or "middle
   management") at sites. 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 any 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,
   PC's or other devices that have access to the Internet. A site may
   be a end user of Internet services or a service provider such as a
   regional 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.

   The "Internet" is those set of networks and machines that use the
   TCP/IP protocol suite, connected through gateways, and sharing a
   common name and address spaces [1].

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

   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 IETF Security Policy Working Group (SPWG) is working on a set of
   recommended security policy guidelines for the Internet [23]. These
   guidelines may be adopted as policy by regional networks or owners of
   other resources. This handbook should be a useful tool to help sites
   implement those policies as desired or required. However, even
   implementing the proposed policies isn't enough to secure a site.
   The proposed Internet policies deal only with network access
   security. It says nothing about how sites should deal with local
   security issues.

1.5 Scope

   This document covers issues about what a computer security policy
   should contain, what kinds of procedures are need to enforce
   security, and some recommendations about how to deal with the
   problem. When developing a security policy, close attention should
   be made not only on the security needs and requirements of the local
   network, but also the security needs and requirements of the other
   interconnected networks.

This is not a cookbook for computer security. Each site has
different needs; the security needs of a corporation might well be
different than the security needs of an academic institution. Any
security plan has to conform to the needs and culture of the site.

This handbook does not cover details of how to do risk assessment,
contingency planning, or physical security. These things are
essential in setting and implementing effective security policy, but
this document leaves treatment of those issues to other documents.
We will try to provide some pointers in that direction.

This document also doesn't talk about how to design or implement
secure systems or programs.

1.6 Why Do We Need Security Policies and Procedures?

For most sites, the interest in computer security is proportional to
the perception of risk and threats.

The world of computers has changed dramatically over the past
twenty-five years. Twenty-five years ago, most computers were
centralized and managed by data centers. Computers were kept in
locked rooms and staffs of people made sure they were carefully
managed and physically secured. Links outside a site were unusual.
Computer security threats were rare, and were basically concerned
with insiders: authorized users misusing accounts, theft and
vandalism, and so forth. These threats were well understood and
dealt with using standard techniques: computers behind locked doors,
and accounting for all resources.

Computing in the 1990's is radically different. Many systems are in
private offices and labs, often managed by individuals or persons
employed outside a computer center. Many systems are connected into
the Internet, and from there around the world: the United States,
Europe, Asia, and Australia are all connected together.

Security threats are different today. The time honored advice says
"don't write your password down and put it in your desk" lest someone
find it. With world-wide Internet connections, someone could get
into your system from the other side of the world and steal your
password in the middle of the night when your building is locked up.
Viruses and worms can be passed from machine to machine. The
Internet allows the electronic equivalent of the thief who looks for
open windows and doors; now a person can check hundreds of machines
for vulnerabilities in a few hours.

System administrators and decision makers have to understand the
security threats that exist, what the risk and cost of a problem

   would be, and what kind of action they want to take (if any) to
   prevent and respond to security threats.

   As an illustration of some of the issues that need to be dealt with
   in security problems, consider the following scenarios (thanks to
   Russell Brand [2, BRAND] for these):

        - A system programmer gets a call reporting that a
          major underground cracker newsletter is being
          distributed from the administrative machine at his
          center to five thousand sites in the US and
          Western Europe.

          Eight weeks later, the authorities call to inform
          you the information in one of these newsletters
          was used to disable "911" in a major city for
          five hours.

        - A user calls in to report that he can't login to his
          account at 3 o'clock in the morning on a Saturday. The
          system staffer can't login either. After rebooting to
          single user mode, he finds that password file is empty.
          By Monday morning, your staff determines that a number
          of privileged file transfers took place between this
          machine and a local university.

          Tuesday morning a copy of the deleted password file is
          found on the university machine along with password
          files for a dozen other machines.

          A week later you find that your system initialization
          files had been altered in a hostile fashion.

        - You receive a call saying that a breakin to a government
          lab occurred from one of your center's machines. You
          are requested to provide accounting files to help
          trackdown the attacker.

          A week later you are given a list of machines at your
          site that have been broken into.

         - A reporter calls up asking about the breakin at your
           center. You haven't heard of any such breakin.

          Three days later, you learn that there was a breakin.
          The center director had his wife's name as a password.

      - A change in system binaries is detected.

        The day that it is corrected, they again are changed.
        This repeats itself for some weeks.

      - If an intruder is found on your system, should you
        leave the system open to monitor the situation or should
        you close down the holes and open them up again later?

      - If an intruder is using your site, should you call law
        enforcement? Who makes that decision? If law enforcement asks
        you to leave your site open, who makes that decision?

      - What steps should be taken if another site calls you and says
        they see activity coming from an account on your system? What
        if the account is owned by a local manager?

1.7 Basic Approach

   Setting security policies and procedures really means developing a
   plan for how to deal with computer security. One way to approach
   this task is suggested by Fites, et. al. [3, FITES]:

      - Look at what you are trying to protect.
      - Look at what you need to protect it from.
      - Determine how likely the threats are.
      - Implement measures which will protect your assets in a
         cost-effective manner.
      - Review the process continuously, and improve things every time
         a weakness is found.

   This handbook will concentrate mostly on the last two steps, but the
   first three are critically important to making effective decisions
   about security. One old truism in security is that the cost of
   protecting yourself against a threat should be less than the cost
   recovering if the threat were to strike you. Without reasonable
   knowledge of what you are protecting and what the likely threats are,
   following this rule could be difficult.

1.8 Organization of this Document

   This document is organized into seven parts in addition to this
   introduction.

   The basic form of each section is to discuss issues that a site might
   want to consider in creating a computer security policy and setting
   procedures to implement that policy. In some cases, possible options
   are discussed along with the some of the ramifications of those

choices. As far as possible, this document tries not to dictate the
choices a site should make, since these depend on local
circumstances. Some of the issues brought up may not apply to all
sites. Nonetheless, all sites should at least consider the issues
brought up here to ensure that they do not miss some important area.

The overall flow of the document is to discuss policy issues followed
by the issues that come up in creating procedures to implement the
policies.

Section 2 discusses setting official site policies for access to
computing resources. It also goes into the issue of what happens
when the policy is violated. The policies will drive the procedures
that need to be created, so decision makers will need to make choices
about policies before many of the procedural issues in following
sections can be dealt with. A key part of creating policies is doing
some kind of risk assessment to decide what really needs to be
protected and the level of resources that should be applied to
protect them.

Once policies are in place, procedures to prevent future security
problems should be established. Section 3 defines and suggests
actions to take when unauthorized activity is suspected. Resources
to prevent secruity breaches are also discussed.

Section 4 discusses types of procedures to prevent security problems.
Prevention is a key to security; as an example, the Computer
Emergency Response Team/Coordination Center (CERT/CC) at Carnegie-
Mellon University (CMU) estimates that 80% or more of the problems
they see have to do with poorly chosen passwords.

Section 5 discusses incident handling: what kinds of issues does a
site face when someone violates the security policy. Many decisions
will have to made on the spot as the incident occurs, but many of the
options and issues can be discussed in advance. At very least,
responsibilities and methods of communication can be established
before an incident. Again, the choices here are influenced by the
policies discussed in section 2.

Section 6 deals with what happens after a security violation has been
dealt with. Security planning is an on-going cycle; just after an
incident has occurred is an excellent opportunity to improve policies
and procedures.

The rest of the document provides references and an annotated
bibliography.

2. Establishing Official Site Policy on Computer Security

2.1 Brief Overview

   2.1.1 Organization Issues

   The goal in developing an official site policy on computer
   security is to define the organization's expectations of proper
   computer and network use and to define procedures to prevent and
   respond to security incidents. In order to do this, aspects of
   the particular organization must be considered.

   First, the goals and direction of the organization should be
   considered. For example, a military base may have very different
   security concerns from a those of a university.

   Second, the site security policy developed must conform to
   existing policies, rules, regulations and laws that the
   organization is subject to. Therefore it will be necessary to
   identify these and take them into consideration while developing
   the policy.

   Third, unless the local network is completely isolated and
   standalone, it is necessary to consider security implications in a
   more global context. The policy should address the issues when
   local security problems develop as a result of a remote site as
   well as when problems occur on remote systems as a result of a
   local host or user.

   2.1.2 Who Makes the Policy?

   Policy creation must be a joint effort by technical personnel, who
   understand the full ramifications of the proposed policy and the
   implementation of the policy, and by decision makers who have the
   power to enforce the policy. A policy which is neither
   implementable nor enforceable is useless.

   Since a computer security policy can affect everyone in an
   organization, it is worth taking some care to make sure you have
   the right level of authority in on the policy decisions. Though a
   particular group (such as a campus information services group) may
   have responsibility for enforcing a policy, an even higher group
   may have to support and approve the policy.

   2.1.3 Who is Involved?

   Establishing a site policy has the potential for involving every
   computer user at the site in a variety of ways. Computer users

may be responsible for personal password administration. Systems
managers are obligated to fix security holes and to oversee the
system.

It is critical to get the right set of people involved at the
start of the process. There may already be groups concerned with
security who would consider a computer security policy to be their
area. Some of the types of groups that might be involved include
auditing/control, organizations that deal with physical security,
campus information systems groups, and so forth. Asking these
types of groups to "buy in" from the start can help facilitate the
acceptance of the policy.

2.1.4 Responsibilities

A key element of a computer security policy is making sure
everyone knows their own responsibility for maintaining security.
A computer security policy cannot anticipate all possibilities;
however, it can ensure that each kind of problem does have someone
assigned to deal with it.

There may be levels of responsibility associated with a policy on
computer security. At one level, each user of a computing
resource may have a responsibility to protect his account. A user
who allows his account to be compromised increases the chances of
compromising other accounts or resources.

System managers may form another responsibility level: they must
help to ensure the security of the computer system. Network
managers may reside at yet another level.

2.2 Risk Assessment

2.2.1 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. Is is the
process of examining all of your risks, and ranking those risks by
level of severity. This process involves making cost-effective

decisions on what you want to protect. The old security adage
says that you should not spend more to protect something than it
is actually worth.

A full treatment of risk analysis is outside the scope of this
document. [3, FITES] and [16, PFLEEGER] 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.

  2.2.2 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 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 [16, PFLEEGER,
page 459]; 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, people needed to run systems.

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

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

   2.2.3 Identifying the Threats

      Once the assets requiring protection are identified, it is
      necessary to identify threats to those assests. 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 sections describe a few of the possible threats.

      2.2.3.1 Unauthorized Access

         A common threat that concerns many sites is unauthorized access
         to computing facilities. Unauthorized access takes many forms.
         One means of unauthorized access is the use of another user's
         account to gain access to a system. The use of any computer
         resource without prior permission may be considered
         unauthorized access to computing facilities.

         The seriousness of an unauthorized access will vary from site
         to site. For some sites, the mere act of granting access to an
         unauthorized user may cause irreparable harm by negative media
         coverage. For other sites, an unauthorized access opens the
         door to other security threats. In addition, some sites may be
         more frequent targets than others; hence the risk from
         unauthorized access will vary from site to site. The Computer
         Emergency Response Team (CERT - see section 3.9.7.3.1) has
         observed that well-known universities, government sites, and
         military sites seem to attract more intruders.

      2.2.3.2 Disclosure of Information

         Another common threat is disclosure of information. Determine
         the value or sensitivity of the information stored on your
         computers. Disclosure of a password file might allow for
         future unauthorized accesses. A glimpse of a proposal may give
         a competitor an unfair advantage. A technical paper may
         contain years of valuable research.

      2.2.3.3 Denial of Service

         Computers and networks provide valuable services to their
         users. Many people rely on these services in order to perform
         their jobs efficiently. When these services are not available
         when called upon, a loss in productivity results.

         Denial of service comes in many forms and might affect users in
         a number of ways. A network may be rendered unusable by a

rogue packet, jamming, or by a disabled network component. A
virus might slow down or cripple a computer system. Each site
should determine which services are essential, and for each of
these services determine the affect to the site if that service
were to become disabled.

2.3 Policy Issues

   There are a number of issues that must be addressed when developing a
   security policy. These are:

      1. Who is allowed to use the resources?
      2. What is the proper use of the resources?
      3. Who is authorized to grant access and approve usage?
      4. Who may have system administration privileges?
      5. What are the user's rights and responsibilities?
      6. What are the rights and responsibilities of the
          system administrator vs. those of the user?
      7. What do you do with sensitive information?

   These issues will be discussed below. In addition you may wish to
   include a section in your policy concerning ethical use of computing
   resources. Parker, Swope and Baker [17, PARKER90] and Forester and
   Morrison [18, FORESTER] are two useful references that address
   ethical issues.

   2.3.1 Who is Allowed to use the Resources?

      One step you must take in developing your security policy is
      defining who is allowed to use your system and services. The
      policy should explicitly state who is authorized to use what
      resources.

   2.3.2 What is the Proper Use of the Resources?

      After determining who is allowed access to system resources it is
      necessary to provide guidelines for the acceptable use of the
      resources. You may have different guidelines for different types
      of users (i.e., students, faculty, external users). The policy
      should state what is acceptable use as well as unacceptable use.
      It should also include types of use that may be restricted.

      Define limits to access and authority. You will need to consider
      the level of access various users will have and what resources
      will be available or restricted to various groups of people.

      Your acceptable use policy should clearly state that individual
      users are responsible for their actions. Their responsibility

exists regardless of the security mechanisms that are in place.
It should be clearly stated that breaking into accounts or
bypassing security is not permitted.

The following points should be covered when developing an
acceptable use policy:

     o Is breaking into accounts permitted?
     o Is cracking passwords permitted?
     o Is disrupting service permitted?
     o Should users assume that a file being world-readable
       grants them the authorization to read it?
     o Should users be permitted to modify files that are
       not their own even if they happen to have write
       permission?
     o Should users share accounts?

The answer to most of these questions will be "no".

You may wish to incorporate a statement in your policies
concerning copyrighted and licensed software. Licensing
agreements with vendors may require some sort of effort on your
part to ensure that the license is not violated. In addition, you
may wish to inform users that the copying of copyrighted software
may be a violation of the copyright laws, and is not permitted.

Specifically concerning copyrighted and/or licensed software, you
may wish to include the following information:

     o Copyrighted and licensed software may not be duplicated
       unless it is explicitly stated that you may do so.
     o Methods of conveying information on the
       copyright/licensed status of software.
     o When in doubt, DON'T COPY.

Your acceptable use policy is very important. A policy which does
not clearly state what is not permitted may leave you unable to
prove that a user violated policy.

There are exception cases like tiger teams and users or
administrators wishing for "licenses to hack" -- you may face the
situation where users will want to "hack" on your services for
security research purposes. You should develop a policy that will
determine whether you will permit this type of research on your
services and if so, what your guidelines for such research will
be.

Points you may wish to cover in this area:

o Whether it is permitted at all.
o What type of activity is permitted: breaking in, releasing
  worms, releasing viruses, etc..
o What type of controls must be in place to ensure that it
  does not get out of control (e.g., separate a segment of
  your network for these tests).
o How you will protect other users from being victims of
  these activities, including external users and networks.
o The process for obtaining permission to conduct these
  tests.

In cases where you do permit these activities, you should isolate
the portions of the network that are being tested from your main
network. Worms and viruses should never be released on a live
network.

You may also wish to employ, contract, or otherwise solicit one or
more people or organizations to evaluate the security of your
services, of which may include "hacking". You may wish to provide
for this in your policy.

2.3.3 Who Is Authorized to Grant Access and Approve Usage?

Your policy should state who is authorized to grant access to your
services. Further, it must be determined what type of access they
are permitted to give. If you do not have control over who is
granted access to your system, you will not have control over who
is using your system. Controlling who has the authorization to
grant access will also enable you to know who was or was not
granting access if problems develop later.

There are many schemes that can be developed to control the
distribution of access to your services. The following are the
factors that you must consider when determining who will
distribute access to your services:

    o Will you be distributing access from a centralized
      point or at various points?

You can have a centralized distribution point to a distributed
system where various sites or departments independently authorize
access. The trade off is between security and convenience. The
more centralized, the easier to secure.

    o What methods will you use for creating accounts and
      terminating access?

From a security standpoint, you need to examine the mechanism that

you will be using to create accounts. In the least restrictive
case, the people who are authorized to grant access would be able
to go into the system directly and create an account by hand or
through vendor supplied mechanisms. Generally, these mechanisms
place a great deal of trust in the person running them, and the
person running them usually has a large amount of privileges. If
this is the choice you make, you need to select someone who is
trustworthy to perform this task. The opposite solution is to
have an integrated system that the people authorized to create
accounts run, or the users themselves may actually run. Be aware
that even in the restrictive case of having a mechanized facility
to create accounts does not remove the potential for abuse.

You should have specific procedures developed for the creation of
accounts. These procedures should be well documented to prevent
confusion and reduce mistakes. A security vulnerability in the
account authorization process is not only possible through abuse,
but is also possible if a mistake is made. Having clear and well
documented procedure will help ensure that these mistakes won't
happen. You should also be sure that the people who will be
following these procedures understand them.

The granting of access to users is one of the most vulnerable of
times. You should ensure that the selection of an initial
password cannot be easily guessed. You should avoid using an
initial password that is a function of the username, is part of
the user's name, or some algorithmically generated password that
can easily be guessed. In addition, you should not permit users
to continue to use the initial password indefinitely. If
possible, you should force users to change the initial password
the first time they login. Consider that some users may never
even login, leaving their password vulnerable indefinitely. Some
sites choose to disable accounts that have never been accessed,
and force the owner to reauthorize opening the account.

2.3.4 Who May Have System Administration Privileges?

One security decision that needs to be made very carefully is who
will have access to system administrator privileges and passwords
for your services. Obviously, the system administrators will need
access, but inevitably other users will request special
privileges. The policy should address this issue. Restricting
privileges is one way to deal with threats from local users. The
challenge is to balance restricting access to these to protect
security with giving people who need these privileges access so
that they can perform their tasks. One approach that can be taken
is to grant only enough privilege to accomplish the necessary
tasks.

Additionally, people holding special privileges should be
accountable to some authority and this should also be identified
within the site's security policy. If the people you grant
privileges to are not accountable, you run the risk of losing
control of your system and will have difficulty managing a
compromise in security.

2.3.5 What Are The Users' Rights and Responsibilities?

The policy should incorporate a statement on the users' rights and
responsibilities concerning the use of the site's computer systems
and services. It should be clearly stated that users are
responsible for understanding and respecting the security rules of
the systems they are using. The following is a list of topics
that you may wish to cover in this area of the policy:

   o What guidelines you have regarding resource consumption
     (whether users are restricted, and if so, what the
     restrictions are).
   o What might constitute abuse in terms of system performance.
   o Whether users are permitted to share accounts or let others
     use their accounts.
   o How "secret" users should keep their passwords.
   o How often users should change their passwords and any other
     password restrictions or requirements.
   o Whether you provide backups or expect the users to create
     their own.
   o Disclosure of information that may be proprietary.
   o Statement on Electronic Mail Privacy (Electronic
     Communications Privacy Act).
   o Your policy concerning controversial mail or postings to
     mailing lists or discussion groups (obscenity, harassment,
     etc.).
   o Policy on electronic communications: mail forging, etc.

The Electronic Mail Association sponsored a white paper on the
privacy of electronic mail in companies [4]. Their basic
recommendation is that every site should have a policy on the
protection of employee privacy. They also recommend that
organizations establish privacy policies that deal with all media,
rather than singling out electronic mail.

They suggest five criteria for evaluating any policy:

   1. Does the policy comply with law and with duties to
      third parties?

   2. Does the policy unnecessarily compromise the interest of

              the employee, the employer or third parties?

          3. Is the policy workable as a practical matter and likely to
             be enforced?

          4. Does the policy deal appropriately with all different
             forms of communications and record keeping with the office?

          5. Has the policy been announced in advance and agreed to by
             all concerned?

   2.3.6 What Are The Rights and Responsibilities of System
          Administrators Versus Rights of Users

      There is a tradeoff between a user's right to absolute privacy and
      the need of system administrators to gather sufficient information
      to diagnose problems. There is also a distinction between a
      system administrator's need to gather information to diagnose
      problems and investigating security violations. The policy should
      specify to what degree system administrators can examine user
      files to diagnose problems or for other purposes, and what rights
      you grant to the users. You may also wish to make a statement
      concerning system administrators' obligation to maintaining the
      privacy of information viewed under these circumstances. A few
      questions that should be answered are:

          o Can an administrator monitor or read a user's files
            for any reason?
          o What are the liabilities?
          o Do network administrators have the right to examine
            network or host traffic?

   2.3.7 What To Do With Sensitive Information

      Before granting users access to your services, you need to
      determine at what level you will provide for the security of data
      on your systems. By determining this, you are determining the
      level of sensitivity of data that users should store on your
      systems. You do not want users to store very sensitive
      information on a system that you are not going to secure very
      well. You need to tell users who might store sensitive
      information what services, if any, are appropriate for the storage
      of sensitive information. This part should include storing of
      data in different ways (disk, magnetic tape, file servers, etc.).
      Your policy in this area needs to be coordinated with the policy
      concerning the rights of system administrators versus users (see
      section 2.3.6).

2.4 What Happens When the Policy is Violated

   It is obvious that when any type of official policy is defined, be it
   related to computer security or not, it will eventually be broken.
   The violation may occur due to an individual's negligence, accidental
   mistake, having not been properly informed of the current policy, or
   not understanding the current policy. It is equally possible that an
   individual (or group of individuals) may knowingly perform an act
   that is in direct violation of the defined policy.

   When a policy violation has been detected, the immediate course of
   action should be pre-defined to ensure prompt and proper enforcement.
   An investigation should be performed to determine how and why the
   violation occurred. Then the appropriate corrective action should be
   executed. The type and severity of action taken varies depending on
   the type of violation that occurred.

   2.4.1 Determining the Response to Policy Violations

      Violations to policy may be committed by a wide variety of users.
      Some may be local users and others may be from outside the local
      environment. Sites may find it helpful to define what it
      considers "insiders" and "outsiders" based upon administrative,
      legal or political boundaries. These boundaries imply what type
      of action must be taken to correct the offending party; from a
      written reprimand to pressing legal charges. So, not only do you
      need to define actions based on the type of violation, you also
      need to have a clearly defined series of actions based on the kind
      of user violating your computer security policy. This all seems
      rather complicated, but should be addressed long before it becomes
      necessary as the result of a violation.

      One point to remember about your policy is that proper education
      is your best defense. For the outsiders who are using your
      computer legally, it is your responsibility to verify that these
      individuals are aware of the policies that you have set forth.
      Having this proof may assist you in the future if legal action
      becomes necessary.

      As for users who are using your computer illegally, the problem is
      basically the same. What type of user violated the policy and how
      and why did they do it? Depending on the results of your
      investigation, you may just prefer to "plug" the hole in your
      computer security and chalk it up to experience. Or if a
      significant amount of loss was incurred, you may wish to take more
      drastic action.

2.4.2 What to do When Local Users Violate the Policy of a Remote
       Site

   In the event that a local user violates the security policy of a
   remote site, the local site should have a clearly defined set of
   administrative actions to take concerning that local user. The
   site should also be prepared to protect itself against possible
   actions by the remote site. These situations involve legal issues
   which should be addressed when forming the security policy.

2.4.3 Defining Contacts and Responsibilities to Outside
       Organizations

   The local security policy should include procedures for
   interaction with outside organizations. These include law
   enforcement agencies, other sites, external response team
   organizations (e.g., the CERT, CIAC) and various press agencies.
   The procedure should state who is authorized to make such contact
   and how it should be handled. Some questions to be answered
   include:

      o Who may talk to the press?
      o When do you contact law enforcement and investigative agencies?
      o If a connection is made from a remote site, is the
        system manager authorized to contact that site?
      o Can data be released? What kind?

   Detailed contact information should be readily available along
   with clearly defined procedures to follow.

2.4.4 What are the Responsibilities to our Neighbors and Other
       Internet Sites?

   The Security Policy Working Group within the IETF is working on a
   document entitled, "Policy Guidelines for the Secure Operation of
   the Internet" [23]. It addresses the issue that the Internet is a
   cooperative venture and that sites are expected to provide mutual
   security assistance. This should be addressed when developing a
   site's policy. The major issue to be determined is how much
   information should be released. This will vary from site to site
   according to the type of site (e.g., military, education,
   commercial) as well as the type of security violation that
   occurred.

2.4.5 Issues for Incident Handling Procedures

   Along with statements of policy, the document being prepared
   should include procedures for incident handling. This is covered

       in detail in the next chapter. There should be procedures
       available that cover all facets of policy violation.

2.5 Locking In or Out

    Whenever a site suffers an incident which may compromise computer
    security, the strategies for reacting may be influenced by two
    opposing pressures.

    If management fears that the site is sufficiently vulnerable, it may
    choose a "Protect and Proceed" strategy. This approach will have as
    its primary goal the protection and preservation of the site
    facilities and to provide for normalcy for its users as quickly as
    possible. Attempts will be made to actively interfere with the
    intruder's processes, prevent further access and begin immediate
    damage assessment and recovery. This process may involve shutting
    down the facilities, closing off access to the network, or other
    drastic measures. The drawback is that unless the intruder is
    identified directly, they may come back into the site via a different
    path, or may attack another site.

    The alternate approach, "Pursue and Prosecute", adopts the opposite
    philosophy and goals. The primary goal is to allow intruders to
    continue their activities at the site until the site can identify the
    responsible persons. This approach is endorsed by law enforcement
    agencies and prosecutors. The drawback is that the agencies cannot
    exempt a site from possible user lawsuits if damage is done to their
    systems and data.

    Prosecution is not the only outcome possible if the intruder is
    identified. If the culprit is an employee or a student, the
    organization may choose to take disciplinary actions. The computer
    security policy needs to spell out the choices and how they will be
    selected if an intruder is caught.

    Careful consideration must be made by site management regarding their
    approach to this issue before the problem occurs. The strategy
    adopted might depend upon each circumstance. Or there may be a
    global policy which mandates one approach in all circumstances. The
    pros and cons must be examined thoroughly and the users of the
    facilities must be made aware of the policy so that they understand
    their vulnerabilities no matter which approach is taken.

    The following are checklists to help a site determine which strategy
    to adopt: "Protect and Proceed" or "Pursue and Prosecute".

Protect and Proceed

    1. If assets are not well protected.

    2. If continued penetration could result in great
       financial risk.

    3. If the possibility or willingness to prosecute
       is not present.

    4. If user base is unknown.

    5. If users are unsophisticated and their work is
       vulnerable.

    6. If the site is vulnerable to lawsuits from users, e.g.,
       if their resources are undermined.

Pursue and Prosecute

    1. If assets and systems are well protected.

    2. If good backups are available.

    3. If the risk to the assets is outweighed by the
       disruption caused by the present and possibly future
       penetrations.

    4. If this is a concentrated attack occurring with great
       frequency and intensity.

    5. If the site has a natural attraction to intruders, and
       consequently regularly attracts intruders.

    6. If the site is willing to incur the financial (or other)
       risk to assets by allowing the penetrator continue.

    7. If intruder access can be controlled.

    8. If the monitoring tools are sufficiently well-developed
       to make the pursuit worthwhile.

    9. If the support staff is sufficiently clever and knowledgable
       about the operating system, related utilities, and systems
       to make the pursuit worthwhile.

    10. If there is willingness on the part of management to
        prosecute.

11. If the system adminitrators know in general what kind of
    evidence would lead to prosecution.

12. If there is established contact with knowledgeable law
    enforcement.

13. If there is a site representative versed in the relevant
    legal issues.

14. If the site is prepared for possible legal action from
    its own users if their data or systems become compromised
    during the pursuit.

2.6 Interpreting the Policy

It is important to define who will interpret the policy. This could
be an individual or a committee. No matter how well written, the
policy will require interpretation from time to time and this body
would serve to review, interpret, and revise the policy as needed.

2.7 Publicizing the Policy

Once the site security policy has been written and established, a
vigorous process should be engaged to ensure that the policy
statement is widely and thoroughly disseminated and discussed. A
mailing of the policy should not be considered sufficient. A period
for comments should be allowed before the policy becomes effective to
ensure that all affected users have a chance to state their reactions
and discuss any unforeseen ramifications. Ideally, the policy should
strike a balance between protection and productivity.

Meetings should be held to elicit these comments, and also to ensure
that the policy is correctly understood. (Policy promulgators are
not necessarily noted for their skill with the language.) These
meetings should involve higher management as well as line employees.
Security is a collective effort.

In addition to the initial efforts to publicize the policy, it is
essential for the site to maintain a continual awareness of its
computer security policy. Current users may need periodic reminders
New users should have the policy included as part of their site
introduction packet. As a condition for using the site facilities,
it may be advisable to have them sign a statement that they have read
and understood the policy. Should any of these users require legal
action for serious policy violations, this signed statement might
prove to be a valuable aid.

3. Establishing Procedures to Prevent Security Problems

   The security policy defines what needs to be protected. This section
   discusses security procedures which specify what steps will be used
   to carry out the security policy.

3.1 Security Policy Defines What Needs to be Protected

   The security policy defines the WHAT's: what needs to be protected,
   what is most important, what the priorities are, and what the general
   approach to dealing with security problems should be.

   The security policy by itself doesn't say HOW things are protected.
   That is the role of security procedures, which this section
   discusses. The security policy should be a high level document,
   giving general strategy. The security procedures need to set out, in
   detail, the precise steps your site will take to protect itself.

   The security policy should include a general risk assessment of the
   types of threats a site is mostly likely to face and the consequences
   of those threats (see section 2.2). Part of doing a risk assessment
   will include creating a general list of assets that should be
   protected (section 2.2.2). This information is critical in devising
   cost-effective procedures.

   It is often tempting to start creating security procedures by
   deciding on different mechanisms first: "our site should have logging
   on all hosts, call-back modems, and smart cards for all users." This
   approach could lead to some areas that have too much protection for
   the risk they face, and other areas that aren't protected enough.
   Starting with the security policy and the risks it outlines should
   ensure that the procedures provide the right level of protect for all
   assets.

3.2 Identifying Possible Problems

   To determine risk, vulnerabilities must be identified. Part of the
   purpose of the policy is to aid in shoring up the vulnerabilities and
   thus to decrease the risk in as many areas as possible. Several of
   the more popular problem areas are presented in sections below. This
   list is by no means complete. In addition, each site is likely to
   have a few unique vulnerabilities.

   3.2.1 Access Points

      Access points are typically used for entry by unauthorized users.
      Having many access points increases the risk of access to an
      organization's computer and network facilities.

Network links to networks outside the organization allow access
into the organization for all others connected to that external
network. A network link typically provides access to a large
number of network services, and each service has a potential to be
compromised.

Dialup lines, depending on their configuration, may provide access
merely to a login port of a single system. If connected to a
terminal server, the dialup line may give access to the entire
network.

Terminal servers themselves can be a source of problem. Many
terminal servers do not require any kind of authentication.
Intruders often use terminal servers to disguise their actions,
dialing in on a local phone and then using the terminal server to
go out to the local network. Some terminal servers are configured
so that intruders can TELNET [19] in from outside the network, and
then TELNET back out again, again serving to make it difficult to
trace them.

3.2.2 Misconfigured Systems

Misconfigured systems form a large percentage of security holes.
Today's operating systems and their associated software have
become so complex that understanding how the system works has
become a full-time job. Often, systems managers will be non-
specialists chosen from the current organization's staff.

Vendors are also partly responsible for misconfigured systems. To
make the system installation process easier, vendors occasionally
choose initial configurations that are not secure in all
environments.

3.2.3 Software Bugs

Software will never be bug free. Publicly known security bugs are
common methods of unauthorized entry. Part of the solution to
this problem is to be aware of the security problems and to update
the software when problems are detected. When bugs are found,
they should be reported to the vendor so that a solution to the
problem can be implemented and distributed.

3.2.4 "Insider" Threats

An insider to the organization may be a considerable threat to the
security of the computer systems. Insiders often have direct
access to the computer and network hardware components. The
ability to access the components of a system makes most systems

easier to compromise. Most desktop workstations can be easily
manipulated so that they grant privileged access. Access to a
local area network provides the ability to view possibly sensitive
data traversing the network.

3.3 Choose Controls to Protect Assets in a Cost-Effective Way

   After establishing what is to be protected, and assessing the risks
   these assets face, it is necessary to decide how to implement the
   controls which protect these assets. The controls and protection
   mechanisms should be selected in a way so as to adequately counter
   the threats found during risk assessment, and to implement those
   controls in a cost effective manner. It makes little sense to spend
   an exorbitant sum of money and overly constrict the user base if the
   risk of exposure is very small.

   3.3.1 Choose the Right Set of Controls

      The controls that are selected represent the physical embodiment
      of your security policy. They are the first and primary line of
      defense in the protection of your assets. It is therefore most
      important to ensure that the controls that you select are the
      right set of controls. If the major threat to your system is
      outside penetrators, it probably doesn't make much sense to use
      biometric devices to authenticate your regular system users. On
      the other hand, if the major threat is unauthorized use of
      computing resources by regular system users, you'll probably want
      to establish very rigorous automated accounting procedures.

   3.3.2 Use Common Sense

      Common sense is the most appropriate tool that can be used to
      establish your security policy. Elaborate security schemes and
      mechanisms are impressive, and they do have their place, yet there
      is little point in investing money and time on an elaborate
      implementation scheme if the simple controls are forgotten. For
      example, no matter how elaborate a system you put into place on
      top of existing security controls, a single user with a poor
      password can still leave your system open to attack.

3.4 Use Multiple Strategies to Protect Assets

   Another method of protecting assets is to use multiple strategies.
   In this way, if one strategy fails or is circumvented, another
   strategy comes into play to continue protecting the asset. By using
   several simpler strategies, a system can often be made more secure
   than if one very sophisticated method were used in its place. For
   example, dial-back modems can be used in conjunction with traditional

logon mechanisms. Many similar approaches could be devised that
provide several levels of protection for assets. However, it's very
easy to go overboard with extra mechanisms. One must keep in mind
exactly what it is that needs to be protected.

3.5 Physical Security

   It is a given in computer security if the system itself is not
   physically secure, nothing else about the system can be considered
   secure. With physical access to a machine, an intruder can halt the
   machine, bring it back up in privileged mode, replace or alter the
   disk, plant Trojan horse programs (see section 2.13.9.2), or take any
   number of other undesirable (and hard to prevent) actions.

   Critical communications links, important servers, and other key
   machines should be located in physically secure areas. Some security
   systems (such as Kerberos) require that the machine be physically
   secure.

   If you cannot physically secure machines, care should be taken about
   trusting those machines. Sites should consider limiting access from
   non-secure machines to more secure machines. In particular, allowing
   trusted access (e.g., the BSD Unix remote commands such as rsh) from
   these kinds of hosts is particularly risky.

   For machines that seem or are intended to be physically secure, care
   should be taken about who has access to the machines. Remember that
   custodial and maintenance staff often have keys to rooms.

3.6   Procedures to Recognize Unauthorized Activity

   Several simple procedures can be used to detect most unauthorized
   uses of a computer system. These procedures use tools provided with
   the operating system by the vendor, or tools publicly available from
   other sources.

   3.6.1 Monitoring System Use

      System monitoring can be done either by a system administrator, or
      by software written for the purpose. Monitoring a system involves
      looking at several parts of the system and searching for anything
      unusual. Some of the easier ways to do this are described in this
      section.

      The most important thing about monitoring system use is that it be
      done on a regular basis. Picking one day out of the month to
      monitor the system is pointless, since a security breach can be
      isolated to a matter of hours. Only by maintaining a constant

vigil can you expect to detect security violations in time to
react to them.

3.6.2 Tools for Monitoring the System

This section describes tools and methods for monitoring a system
against unauthorized access and use.

3.6.2.1 Logging

Most operating systems store numerous bits of information in
log files. Examination of these log files on a regular basis
is often the first line of defense in detecting unauthorized
use of the system.

    - Compare lists of currently logged in users and past
      login histories. Most users typically log in and out
      at roughly the same time each day. An account logged
      in outside the "normal" time for the account may be in
      use by an intruder.

    - Many systems maintain accounting records for billing
      purposes. These records can also be used to determine
      usage patterns for the system; unusual accounting records
      may indicate unauthorized use of the system.

    - System logging facilities, such as the UNIX "syslog"
      utility, should be checked for unusual error messages
      from system software. For example, a large number of
      failed login attempts in a short period of time may
      indicate someone trying to guess passwords.

    - Operating system commands which list currently executing
      processes can be used to detect users running programs
      they are not authorized to use, as well as to detect
      unauthorized programs which have been started by an
      intruder.

3.6.2.2 Monitoring Software

Other monitoring tools can easily be constructed using standard
operating system software, by using several, often unrelated,
programs together. For example, checklists of file ownerships
and permission settings can be constructed (for example, with
"ls" and "find" on UNIX) and stored off-line. These lists can
then be reconstructed periodically and compared against the
master checklist (on UNIX, by using the "diff" utility).
Differences may indicate that unauthorized modifications have

been made to the system.

Still other tools are available from third-party vendors and
public software distribution sites. Section 3.9.9 lists
several sources from which you can learn what tools are
available and how to get them.

3.6.2.3 Other Tools

Other tools can also be used to monitor systems for security
violations, although this is not their primary purpose. For
example, network monitors can be used to detect and log
connections from unknown sites.

3.6.3 Vary the Monitoring Schedule

The task of system monitoring is not as daunting as it may seem.
System administrators can execute many of the commands used for
monitoring periodically throughout the day during idle moments
(e.g., while talking on the telephone), rather than spending fixed
periods of each day monitoring the system. By executing the
commands frequently, you will rapidly become used to seeing
"normal" output, and will easily spot things which are out of the
ordinary. In addition, by running various monitoring commands at
different times throughout the day, you make it hard for an
intruder to predict your actions. For example, if an intruder
knows that each day at 5:00 p.m. the system is checked to see that
everyone has logged off, he will simply wait until after the check
has completed before logging in. But the intruder cannot guess
when a system administrator might type a command to display all
logged-in users, and thus he runs a much greater risk of
detection.

Despite the advantages that regular system monitoring provides,
some intruders will be aware of the standard logging mechanisms in
use on systems they are attacking. They will actively pursue and
attempt to disable monitoring mechanisms. Regular monitoring
therefore is useful in detecting intruders, but does not provide
any guarantee that your system is secure, nor should monitoring be
considered an infallible method of detecting unauthorized use.

3.7 Define Actions to Take When Unauthorized Activity is Suspected

Sections 2.4 and 2.5 discussed the course of action a site should
take when it suspects its systems are being abused. The computer
security policy should state the general approach towards dealing
with these problems.

The procedures for dealing with these types of problems should be
written down. Who has authority to decide what actions will be
taken? Should law enforcement be involved? Should your
organization cooperate with other sites in trying to track down an
intruder? Answers to all the questions in section 2.4 should be
part of the incident handling procedures.

Whether you decide to lock out or pursue intruders, you should
have tools and procedures ready to apply. It is best to work up
these tools and procedures before you need them. Don't wait until
an intruder is on your system to figure out how to track the
intruder's actions; you will be busy enough if an intruder
strikes.

3.8 Communicating Security Policy

   Security policies, in order to be effective, must be communicated to
   both the users of the system and the system maintainers. This
   section describes what these people should be told, and how to tell
   them.

   3.8.1 Educating the Users

      Users should be made aware of how the computer systems are
      expected to be used, and how to protect themselves from
      unauthorized users.

      3.8.1.1 Proper Account/Workstation Use

         All users should be informed about what is considered the
         "proper" use of their account or workstation ("proper" use is
         discussed in section 2.3.2). This can most easily be done at
         the time a user receives their account, by giving them a policy
         statement. Proper use policies typically dictate things such
         as whether or not the account or workstation may be used for
         personal activities (such as checkbook balancing or letter
         writing), whether profit-making activities are allowed, whether
         game playing is permitted, and so on. These policy statements
         may also be used to summarize how the computer facility is
         licensed and what software licenses are held by the
         institution; for example, many universities have educational
         licenses which explicitly prohibit commercial uses of the
         system. A more complete list of items to consider when writing
         a policy statement is given in section 2.3.

      3.8.1.2 Account/Workstation Management Procedures

         Each user should be told how to properly manage their account

and workstation. This includes explaining how to protect files
stored on the system, how to log out or lock the terminal or
workstation, and so on. Much of this information is typically
covered in the "beginning user" documentation provided by the
operating system vendor, although many sites elect to
supplement this material with local information.

If your site offers dial-up modem access to the computer
systems, special care must be taken to inform users of the
security problems inherent in providing this access. Issues
such as making sure to log out before hanging up the modem
should be covered when the user is initially given dial-up
access.

Likewise, access to the systems via local and wide-area
networks presents its own set of security problems which users
should be made aware of. Files which grant "trusted host" or
"trusted user" status to remote systems and users should be
carefully explained.

3.8.1.3 Determining Account Misuse

Users should be told how to detect unauthorized access to their
account. If the system prints the last login time when a user
logs in, he or she should be told to check that time and note
whether or not it agrees with the last time he or she actually
logged in.

Command interpreters on some systems (e.g., the UNIX C shell)
maintain histories of the last several commands executed.
Users should check these histories to be sure someone has not
executed other commands with their account.

3.8.1.4 Problem Reporting Procedures

A procedure should be developed to enable users to report
suspected misuse of their accounts or other misuse they may
have noticed. This can be done either by providing the name
and telephone number of a system administrator who manages
security of the computer system, or by creating an electronic
mail address (e.g., "security") to which users can address
their problems.

3.8.2 Educating the Host Administrators

In many organizations, computer systems are administered by a wide
variety of people. These administrators must know how to protect
their own systems from attack and unauthorized use, as well as how

to communicate successful penetration of their systems to other
administrators as a warning.

3.8.2.1 Account Management Procedures

   Care must be taken when installing accounts on the system in
   order to make them secure. When installing a system from
   distribution media, the password file should be examined for
   "standard" accounts provided by the vendor. Many vendors
   provide accounts for use by system services or field service
   personnel. These accounts typically have either no password or
   one which is common knowledge. These accounts should be given
   new passwords if they are needed, or disabled or deleted from
   the system if they are not.

   Accounts without passwords are generally very dangerous since
   they allow anyone to access the system. Even accounts which do
   not execute a command interpreter (e.g., accounts which exist
   only to see who is logged in to the system) can be compromised
   if set up incorrectly. A related concept, that of "anonymous"
   file transfer (FTP) [20], allows users from all over the
   network to access your system to retrieve files from (usually)
   a protected disk area. You should carefully weigh the benefits
   that an account without a password provides against the
   security risks of providing such access to your system.

   If the operating system provides a "shadow" password facility
   which stores passwords in a separate file accessible only to
   privileged users, this facility should be used. System V UNIX,
   SunOS 4.0 and above, and versions of Berkeley UNIX after 4.3BSD
   Tahoe, as well as others, provide this feature. It protects
   passwords by hiding their encrypted values from unprivileged
   users. This prevents an attacker from copying your password
   file to his or her machine and then attempting to break the
   passwords at his or her leisure.

   Keep track of who has access to privileged user accounts (e.g.,
   "root" on UNIX or "MAINT" on VMS). Whenever a privileged user
   leaves the organization or no longer has need of the privileged
   account, the passwords on all privileged accounts should be
   changed.

3.8.2.2 Configuration Management Procedures

   When installing a system from the distribution media or when
   installing third-party software, it is important to check the
   installation carefully. Many installation procedures assume a
   "trusted" site, and hence will install files with world write

permission enabled, or otherwise compromise the security of
files.

Network services should also be examined carefully when first
installed. Many vendors provide default network permission
files which imply that all outside hosts are to be "trusted",
which is rarely the case when connected to wide-area networks
such as the Internet.

Many intruders collect information on the vulnerabilities of
particular system versions. The older a system, the more
likely it is that there are security problems in that version
which have since been fixed by the vendor in a later release.
For this reason, it is important to weigh the risks of not
upgrading to a new operating system release (thus leaving
security holes unplugged) against the cost of upgrading to the
new software (possibly breaking third-party software, etc.).
Bug fixes from the vendor should be weighed in a similar
fashion, with the added note that "security" fixes from a
vendor usually address fairly serious security problems.

Other bug fixes, received via network mailing lists and the
like, should usually be installed, but not without careful
examination. Never install a bug fix unless you're sure you
know what the consequences of the fix are - there's always the
possibility that an intruder has suggested a "fix" which
actually gives him or her access to your system.

3.8.2.3 Recovery Procedures - Backups

It is impossible to overemphasize the need for a good backup
strategy. File system backups not only protect you in the
event of hardware failure or accidental deletions, but they
also protect you against unauthorized changes made by an
intruder. Without a copy of your data the way it's "supposed"
to be, it can be difficult to undo something an attacker has
done.

Backups, especially if run daily, can also be useful in
providing a history of an intruder's activities. Looking
through old backups can establish when your system was first
penetrated. Intruders may leave files around which, although
deleted later, are captured on the backup tapes. Backups can
also be used to document an intruder's activities to law
enforcement agencies if necessary.

A good backup strategy will dump the entire system to tape at
least once a month. Partial (or "incremental") dumps should be

done at least twice a week, and ideally they should be done
daily. Commands specifically designed for performing file
system backups (e.g., UNIX "dump" or VMS "BACKUP") should be
used in preference to other file copying commands, since these
tools are designed with the express intent of restoring a
system to a known state.

3.8.2.4 Problem Reporting Procedures

As with users, system administrators should have a defined
procedure for reporting security problems. In large
installations, this is often done by creating an electronic
mail alias which contains the names of all system
administrators in the organization. Other methods include
setting up some sort of response team similar to the CERT, or
establishing a "hotline" serviced by an existing support group.

3.9 Resources to Prevent Security Breaches

This section discusses software, hardware, and procedural resources
that can be used to support your site security policy.

3.9.1 Network Connections and Firewalls

A "firewall" is put in place in a building to provide a point of
resistance to the entry of flames into another area. Similarly, a
secretary's desk and reception area provides a point of
controlling access to other office spaces. This same technique
can be applied to a computer site, particularly as it pertains to
network connections.

Some sites will be connected only to other sites within the same
organization and will not have the ability to connect to other
networks. Sites such as these are less susceptible to threats
from outside their own organization, although intrusions may still
occur via paths such as dial-up modems. On the other hand, many
other organizations will be connected to other sites via much
larger networks, such as the Internet. These sites are
susceptible to the entire range of threats associated with a
networked environment.

The risks of connecting to outside networks must be weighed
against the benefits. It may be desirable to limit connection to
outside networks to those hosts which do not store sensitive
material, keeping "vital" machines (such as those which maintain
company payroll or inventory systems) isolated. If there is a
need to participate in a Wide Area Network (WAN), consider
restricting all access to your local network through a single

system. That is, all access to or from your own local network
must be made through a single host computer that acts as a
firewall between you and the outside world. This firewall system
should be rigorously controlled and password protected, and
external users accessing it should also be constrained by
restricting the functionality available to remote users. By using
this approach, your site could relax some of the internal security
controls on your local net, but still be afforded the protection
of a rigorously controlled host front end.

Note that even with a firewall system, compromise of the firewall
could result in compromise of the network behind the firewall.
Work has been done in some areas to construct a firewall which
even when compromised, still protects the local network [6,
CHESWICK].

3.9.2 Confidentiality

Confidentiality, the act of keeping things hidden or secret, is
one of the primary goals of computer security practitioners.
Several mechanisms are provided by most modern operating systems
to enable users to control the dissemination of information.
Depending upon where you work, you may have a site where
everything is protected, or a site where all information is
usually regarded as public, or something in-between. Most sites
lean toward the in-between, at least until some penetration has
occurred.

Generally, there are three instances in which information is
vulnerable to disclosure: when the information is stored on a
computer system, when the information is in transit to another
system (on the network), and when the information is stored on
backup tapes.

The first of these cases is controlled by file permissions, access
control lists, and other similar mechanisms. The last can be
controlled by restricting access to the backup tapes (by locking
them in a safe, for example). All three cases can be helped by
using encryption mechanisms.

3.9.2.1 Encryption (hardware and software)

Encryption is the process of taking information that exists in
some readable form and converting it into a non-readable form.
There are several types of commercially available encryption
packages in both hardware and software forms. Hardware
encryption engines have the advantage that they are much faster
than the software equivalent, yet because they are faster, they

are of greater potential benefit to an attacker who wants to
execute a brute-force attack on your encrypted information.

The advantage of using encryption is that, even if other access
control mechanisms (passwords, file permissions, etc.) are
compromised by an intruder, the data is still unusable.
Naturally, encryption keys and the like should be protected at
least as well as account passwords.

Information in transit (over a network) may be vulnerable to
interception as well. Several solutions to this exist, ranging
from simply encrypting files before transferring them (end-to-
end encryption) to special network hardware which encrypts
everything it sends without user intervention (secure links).
The Internet as a whole does not use secure links, thus end-
to-end encryption must be used if encryption is desired across
the Internet.

3.9.2.1.1 Data Encryption Standard (DES)

   DES is perhaps the most widely used data encryption
   mechanism today. Many hardware and software implementations
   exist, and some commercial computers are provided with a
   software version. DES transforms plain text information
   into encrypted data (or ciphertext) by means of a special
   algorithm and "seed" value called a key. So long as the key
   is retained (or remembered) by the original user, the
   ciphertext can be restored to the original plain text.

   One of the pitfalls of all encryption systems is the need to
   remember the key under which a thing was encrypted (this is
   not unlike the password problem discussed elsewhere in this
   document). If the key is written down, it becomes less
   secure. If forgotten, there is little (if any) hope of
   recovering the original data.

   Most UNIX systems provide a DES command that enables a user
   to encrypt data using the DES algorithm.

3.9.2.1.2 Crypt

   Similar to the DES command, the UNIX "crypt" command allows
   a user to encrypt data. Unfortunately, the algorithm used
   by "crypt" is very insecure (based on the World War II
   "Enigma" device), and files encrypted with this command can
   be decrypted easily in a matter of a few hours. Generally,
   use of the "crypt" command should be avoided for any but the
   most trivial encryption tasks.

3.9.2.2 Privacy Enhanced Mail

Electronic mail normally transits the network in the clear
(i.e., anyone can read it). This is obviously not the optimal
solution. Privacy enhanced mail provides a means to
automatically encrypt electronic mail messages so that a person
eavesdropping at a mail distribution node is not (easily)
capable of reading them. Several privacy enhanced mail
packages are currently being developed and deployed on the
Internet.

The Internet Activities Board Privacy Task Force has defined a
draft standard, elective protocol for use in implementing
privacy enhanced mail. This protocol is defined in RFCs 1113,
1114, and 1115 [7,8,9]. Please refer to the current edition of
the "IAB Official Protocol Standards" (currently, RFC 1200
[21]) for the standardization state and status of these
protocols.

3.9.3 Origin Authentication

We mostly take it on faith that the header of an electronic mail
message truly indicates the originator of a message. However, it
iseasy to "spoof", or forge the source of a mail message. Origin
authentication provides a means to be certain of the originator of
a message or other object in the same way that a Notary Public
assures a signature on a legal document. This is done by means of
a "Public Key" cryptosystem.

A public key cryptosystem differs from a private key cryptosystem
in several ways. First, a public key system uses two keys, a
Public Key that anyone can use (hence the name) and a Private Key
that only the originator of a message uses. The originator uses
the private key to encrypt the message (as in DES). The receiver,
who has obtained the public key for the originator, may then
decrypt the message.

In this scheme, the public key is used to authenticate the
originator's use of his or her private key, and hence the identity
of the originator is more rigorously proven. The most widely
known implementation of a public key cryptosystem is the RSA
system [26]. The Internet standard for privacy enhanced mail
makes use of the RSA system.

3.9.4 Information Integrity

Information integrity refers to the state of information such that
it is complete, correct, and unchanged from the last time in which

it was verified to be in an "integral" state. The value of
information integrity to a site will vary. For example, it is
more important for military and government installations to
prevent the "disclosure" of classified information, whether it is
right or wrong. A bank, on the other hand, is far more concerned
with whether the account information maintained for its customers
is complete and accurate.

Numerous computer system mechanisms, as well as procedural
controls, have an influence on the integrity of system
information. Traditional access control mechanisms maintain
controls over who can access system information. These mechanisms
alone are not sufficient in some cases to provide the degree of
integrity required. Some other mechanisms are briefly discussed
below.

It should be noted that there are other aspects to maintaining
system integrity besides these mechanisms, such as two-person
controls, and integrity validation procedures. These are beyond
the scope of this document.

3.9.4.1 Checksums

   Easily the simplest mechanism, a simple checksum routine can
   compute a value for a system file and compare it with the last
   known value. If the two are equal, the file is probably
   unchanged. If not, the file has been changed by some unknown
   means.

   Though it is the easiest to implement, the checksum scheme
   suffers from a serious failing in that it is not very
   sophisticated and a determined attacker could easily add enough
   characters to the file to eventually obtain the correct value.

   A specific type of checksum, called a CRC checksum, is
   considerably more robust than a simple checksum. It is only
   slightly more difficult to implement and provides a better
   degree of catching errors. It too, however, suffers from the
   possibility of compromise by an attacker.

   Checksums may be used to detect the altering of information.
   However, they do not actively guard against changes being made.
   For this, other mechanisms such as access controls and
   encryption should be used.

3.9.4.2 Cryptographic Checksums

   Cryptographic checksums (also called cryptosealing) involve
   breaking a file up into smaller chunks, calculating a (CRC)
   checksum for each chunk, and adding the CRCs together.
   Depending upon the exact algorithm used, this can result in a
   nearly unbreakable method of determining whether a file has
   been changed. This mechanism suffers from the fact that it is
   sometimes computationally intensive and may be prohibitive
   except in cases where the utmost integrity protection is
   desired.

   Another related mechanism, called a one-way hash function (or a
   Manipulation Detection Code (MDC)) can also be used to uniquely
   identify a file. The idea behind these functions is that no
   two inputs can produce the same output, thus a modified file
   will not have the same hash value. One-way hash functions can
   be implemented efficiently on a wide variety of systems, making
   unbreakable integrity checks possible. (Snefru, a one-way hash
   function available via USENET as well as the Internet is just
   one example of an efficient one-way hash function.) [10]

3.9.5 Limiting Network Access

   The dominant network protocols in use on the Internet, IP (RFC
   791) [11], TCP (RFC 793) [12], and UDP (RFC 768) [13], carry
   certain control information which can be used to restrict access
   to certain hosts or networks within an organization.

   The IP packet header contains the network addresses of both the
   sender and recipient of the packet. Further, the TCP and UDP
   protocols provide the notion of a "port", which identifies the
   endpoint (usually a network server) of a communications path. In
   some instances, it may be desirable to deny access to a specific
   TCP or UDP port, or even to certain hosts and networks altogether.

3.9.5.1 Gateway Routing Tables

   One of the simplest approaches to preventing unwanted network
   connections is to simply remove certain networks from a
   gateway's routing tables. This makes it "impossible" for a
   host to send packets to these networks. (Most protocols
   require bidirectional packet flow even for unidirectional data
   flow, thus breaking one side of the route is usually
   sufficient.)

   This approach is commonly taken in "firewall" systems by
   preventing the firewall from advertising local routes to the

outside world. The approach is deficient in that it often
prevents "too much" (e.g., in order to prevent access to one
system on the network, access to all systems on the network is
disabled).

3.9.5.2 Router Packet Filtering

Many commercially available gateway systems (more correctly
called routers) provide the ability to filter packets based not
only on sources or destinations, but also on source-destination
combinations. This mechanism can be used to deny access to a
specific host, network, or subnet from any other host, network,
or subnet.

Gateway systems from some vendors (e.g., cisco Systems) support
an even more complex scheme, allowing finer control over source
and destination addresses. Via the use of address masks, one
can deny access to all but one host on a particular network.
The cisco Systems also allow packet screening based on IP
protocol type and TCP or UDP port numbers [14].

This can also be circumvented by "source routing" packets
destined for the "secret" network. Source routed packets may
be filtered out by gateways, but this may restrict other
legitimate activities, such as diagnosing routing problems.

3.9.6 Authentication Systems

Authentication refers to the process of proving a claimed identity
to the satisfaction of some permission-granting authority.
Authentication systems are hardware, software, or procedural
mechanisms that enable a user to obtain access to computing
resources. At the simplest level, the system administrator who
adds new user accounts to the system is part of the system
authentication mechanism. At the other end of the spectrum,
fingerprint readers or retinal scanners provide a very high-tech
solution to establishing a potential user's identity. Without
establishing and proving a user's identity prior to establishing a
session, your site's computers are vulnerable to any sort of
attack.

Typically, a user authenticates himself or herself to the system
by entering a password in response to a prompt.
Challenge/Response mechanisms improve upon passwords by prompting
the user for some piece of information shared by both the computer
and the user (such as mother's maiden name, etc.).

3.9.6.1 Kerberos

   Kerberos, named after the dog who in mythology is said to stand
   at the gates of Hades, is a collection of software used in a
   large network to establish a user's claimed identity.
   Developed at the Massachusetts Institute of Technology (MIT),
   it uses a combination of encryption and distributed databases
   so that a user at a campus facility can login and start a
   session from any computer located on the campus. This has
   clear advantages in certain environments where there are a
   large number of potential users who may establish a connection
   from any one of a large number of workstations. Some vendors
   are now incorporating Kerberos into their systems.

   It should be noted that while Kerberos makes several advances
   in the area of authentication, some security weaknesses in the
   protocol still remain [15].

3.9.6.2 Smart Cards

   Several systems use "smart cards" (a small calculator-like
   device) to help authenticate users. These systems depend on
   the user having an object in their possession. One such system
   involves a new password procedure that require a user to enter
   a value obtained from a "smart card" when asked for a password
   by the computer. Typically, the host machine will give the
   user some piece of information that is entered into the
   keyboard of the smart card. The smart card will display a
   response which must then be entered into the computer before
   the session will be established. Another such system involves
   a smart card which displays a number which changes over time,
   but which is synchronized with the authentication software on
   the computer.

   This is a better way of dealing with authentication than with
   the traditional password approach. On the other hand, some say
   it's inconvenient to carry the smart card. Start-up costs are
   likely to be high as well.

3.9.7 Books, Lists, and Informational Sources

   There are many good sources for information regarding computer
   security. The annotated bibliography at the end of this document
   can provide you with a good start. In addition, information can
   be obtained from a variety of other sources, some of which are
   described in this section.

3.9.7.1 Security Mailing Lists

The UNIX Security mailing list exists to notify system
administrators of security problems before they become common
knowledge, and to provide security enhancement information. It
is a restricted-access list, open only to people who can be
verified as being principal systems people at a site. Requests
to join the list must be sent by either the site contact listed
in the Defense Data Network's Network Information Center's (DDN
NIC) WHOIS database, or from the "root" account on one of the
major site machines. You must include the destination address
you want on the list, an indication of whether you want to be
on the mail reflector list or receive weekly digests, the
electronic mail address and voice telephone number of the site
contact if it isn't you, and the name, address, and telephone
number of your organization. This information should be sent
to SECURITY-REQUEST@CPD.COM.

The RISKS digest is a component of the ACM Committee on
Computers and Public Policy, moderated by Peter G. Neumann. It
is a discussion forum on risks to the public in computers and
related systems, and along with discussing computer security
and privacy issues, has discussed such subjects as the Stark
incident, the shooting down of the Iranian airliner in the
Persian Gulf (as it relates to the computerized weapons
systems), problems in air and railroad traffic control systems,
software engineering, and so on. To join the mailing list,
send a message to RISKS-REQUEST@CSL.SRI.COM. This list is also
available in the USENET newsgroup "comp.risks".

The VIRUS-L list is a forum for the discussion of computer
virus experiences, protection software, and related topics.
The list is open to the public, and is implemented as a
moderated digest. Most of the information is related to
personal computers, although some of it may be applicable to
larger systems. To subscribe, send the line:

    SUB VIRUS-L your full name

to the address LISTSERV%LEHIIBM1.BITNET@MITVMA.MIT.EDU. This
list is also available via the USENET newsgroup "comp.virus".

The Computer Underground Digest "is an open forum dedicated to
sharing information among computerists and to the presentation
and debate of diverse views." While not directly a security
list, it does contain discussions about privacy and other
security related topics. The list can be read on USENET as
alt.society.cu-digest, or to join the mailing list, send mail

to Gordon Myer (TK0JUT2%NIU.bitnet@mitvma.mit.edu).
Submissions may be mailed to: cud@chinacat.unicom.com.

3.9.7.2 Networking Mailing Lists

The TCP-IP mailing list is intended to act as a discussion
forum for developers and maintainers of implementations of the
TCP/IP protocol suite. It also discusses network-related
security problems when they involve programs providing network
services, such as "Sendmail". To join the TCP-IP list, send a
message to TCP-IP-REQUEST@NISC.SRI.COM. This list is also
available in the USENET newsgroup "comp.protocols.tcp-ip".

SUN-NETS is a discussion list for items pertaining to
networking on Sun systems. Much of the discussion is related
to NFS, NIS (formally Yellow Pages), and name servers. To
subscribe, send a message to SUN-NETS-REQUEST@UMIACS.UMD.EDU.

The USENET groups misc.security and alt.security also discuss
security issues. misc.security is a moderated group and also
includes discussions of physical security and locks.
alt.security is unmoderated.

3.9.7.3 Response Teams

Several organizations have formed special groups of people to
deal with computer security problems. These teams collect
information about possible security holes and disseminate it to
the proper people, track intruders, and assist in recovery from
security violations. The teams typically have both electronic
mail distribution lists as well as a special telephone number
which can be called for information or to report a problem.
Many of these teams are members of the CERT System, which is
coordinated by the National Institute of Standards and
Technology (NIST), and exists to facilitate the exchange of
information between the various teams.

3.9.7.3.1 DARPA Computer Emergency Response Team

The Computer Emergency Response Team/Coordination Center
(CERT/CC) was established in December 1988 by the Defense
Advanced Research Projects Agency (DARPA) to address
computer security concerns of research users of the
Internet. It is operated by the Software Engineering
Institute (SEI) at Carnegie-Mellon University (CMU). The
CERT can immediately confer with experts to diagnose and
solve security problems, and also establish and maintain
communications with the affected computer users and

government authorities as appropriate.

The CERT/CC serves as a clearing house for the
identification and repair of security vulnerabilities,
informal assessments of existing systems, improvement of
emergency response capability, and both vendor and user
security awareness. In addition, the team works with
vendors of various systems in order to coordinate the fixes
for security problems.

The CERT/CC sends out security advisories to the CERT-
ADVISORY mailing list whenever appropriate. They also
operate a 24-hour hotline that can be called to report
security problems (e.g., someone breaking into your system),
as well as to obtain current (and accurate) information
about rumored security problems.

To join the CERT-ADVISORY mailing list, send a message to
CERT@CERT.SEI.CMU.EDU and ask to be added to the mailing
list. The material sent to this list also appears in the
USENET newsgroup "comp.security.announce". Past advisories
are available for anonymous FTP from the host
CERT.SEI.CMU.EDU. The 24-hour hotline number is (412) 268-
7090.

The CERT/CC also maintains a CERT-TOOLS list to encourage
the exchange of information on tools and techniques that
increase the secure operation of Internet systems. The
CERT/CC does not review or endorse the tools described on
the list. To subscribe, send a message to CERT-TOOLS-
REQUEST@CERT.SEI.CMU.EDU and ask to be added to the mailing
list.

The CERT/CC maintains other generally useful security
information for anonymous FTP from CERT.SEI.CMU.EDU. Get
the README file for a list of what is available.

For more information, contact:

   CERT
   Software Engineering Institute
   Carnegie Mellon University
   Pittsburgh, PA 15213-3890

   (412) 268-7090
   cert@cert.sei.cmu.edu.

3.9.7.3.2 DDN Security Coordination Center

   For DDN users, the Security Coordination Center (SCC) serves
   a function similar to CERT. The SCC is the DDN's clearing-
   house for host/user security problems and fixes, and works
   with the DDN Network Security Officer. The SCC also
   distributes the DDN Security Bulletin, which communicates
   information on network and host security exposures, fixes,
   and concerns to security and management personnel at DDN
   facilities. It is available online, via kermit or anonymous
   FTP, from the host NIC.DDN.MIL, in SCC:DDN-SECURITY-yy-
   nn.TXT (where "yy" is the year and "nn" is the bulletin
   number). The SCC provides immediate assistance with DDN-
   related host security problems; call (800) 235-3155 (6:00
   a.m. to 5:00 p.m. Pacific Time) or send email to
   SCC@NIC.DDN.MIL. For 24 hour coverage, call the MILNET
   Trouble Desk (800) 451-7413 or AUTOVON 231-1713.

3.9.7.3.3 NIST Computer Security Resource and Response Center

   The National Institute of Standards and Technology (NIST)
   has responsibility within the U.S. Federal Government for
   computer science and technology activities. NIST has played
   a strong role in organizing the CERT System and is now
   serving as the CERT System Secretariat. NIST also operates
   a Computer Security Resource and Response Center (CSRC) to
   provide help and information regarding computer security
   events and incidents, as well as to raise awareness about
   computer security vulnerabilities.

   The CSRC team operates a 24-hour hotline, at (301) 975-5200.
   For individuals with access to the Internet, on-line
   publications and computer security information can be
   obtained via anonymous FTP from the host CSRC.NCSL.NIST.GOV
   (129.6.48.87). NIST also operates a personal computer
   bulletin board that contains information regarding computer
   viruses as well as other aspects of computer security. To
   access this board, set your modem to 300/1200/2400 BPS, 1
   stop bit, no parity, and 8-bit characters, and call (301)
   948-5717. All users are given full access to the board
   immediately upon registering.

   NIST has produced several special publications related to
   computer security and computer viruses in particular; some
   of these publications are downloadable. For further
   information, contact NIST at the following address:

Computer Security Resource and Response Center
A-216 Technology
Gaithersburg, MD 20899
Telephone: (301) 975-3359
Electronic Mail: CSRC@nist.gov

3.9.7.3.4 DOE Computer Incident Advisory Capability (CIAC)

CIAC is the Department of Energy's (DOE's) Computer Incident
Advisory Capability. CIAC is a four-person team of computer
scientists from Lawrence Livermore National Laboratory
(LLNL) charged with the primary responsibility of assisting
DOE sites faced with computer security incidents (e.g.,
intruder attacks, virus infections, worm attacks, etc.).
This capability is available to DOE sites on a 24-hour-a-day
basis.

CIAC was formed to provide a centralized response capability
(including technical assistance), to keep sites informed of
current events, to deal proactively with computer security
issues, and to maintain liaisons with other response teams
and agencies. CIAC's charter is to assist sites (through
direct technical assistance, providing information, or
referring inquiries to other technical experts), serve as a
clearinghouse for information about threats/known
incidents/vulnerabilities, develop guidelines for incident
handling, develop software for responding to
events/incidents, analyze events and trends, conduct
training and awareness activities, and alert and advise
sites about vulnerabilities and potential attacks.

CIAC's business hours phone number is (415) 422-8193 or FTS
532-8193. CIAC's e-mail address is CIAC@TIGER.LLNL.GOV.

3.9.7.3.5 NASA Ames Computer Network Security Response Team

The Computer Network Security Response Team (CNSRT) is NASA
Ames Research Center's local version of the DARPA CERT.
Formed in August of 1989, the team has a constituency that
is primarily Ames users, but it is also involved in
assisting other NASA Centers and federal agencies. CNSRT
maintains liaisons with the DOE's CIAC team and the DARPA
CERT. It is also a charter member of the CERT System. The
team may be reached by 24 hour pager at (415) 694-0571, or
by electronic mail to CNSRT@AMES.ARC.NASA.GOV.

3.9.7.4 DDN Management Bulletins

The DDN Management Bulletin is distributed electronically by
the DDN NIC under contract to the Defense Communications Agency
(DCA). It is a means of communicating official policy,
procedures, and other information of concern to management
personnel at DDN facilities.

The DDN Security Bulletin is distributed electronically by the
DDN SCC, also under contract to DCA, as a means of
communicating information on network and host security
exposures, fixes, and concerns to security and management
personnel at DDN facilities.

Anyone may join the mailing lists for these two bulletins by
sending a message to NIC@NIC.DDN.MIL and asking to be placed on
the mailing lists. These messages are also posted to the
USENET newsgroup "ddn.mgt-bulletin". For additional
information, see section 8.7.

3.9.7.5 System Administration List

The SYSADM-LIST is a list pertaining exclusively to UNIX system
administration. Mail requests to be added to the list to
SYSADM-LIST-REQUEST@SYSADMIN.COM.

3.9.7.6 Vendor Specific System Lists

The SUN-SPOTS and SUN-MANAGERS lists are discussion groups for
users and administrators of systems supplied by Sun
Microsystems. SUN-SPOTS is a fairly general list, discussing
everything from hardware configurations to simple UNIX
questions. To subscribe, send a message to SUN-SPOTS-
REQUEST@RICE.EDU. This list is also available in the USENET
newsgroup "comp.sys.sun". SUN-MANAGERS is a discussion list
for Sun system administrators and covers all aspects of Sun
system administration. To subscribe, send a message to SUN-
MANAGERS-REQUEST@EECS.NWU.EDU.

The APOLLO list discusses the HP/Apollo system and its
software. To subscribe, send a message to APOLLO-
REQUEST@UMIX.CC.UMICH.EDU. APOLLO-L is a similar list which
can be subscribed to by sending

    SUB APOLLO-L your full name

to LISTSERV%UMRVMB.BITNET@VM1.NODAK.EDU.

HPMINI-L pertains to the Hewlett-Packard 9000 series and HP/UX
operating system. To subscribe, send

    SUB HPMINI-L your full name

to LISTSERV%UAFSYSB.BITNET@VM1.NODAK.EDU.

INFO-IBMPC discusses IBM PCs and compatibles, as well as MS-
DOS. To subscribe, send a note to INFO-IBMPC-REQUEST@WSMR-
SIMTEL20.ARMY.MIL.

There are numerous other mailing lists for nearly every popular
computer or workstation in use today. For a complete list,
obtain the file "netinfo/interest-groups" via anonymous FTP
from the host FTP.NISC.SRI.COM.

3.9.7.7 Professional Societies and Journals

The IEEE Technical Committee on Security & Privacy publishes a
quarterly magazine, "CIPHER".

    IEEE Computer Society,
    1730 Massachusetts Ave. N.W.
    Washington, DC   2036-1903

The ACM SigSAC (Special Interest Group on Security, Audit, and
Controls) publishes a quarterly magazine, "SIGSAC Review".

    Association for Computing Machinery
    11 West 42nd St.
    New York, N.Y.   10036

The Information Systems Security Association publishes a
quarterly magazine called "ISSA Access".

    Information Systems Security Association
    P.O. Box 9457
    Newport Beach, CA   92658

"Computers and Security" is an "international journal for the
professional involved with computer security, audit and
control, and data integrity."

$266/year, 8 issues (1990)

Elsevier Advanced Technology
Journal Information Center
655 Avenue of the Americas
New York, NY 10010

The "Data Security Letter" is published "to help data security
professionals by providing inside information and knowledgable
analysis of developments in computer and communications
security."

$690/year, 9 issues (1990)

Data Security Letter
P.O. Box 1593
Palo Alto, CA 94302

3.9.8 Problem Reporting Tools

3.9.8.1 Auditing

Auditing is an important tool that can be used to enhance the
security of your installation. Not only does it give you a
means of identifying who has accessed your system (and may have
done something to it) but it also gives you an indication of
how your system is being used (or abused) by authorized users
and attackers alike. In addition, the audit trail
traditionally kept by computer systems can become an invaluable
piece of evidence should your system be penetrated.

3.9.8.1.1 Verify Security

An audit trail shows how the system is being used from day
to day. Depending upon how your site audit log is
configured, your log files should show a range of access
attempts that can show what normal system usage should look
like. Deviation from that normal usage could be the result
of penetration from an outside source using an old or stale
user account. Observing a deviation in logins, for example,
could be your first indication that something unusual is
happening.

3.9.8.1.2 Verify Software Configurations

One of the ruses used by attackers to gain access to a
system is by the insertion of a so-called Trojan Horse
program. A Trojan Horse program can be a program that does

something useful, or merely something interesting. It
always does something unexpected, like steal passwords or
copy files without your knowledge [25]. Imagine a Trojan
login program that prompts for username and password in the
usual way, but also writes that information to a special
file that the attacker can come back and read at will.
Imagine a Trojan Editor program that, despite the file
permissions you have given your files, makes copies of
everything in your directory space without you knowing about
it.

This points out the need for configuration management of the
software that runs on a system, not as it is being
developed, but as it is in actual operation. Techniques for
doing this range from checking each command every time it is
executed against some criterion (such as a cryptoseal,
described above) or merely checking the date and time stamp
of the executable. Another technique might be to check each
command in batch mode at midnight.

3.9.8.2 Tools

COPS is a security tool for system administrators that checks
for numerous common security problems on UNIX systems [27].
COPS is a collection of shell scripts and C programs that can
easily be run on almost any UNIX variant. Among other things,
it checks the following items and sends the results to the
system administrator:

- Checks "/dev/kmem" and other devices for world
  read/writability.

- Checks special or important files and directories for
  "bad" modes (world writable, etc.).

- Checks for easily-guessed passwords.

- Checks for duplicate user ids, invalid fields in the
  password file, etc..

- Checks for duplicate group ids, invalid fields in the
  group file, etc..

- Checks all users' home directories and their ".cshrc",
  ".login", ".profile", and ".rhosts" files for security
  problems.

- Checks all commands in the "/etc/rc" files and "cron"

         files for world writability.

      - Checks for bad "root" paths, NFS file systems exported
        to the world, etc..

      - Includes an expert system that checks to see if a given
        user (usually "root") can be compromised, given that
        certain rules are true.

      - Checks for changes in the setuid status of programs on the
        system.

    The COPS package is available from the "comp.sources.unix"
    archive on "ftp.uu.net", and also from the UNIX-SW repository
    on the MILNET host "wsmr-simtel20.army.mil".

  3.9.9 Communication Among Administrators

    3.9.9.1 Secure Operating Systems

    The following list of products and vendors is adapted from the
    National Computer Security Center's (NCSC) Evaluated Products
    List. They represent those companies who have either received
    an evaluation from the NCSC or are in the process of a product
    evaluation. This list is not complete, but it is
    representative of those operating systems and add on components
    available in the commercial marketplace.

    For a more detailed listing of the current products appearing
    in the NCSC EPL, contact the NCSC at:

       National Computer Security Center
       9800 Savage Road
       Fort George G. Meade, MD 20755-6000
       (301) 859-4458

| | | Version | Evaluation |
Evaluated Product	Vendor	Evaluated	Class
Secure Communications Processor (SCOMP)	Honeywell Information Systems, Inc.	2.1	A1
Multics	Honeywell Information Systems, Inc.	MR11.0	B2
System V/MLS 1.1.2 on UNIX System V 3.1.1 on AT&T 3B2/500and 3B2/600	AT&T	1.1.2	B1
OS 1100	Unisys Corp.	Security Release 1	B1
MPE V/E	Hewlett-Packard Computer Systems Division	G.03.04	C2
AOS/VS on MV/ECLIPSE series	Data General Corp.	7.60	C2
VM/SP or VM/SP HPO with CMS, RACF, DIRMAINT, VMTAPE-MS, ISPF	IBM Corp.	5	C2
MVS/XA with RACF	IBM Corp.	2.2,2.3	C2
AX/VMS	Digital Equipment Corp.	4.3	C2
NOS	Control Data Corp.	NOS Security Eval Product	C2
TOP SECRET	CGA Software Products Group, Inc.	3.0/163	C2
Access Control Facility 2	SKK, Inc.	3.1.3	C2
UTX/32S	Gould, Inc. Computer Systems Division	1.0	C2
A Series MCP/AS with InfoGuard Security Enhancements	Unisys Corp.	3.7	C2
Primos	Prime Computer, Inc.	21.0.1DODC2A	C2
Resource Access Control Facility (RACF)	IBM Corp.	1.5	C1

| | | Version | Candidate |
Candidate Product	Vendor	Evaluated	Class
Boeing MLS LAN	Boeing Aerospace		A1 M1
Trusted XENIX	Trusted Information		
	Systems, Inc.		B2
VSLAN	VERDIX Corp.		B2
System V/MLS	AT&T		B1
VM/SP with RACF	IBM Corp.	5/1.8.2	C2
Wang SVS/OS with CAP	Wang Laboratories, Inc.	1.0	C2

3.9.9.2 Obtaining Fixes for Known Problems

It goes without saying that computer systems have bugs. Even
operating systems, upon which we depend for protection of our
data, have bugs. And since there are bugs, things can be
broken, both maliciously and accidentally. It is important
that whenever bugs are discovered, a should fix be identified
and implemented as soon as possible. This should minimize any
exposure caused by the bug in the first place.

A corollary to the bug problem is: from whom do I obtain the
fixes? Most systems have some support from the manufacturer or
supplier. Fixes coming from that source tend to be implemented
quickly after receipt. Fixes for some problems are often
posted on the network and are left to the system administrators
to incorporate as they can. The problem is that one wants to
have faith that the fix will close the hole and not introduce
any others. We will tend to trust that the manufacturer's
fixes are better than those that are posted on the net.

3.9.9.3 Sun Customer Warning System

Sun Microsystems has established a Customer Warning System
(CWS) for handling security incidents. This is a formal
process which includes:

- Having a well advertised point of contact in Sun
  for reporting security problems.
- Pro-actively alerting customers of worms, viruses,
  or other security holes that could affect their systems.
- Distributing the patch (or work-around) as quickly
  as possible.

They have created an electronic mail address, SECURITY-
ALERT@SUN.COM, which will enable customers to report security
problems. A voice-mail backup is available at (415) 688-9081.
A "Security Contact" can be designated by each customer site;
this person will be contacted by Sun in case of any new
security problems. For more information, contact your Sun
representative.

3.9.9.4 Trusted Archive Servers

Several sites on the Internet maintain large repositories of
public-domain and freely distributable software, and make this
material available for anonymous FTP. This section describes
some of the larger repositories. Note that none of these
servers implements secure checksums or anything else
guaranteeing the integrity of their data. Thus, the notion of
"trust" should be taken as a somewhat limited definition.

3.9.9.4.1 Sun Fixes on UUNET

Sun Microsystems has contracted with UUNET Communications
Services, Inc., to make fixes for bugs in Sun software
available via anonymous FTP. You can access these fixes by
using the "ftp" command to connect to the host FTP.UU.NET.
Then change into the directory "sun-dist/security", and
obtain a directory listing. The file "README" contains a
brief description of what each file in this directory
contains, and what is required to install the fix.

3.9.9.4.2 Berkeley Fixes

The University of California at Berkeley also makes fixes
available via anonymous FTP; these fixes pertain primarily
to the current release of BSD UNIX (currently, release 4.3).
However, even if you are not running their software, these
fixes are still important, since many vendors (Sun, DEC,
Sequent, etc.) base their software on the Berkeley releases.

The Berkeley fixes are available for anonymous FTP from the
host UCBARPA.BERKELEY.EDU in the directory "4.3/ucb-fixes".
The file "INDEX" in this directory describes what each file
contains. They are also available from UUNET (see section
3.9.9.4.3).

Berkeley also distributes new versions of "sendmail" and
"named" from this machine. New versions of these commands
are stored in the "4.3" directory, usually in the files
"sendmail.tar.Z" and "bind.tar.Z", respectively.

3.9.9.4.3 Simtel-20 and UUNET

The two largest general-purpose software repositories on the
Internet are the hosts WSMR-SIMTEL20.ARMY.MIL and
FTP.UU.NET.

WSMR-SIMTEL20.ARMY.MIL is a TOPS-20 machine operated by the
U.S. Army at White Sands Missile Range (WSMR), New Mexico.
The directory "pd2:<unix-c>" contains a large amount of UNIX
software, primarily taken from the "comp.sources"
newsgroups. The directories "pd1:<msdos>" and
"pd2:<msdos2>" contains software for IBM PC systems, and
"pd3:<macintosh>" contains software for the Apple Macintosh.

FTP.UU.NET is operated by UUNET Communications Services,
Inc. in Falls Church, Virginia. This company sells Internet
and USENET access to sites all over the country (and
internationally). The software posted to the following
USENET source newsgroups is stored here, in directories of
the same name:

    comp.sources.games
    comp.sources.misc
    comp.sources.sun
    comp.sources.unix
    comp.sources.x

Numerous other distributions, such as all the freely
distributable Berkeley UNIX source code, Internet Request
for Comments (RFCs), and so on are also stored on this
system.

3.9.9.4.4 Vendors

Many vendors make fixes for bugs in their software available
electronically, either via mailing lists or via anonymous
FTP. You should contact your vendor to find out if they
offer this service, and if so, how to access it. Some
vendors that offer these services include Sun Microsystems
(see above), Digital Equipment Corporation (DEC), the
University of California at Berkeley (see above), and Apple
Computer [5, CURRY].

4. Types of Security Procedures

4.1 System Security Audits

   Most businesses undergo some sort of annual financial auditing as a
   regular part of their business life. Security audits are an
   important part of running any computing environment. Part of the
   security audit should be a review of any policies that concern system
   security, as well as the mechanisms that are put in place to enforce
   them.

   4.1.1   Organize Scheduled Drills

      Although not something that would be done each day or week,
      scheduled drills may be conducted to determine if the procedures
      defined are adequate for the threat to be countered. If your
      major threat is one of natural disaster, then a drill would be
      conducted to verify your backup and recovery mechanisms. On the
      other hand, if your greatest threat is from external intruders
      attempting to penetrate your system, a drill might be conducted to
      actually try a penetration to observe the effect of the policies.

      Drills are a valuable way to test that your policies and
      procedures are effective. On the other hand, drills can be time-
      consuming and disruptive to normal operations. It is important to
      weigh the benefits of the drills against the possible time loss
      which may be associated with them.

   4.1.2 Test Procedures

      If the choice is made to not to use scheduled drills to examine
      your entire security procedure at one time, it is important to
      test individual procedures frequently. Examine your backup
      procedure to make sure you can recover data from the tapes. Check
      log files to be sure that information which is supposed to be
      logged to them is being logged to them, etc..

      When a security audit is mandated, great care should be used in
      devising tests of the security policy. It is important to clearly
      identify what is being tested, how the test will be conducted, and
      results expected from the test. This should all be documented and
      included in or as an adjunct to the security policy document
      itself.

      It is important to test all aspects of the security policy, both
      procedural and automated, with a particular emphasis on the
      automated mechanisms used to enforce the policy. Tests should be
      defined to ensure a comprehensive examination of policy features,

that is, if a test is defined to examine the user logon process,
it should be explicitly stated that both valid and invalid user
names and passwords will be used to demonstrate proper operation
of the logon program.

Keep in mind that there is a limit to the reasonableness of tests.
The purpose of testing is to ensure confidence that the security
policy is being correctly enforced, and not to "prove" the
absoluteness of the system or policy. The goal should be to
obtain some assurance that the reasonable and credible controls
imposed by your security policy are adequate.

4.2 Account Management Procedures

   Procedures to manage accounts are important in preventing
   unauthorized access to your system. It is necessary to decide
   several things: Who may have an account on the system? How long may
   someone have an account without renewing his or her request? How do
   old accounts get removed from the system? The answers to all these
   questions should be explicitly set out in the policy.

   In addition to deciding who may use a system, it may be important to
   determine what each user may use the system for (is personal use
   allowed, for example). If you are connected to an outside network,
   your site or the network management may have rules about what the
   network may be used for. Therefore, it is important for any security
   policy to define an adequate account management procedure for both
   administrators and users. Typically, the system administrator would
   be responsible for creating and deleting user accounts and generally
   maintaining overall control of system use. To some degree, account
   management is also the responsibility of each system user in the
   sense that the user should observe any system messages and events
   that may be indicative of a policy violation. For example, a message
   at logon that indicates the date and time of the last logon should be
   reported by the user if it indicates an unreasonable time of last
   logon.

4.3 Password Management Procedures

   A policy on password management may be important if your site wishes
   to enforce secure passwords. These procedures may range from asking
   or forcing users to change their passwords occasionally to actively
   attempting to break users' passwords and then informing the user of
   how easy it was to do. Another part of password management policy
   covers who may distribute passwords - can users give their passwords
   to other users?

   Section 2.3 discusses some of the policy issues that need to be

decided for proper password management. Regardless of the policies,
password management procedures need to be carefully setup to avoid
disclosing passwords. The choice of initial passwords for accounts
is critical. In some cases, users may never login to activate an
account; thus, the choice of the initial password should not be
easily guessed. Default passwords should never be assigned to
accounts: always create new passwords for each user. If there are
any printed lists of passwords, these should be kept off-line in
secure locations; better yet, don't list passwords.

4.3.1 Password Selection

Perhaps the most vulnerable part of any computer system is the
account password. Any computer system, no matter how secure it is
from network or dial-up attack, Trojan horse programs, and so on,
can be fully exploited by an intruder if he or she can gain access
via a poorly chosen password. It is important to define a good
set of rules for password selection, and distribute these rules to
all users. If possible, the software which sets user passwords
should be modified to enforce as many of the rules as possible.

A sample set of guidelines for password selection is shown below:

- DON'T use your login name in any form (as-is,
  reversed, capitalized, doubled, etc.).

- DON'T use your first, middle, or last name in any form.

- DON'T use your spouse's or child's name.

- DON'T use other information easily obtained about you.
  This includes license plate numbers, telephone numbers,
  social security numbers, the make of your automobile,
  the name of the street you live on, etc..

- DON'T use a password of all digits, or all the same
  letter.

- DON'T use a word contained in English or foreign
  language dictionaries, spelling lists, or other
  lists of words.

- DON'T use a password shorter than six characters.

- DO use a password with mixed-case alphabetics.

- DO use a password with non-alphabetic characters (digits
  or punctuation).

- DO use a password that is easy to remember, so you don't
  have to write it down.

- DO use a password that you can type quickly, without
  having to look at the keyboard.

Methods of selecting a password which adheres to these guidelines
include:

- Choose a line or two from a song or poem, and use the
  first letter of each word.

- Alternate between one consonant and one or two vowels, up
  to seven or eight characters. This provides nonsense
  words which are usually pronounceable, and thus easily
  remembered.

- Choose two short words and concatenate them together with
  a punctuation character between them.

Users should also be told to change their password periodically,
usually every three to six months. This makes sure that an
intruder who has guessed a password will eventually lose access,
as well as invalidating any list of passwords he/she may have
obtained. Many systems enable the system administrator to force
users to change their passwords after an expiration period; this
software should be enabled if your system supports it [5, CURRY].

Some systems provide software which forces users to change their
passwords on a regular basis. Many of these systems also include
password generators which provide the user with a set of passwords
to choose from. The user is not permitted to make up his or her
own password. There are arguments both for and against systems
such as these. On the one hand, by using generated passwords,
users are prevented from selecting insecure passwords. On the
other hand, unless the generator is good at making up easy to
remember passwords, users will begin writing them down in order to
remember them.

4.3.2 Procedures for Changing Passwords

How password changes are handled is important to keeping passwords
secure. Ideally, users should be able to change their own
passwords on-line. (Note that password changing programs are a
favorite target of intruders. See section 4.4 on configuration
management for further information.)

However, there are exception cases which must be handled

carefully. Users may forget passwords and not be able to get onto
the system. The standard procedure is to assign the user a new
password. Care should be taken to make sure that the real person
is requesting the change and gets the new password. One common
trick used by intruders is to call or message to a system
administrator and request a new password. Some external form of
verification should be used before the password is assigned. At
some sites, users are required to show up in person with ID.

There may also be times when many passwords need to be changed.
If a system is compromised by an intruder, the intruder may be
able to steal a password file and take it off the system. Under
these circumstances, one course of action is to change all
passwords on the system. Your site should have procedures for how
this can be done quickly and efficiently. What course you choose
may depend on the urgency of the problem. In the case of a known
attack with damage, you may choose to forcibly disable all
accounts and assign users new passwords before they come back onto
the system. In some places, users are sent a message telling them
that they should change their passwords, perhaps within a certain
time period. If the password isn't changed before the time period
expires, the account is locked.

Users should be aware of what the standard procedure is for
passwords when a security event has occurred. One well-known
spoof reported by the Computer Emergency Response Team (CERT)
involved messages sent to users, supposedly from local system
administrators, requesting them to immediately change their
password to a new value provided in the message [24]. These
messages were not from the administrators, but from intruders
trying to steal accounts. Users should be warned to immediately
report any suspicious requests such as this to site
administrators.

4.4 Configuration Management Procedures

   Configuration management is generally applied to the software
   development process. However, it is certainly applicable in a
   operational sense as well. Consider that the since many of the
   system level programs are intended to enforce the security policy, it
   is important that these be "known" as correct. That is, one should
   not allow system level programs (such as the operating system, etc.)
   to be changed arbitrarily. At very least, the procedures should
   state who is authorized to make changes to systems, under what
   circumstances, and how the changes should be documented.

   In some environments, configuration management is also desirable as
   applied to physical configuration of equipment. Maintaining valid

and authorized hardware configuration should be given due
consideration in your security policy.

   4.4.1 Non-Standard Configurations

      Occasionally, it may be beneficial to have a slightly non-standard
      configuration in order to thwart the "standard" attacks used by
      some intruders. The non-standard parts of the configuration might
      include different password encryption algorithms, different
      configuration file locations, and rewritten or functionally
      limited system commands.

      Non-standard configurations, however, also have their drawbacks.
      By changing the "standard" system, these modifications make
      software maintenance more difficult by requiring extra
      documentation to be written, software modification after operating
      system upgrades, and, usually, someone with special knowledge of
      the changes.

      Because of the drawbacks of non-standard configurations, they are
      often only used in environments with a "firewall" machine (see
      section 3.9.1). The firewall machine is modified in non-standard
      ways since it is susceptible to attack, while internal systems
      behind the firewall are left in their standard configurations.

5. Incident Handling

5.1 Overview

   This section of the document will supply some guidance to be applied
   when a computer security event is in progress on a machine, network,
   site, or multi-site environment. The operative philosophy in the
   event of a breach of computer security, whether it be an external
   intruder attack or a disgruntled employee, is to plan for adverse
   events in advance. There is no substitute for creating contingency
   plans for the types of events described above.

   Traditional computer security, while quite important in the overall
   site security plan, usually falls heavily on protecting systems from
   attack, and perhaps monitoring systems to detect attacks. Little
   attention is usually paid for how to actually handle the attack when
   it 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.

   5.1.1 Have a Plan to Follow in Case of an Incident

      Part of handling an incident is being prepared to respond before
      the incident occurs. This includes establishing a suitable level
      of protections, so that if the incident becomes severe, the damage
      which can occur is limited. Protection includes preparing
      incident handling guidelines or a contingency response 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. Second, part of
      protection is preparing a method of notification, so you will know
      who to call and the relevant phone numbers. It is important, for
      example, to conduct "dry runs," in which your computer security
      personnel, system administrators, and managers simulate handling
      an incident.

      Learning to respond efficiently to an incident is important for
      numerous reasons. The most important benefit is directly to human
      beings--preventing loss of human life. Some computing systems are
      life critical systems, systems on which human life depends (e.g.,
      by controlling some aspect of life-support in a hospital or
      assisting air traffic controllers).

      An important but often overlooked benefit is an economic one.
      Having both technical and managerial personnel respond to an
      incident requires considerable resources, resources which could be
      utilized more profitably if an incident did not require their
      services. If these personnel are trained to handle an incident
      efficiently, less of their time is required to deal with that
      incident.

      A third benefit is protecting classified, sensitive, or
      proprietary information. One of the major dangers of a computer
      security incident is that information may be irrecoverable.
      Efficient incident handling minimizes this danger. When
      classified information is involved, other government regulations
      may apply and must be integrated into any plan for incident
      handling.

      A fourth 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 sued 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 damage to systems, or if the patches or
workarounds themselves damage systems. Knowing about operating
system vulnerabilities and patterns of attacks and then taking
appropriate measures is critical to circumventing possible legal
problems.

5.1.2 Order of Discussion in this Session Suggests an Order for
       a Plan

This chapter is arranged such that a list may be generated from
the Table of Contents to provide a starting point for creating a
policy for handling ongoing incidents. The main points to be
included in a policy for handling incidents are:

    o Overview (what are the goals and objectives in handling the
      incident).
    o Evaluation (how serious is the incident).
    o Notification (who should be notified about the incident).
    o Response (what should the response to the incident be).
    o Legal/Investigative (what are the legal and prosecutorial
      implications of the incident).
    o Documentation Logs (what records should be kept from before,
      during, and after the incident).

Each of these points is important in an overall plan for handling
incidents. The remainder of this chapter will detail the issues
involved in each of these topics, and provide some guidance as to
what should be included in a site policy for handling incidents.

5.1.3 Possible Goals and Incentives for Efficient Incident
       Handling

As in any set of pre-planned procedures, attention must be placed
on a set of goals to be obtained in handling an incident. These
goals will be placed in order of importance depending on the site,
but one such set of goals might be:

    Assure integrity of (life) critical systems.
    Maintain and restore data.
    Maintain and restore service.
    Figure out how it happened.
    Avoid escalation and further incidents.
    Avoid negative publicity.
    Find out who did it.
    Punish the attackers.

It is important to prioritize 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 serve as a starting point for
defining an organization's response:

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

    o Priority two -- protect classified and/or sensitive
      data (as regulated by your site or by government
      regulations).

    o Priority three -- protect other data, including
      proprietary, scientific, managerial and other data,
      because loss of data is costly in terms of resources.

    o 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.

    o Priority five -- minimize disruption of computing
      resources; it is better in many cases to shut a system
      down or disconnect from a network than to risk damage
      to data or systems.

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; the
loss or compromise of data (especially classified data), however,
is usually not an acceptable outcome under any circumstances.

Part of handling an incident is being prepared to respond before
the incident occurs. This includes establishing a suitable level
of protections so that if the incident becomes severe, the damage
which can occur is limited. Protection includes preparing
incident handling guidelines or a contingency response plan for
your organization or site. Written plans eliminate much of the
ambiguity which occurs during an incident, and will lead to a more
appropriate and thorough set of responses. Second, part of
protection is preparing a method of notification so you will know
who to call and how to contact them. For example, every member of

the Department of Energy's CIAC Team carries a card with every
other team member's work and home phone numbers, as well as pager
numbers. Third, your organization or site should establish backup
procedures for every machine and system. Having backups
eliminates much of the threat of even a severe incident, since
backups preclude serious data loss. Fourth, you should set up
secure systems. This involves eliminating vulnerabilities,
establishing an effective password policy, and other procedures,
all of which will be explained later in this document. Finally,
conducting training activities is part of protection. It is
important, for example, to conduct "dry runs," in which your
computer security personnel, system administrators, and managers
simulate handling an incident.

   5.1.4 Local Policies and Regulations Providing Guidance

      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 your site makes about how it responds to incidents
      (as discussed in sections 2.4 and 2.5) 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.

      Section 5.5 also notes that if any legal action is planned, there
      are specific guidelines that must be followed to make sure that
      any information collected can be used as evidence.

5.2 Evaluation

   5.2.1 Is It Real?

      This stage involves determining the exact problem. Of course
      many, if not most, signs often associated with virus infections,
      system intrusions, etc., are simply anomalies such as hardware
      failures. 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. For example, widely available
      software packages can greatly assist someone who thinks there may
      be a virus in a Macintosh computer. 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 do more good in identifying the problem and
any source of attack than most other actions which can be taken at
this stage. 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 which
deserve special attention:

- o System crashes.
- o New user accounts (e.g., the account RUMPLESTILTSKIN
  has unexplainedly been created), or high activity on
  an account that has had virtually no activity for
  months.
- o New files (usually with novel or strange file names,
  such as data.xx or k).
- o Accounting discrepancies (e.g., in a UNIX system you
  might notice that the accounting file called
  /usr/admin/lastlog has shrunk, something that should
  make you very suspicious that there may be an
  intruder).
- o Changes in file lengths or dates (e.g., a user should
  be suspicious if he/she observes that the .EXE files in
  an MS DOS computer have unexplainedly grown
  by over 1800 bytes).
- o Attempts to write to system (e.g., a system manager
  notices that a privileged user in a VMS system is
  attempting to alter RIGHTSLIST.DAT).
- o Data modification or deletion (e.g., files start to
  disappear).
- o Denial of service (e.g., a system manager and all
  other users become locked out of a UNIX system, which
  has been changed to single user mode).
- o Unexplained, poor system performance (e.g., system
  response time becomes unusually slow).
- o Anomalies (e.g., "GOTCHA" is displayed on a display
  terminal or there are frequent unexplained "beeps").
- o Suspicious probes (e.g., there are numerous
  unsuccessful login attempts from another node).
- o Suspicious browsing (e.g., someone becomes a root user
  on a UNIX system and accesses file after file in one
  user's account, then another's).

None of these indications is absolute "proof" that an incident is

occurring, nor are all of these indications normally observed when
an incident occurs. If you observe any of these indications,
however, it is important to suspect that an incident might be
occurring, and act accordingly. There is no formula for
determining with 100 percent accuracy that an incident is
occurring (possible exception: when a virus detection package
indicates that your machine has the nVIR virus and you confirm
this by examining contents of the nVIR resource in your Macintosh
computer, you can be very certain that your machine is infected).
It is best at this point to collaborate with other technical and
computer security personnel to make a decision as a group about
whether an incident is occurring.

5.2.2 Scope

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. In addition, the impact of an incident will
determine its priority in allocating resources to deal with the
event. Without an indication of the scope and impact of the
event, it is difficult to determine a correct response.

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 are:

   o Is this a multi-site incident?
   o Are many computers at your site effected by this
     incident?
   o Is sensitive information involved?
   o What is the entry point of the incident (network,
     phone line, local terminal, etc.)?
   o Is the press involved?
   o What is the potential damage of the incident?
   o What is the estimated time to close out the incident?
   o What resources could be required
     to handle the incident?

5.3 Possible Types of Notification

When you have confirmed that an incident is occurring, the
appropriate personnel must be notified. Who and how this
notification is achieved is very important in keeping the event under
control both from a technical and emotional standpoint.

5.3.1 Explicit

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, or fax) provides
information about the incident that is clear, concise, and fully
qualified. When you are notifying others that will help you to
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 section 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 other information that would help
them resolve a part of the incident.

5.3.2 Factual

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. This is especially true when the press is
involved. When an incident severe enough to gain press attention
is ongoing, it is likely that any false information you provide
will not be substantiated by other sources. This will reflect
badly on the site and may create enough ill-will between the site
and the press to damage the site's public relations.

5.3.3 Choice of Language

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 expectations of damage and negative outcomes of the incident.
It is important to remain calm both in written and spoken
notifications.

Another issue associated with the choice of language is the
notification to non-technical or off-site personnel. It is
important to accurately describe the incident without undue alarm
or confusing messages. 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 notifications cannot be underestimated and may
make the difference between handling the incident properly and
escalating to some higher level of damage.

5.3.4 Notification of Individuals

        o Point of Contact (POC) people (Technical, Administrative,
          Response Teams, Investigative, Legal, Vendors, Service
          providers), and which POCs are visible to whom.
        o Wider community (users).
        o Other sites that might be affected.

   Finally, there is the question of who should be notified during
   and after the incident. There are several classes of individuals
   that need to be considered for notification. These are the
   technical personnel, administration, appropriate response teams
   (such as CERT or CIAC), law enforcement, vendors, and other
   service providers. These issues are important for the central
   point of contact, since that is the person responsible for the
   actual notification of others (see section 5.3.6 for further
   information). A list of people in each of these categories is an
   important time saver for the POC during an incident. It is much
   more difficult to find an appropriate person during an incident
   when many urgent events are ongoing.

   In addition to the people responsible for handling part of the
   incident, there may be other sites affected by the incident (or
   perhaps simply at risk from the incident). A wider community of
   users may also benefit from knowledge of the incident. Often, a
   report of the incident once it is closed out is appropriate for
   publication to the wider user community.

5.3.5 Public Relations - Press Releases

   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. As a
contrast to this consideration, it was discussed above 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:

   o Keep the technical level of detail low. Detailed
     information about the incident may provide enough
     information for copy-cat events or even damage the
     site's ability to prosecute once the event is over.
   o 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.
   o 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.
   o Try not to be forced into a press interview before you are
     prepared. The popular press is famous for the "2am"
     interview, where the hope is to catch the interviewee off
     guard and obtain information otherwise not available.
   o 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.

5.3.6 Who Needs to Get Involved?

There now exists a number of incident response teams (IRTs) such
as the CERT and the CIAC. (See sections 3.9.7.3.1 and 3.9.7.3.4.)
Teams exists for many major government agencies and large
corporations. If such a team is available for your site, the
notification of this team should be of primary importance 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
to a single site, it is possible that the information available
through a response team could help in closing out the incident.

In setting up a site policy for incident handling, it may be
desirable to create an incident handling team (IHT), 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 IHTs.
Once an incident is under way, it is difficult to open a trusted

dialogue between other IHTs if none has existed before.

5.4 Response

A major topic still untouched here is how to actually respond to an
event. The response to an event will fall into the general
categories of containment, eradication, recovery, and follow-up.

Containment

The purpose of containment is to limit the extent of an attack.
For example, it is important to limit the spread of a worm attack
on a network as quickly as possible. An essential part of
containment is decision making (i.e., determining whether to shut
a system down, to disconnect from a network, to monitor system or
network activity, to set traps, to disable functions such as
remote file transfer on a UNIX system, etc.). Sometimes this
decision is trivial; shut the system down if the system is
classified or sensitive, or if proprietary information is at risk!
In other cases, it is worthwhile to risk having some damage to the
system if keeping the system up might enable you to identify an
intruder.

The third stage, containment, 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.
Finally, notification of cognizant authorities should occur during
this stage.

Eradication

Once an incident has been detected, it is important to first think
about containing the incident. Once the incident has been
contained, it is now time to eradicate the cause. Software may be
available to help you in this effort. For example, eradication
software is available to eliminate most viruses which infect small
systems. If any bogus files have been created, it is time to
delete them at this point. In the case of virus infections, it is
important to clean and reformat any disks containing infected
files. Finally, ensure that all backups are clean. Many systems
infected with viruses become periodically reinfected simply
because people do not systematically eradicate the virus from
backups.

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 the case of a network-based
attack, it is important to install patches for any operating
system vulnerability which was exploited.

Follow-up

One of the most important stages of responding to incidents is
also the most often omitted---the follow-up stage. This stage is
important because it helps those involved in handling the incident
develop a set of "lessons learned" (see section 6.3) to improve
future performance in such situations. This stage also provides
information which justifies an organization's computer security
effort to management, and yields information which may be
essential in legal proceedings.

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? A follow-up report is valuable
because it provides a reference to be used in case of other
similar incidents. Creating a formal chronology of events
(including time stamps) is also important for legal reasons.
Similarly, it is also important to as quickly obtain a monetary
estimate of the amount of damage the incident caused in terms of
any loss of software and files, 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 by the FBI, the U.S. Attorney General's
Office, etc..

5.4.1 What Will You Do?

o Restore control.
o Relation to policy.
o Which level of service is needed?
o Monitor activity.
o Constrain or shut down system.

5.4.2 Consider Designating a "Single Point of Contact"

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 "points of
contact" (POC) that are not pulling their efforts together. This
will only add to the confusion of the event, and will probably

lead to additional confusion and wasted or ineffective effort.

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

The point of contact must be a person with the technical expertise
to successfully coordinate the effort of the system managers and
users involved in monitoring and reacting to the attack. Often
the management structure of a site is such that the administrator
of a set of resources is not a technically competent person with
regard to handling the details of the operations of the computers,
but is ultimately responsible for the use of these resources.

Another important function of the POC is to maintain contact with
law enforcement and other external agencies (such as the CIA, DoD,
U.S. Army, or others) to assure that multi-agency involvement
occurs.

Finally, if legal action in the form of prosecution is involved,
the POC may be able to speak for the site in court. The
alternative is to have multiple witnesses that will be hard to
coordinate in a legal sense, and will weaken any case against the
attackers. A single POC may also be the single person in charge
of evidence collected, which will keep the number of people
accounting for evidence to a minimum. 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. The
section below (Legal/Investigative) will provide more details for
consideration on this topic.

5.5 Legal/Investigative

  5.5.1 Establishing Contacts with Investigative Agencies

    It is important to establish contacts with personnel from
    investigative agencies such as the FBI and Secret Service as soon
    as possible, for several reasons. Local law enforcement and local
    security offices or campus police organizations should also be
    informed when appropriate. A primary reason is that once a major
    attack is in progress, there is little time to call various
    personnel in 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 a court of
law. If you don't know in advance how to gather admissible
evidence, your efforts to collect evidence during an incident are
likely to be of no value to the investigative agency with which
you deal. 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 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 lots of effort.

On the other side, 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 anyone.

The balance between supporting investigative activity and limiting
liability is tricky; you'll need to consider the advice of your
council and the damage the intruder is causing (if any) in 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.

  5.5.2 Formal and Informal Legal Procedures

One of the most important considerations in dealing with
investigative agencies is verifying 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 information about incidents,
allowed unauthorized people into their systems, etc., because a
caller has masqueraded as an FBI or Secret Service agent. A
similar consideration is using a secure means of communication.
Because many network attackers can easily reroute 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 (e.g., the phones normally used in the business world)
are also frequent targets for tapping by network intruders, so be
careful!

There is no established set of rules for responding to an incident
when the U.S. Federal Government becomes involved. Except by
court order, no agency can force you to monitor, to disconnect
from the network, to avoid telephone contact with the suspected
attackers, etc.. As discussed in section 5.5.1, you should
consult the matter with your legal counsel, especially before
taking an action that your organization has never taken. The
particular agency involved may ask you to leave an attacked
machine on and to monitor activity on this machine, for example.
Your complying with this request will ensure continued cooperation
of the agency--usually the best route towards finding the source
of the network attacks and, ultimately, terminating these attacks.
Additionally, you may need some information or a favor from the
agency involved in the incident. You are likely to get what you
need only if you have been cooperative. Of particular importance
is avoiding unnecessary or unauthorized disclosure of information
about the incident, including any information furnished by the
agency involved. The trust between your site and the agency
hinges upon your ability to avoid compromising the case the agency
will build; keeping "tight lipped" is imperative.

Sometimes your needs and the needs of an investigative agency will
differ. Your site may want to get back to normal business by
closing an attack route, but the investigative agency may want you
to keep this route open. Similarly, your site may want to close a
compromised system down to avoid the possibility of negative
publicity, but again the investigative agency may want you to
continue monitoring. When there is such a conflict, there may be
a complex set of tradeoffs (e.g., interests of your site's
management, amount of resources you can devote to the problem,
jurisdictional boundaries, etc.). An important guiding principle
is related to what might be called "Internet citizenship" [22,
IAB89, 23] and its responsibilities. Your site can shut a system
down, and this will relieve you of the stress, resource demands,
and danger of negative exposure. The attacker, however, is likely
to simply move on to another system, temporarily leaving others
blind to the attacker's intention and actions until another path
of attack can be detected. Providing that there is no damage to
your systems and others, the most responsible course of action is
to cooperate with the participating agency by leaving your
compromised system on. This will allow monitoring (and,
ultimately, the possibility of terminating the source of the
threat to systems just like yours). On the other hand, if there
is damage to computers illegally accessed through your system, the
choice is more complicated: shutting down the intruder may prevent
further damage to systems, but might make it impossible to track
down the intruder. If there has been damage, the decision about
whether it is important to leave systems up to catch the intruder

should involve all the organizations effected. Further
complicating the issue of network responsibility is the
consideration that if you do not cooperate with the agency
involved, you will be less likely to receive help from that agency
in the future.

5.6 Documentation 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 good
   portion of information you obtain, requiring you to contact the
   source of information once again. This wastes yours and others'
   time, something you can ill afford. At the same time, recording
   details will provide evidence for prosecution efforts, providing the
   case moves in this direction. Documenting an incident also will 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 a follow-up analysis in which you can engage in
   a valuable "lessons learned" exercise.

   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:

      o All system events (audit records).
      o All actions you take (time tagged).
      o All phone 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 you initially
   suspect that an incident will result in prosecution or when an
   investigative agency becomes involved, you need to 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 who can store these copied pages in a secure place (e.g., a
   safe). When you submit information for storage, you should in return
   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.

6. Establishing Post-Incident Procedures

6.1 Overview

   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.

   All four steps should provide feedback to the site security policy
   committee, leading to prompt re-evaluation and amendment of the
   current policy.

6.2 Removing Vulnerabilities

   Removing all vulnerabilities once an incident has occurred is
   difficult. The key to removing vulnerabilities is knowledge and
   understanding of the breach. In some cases, it is prudent to remove
   all access or functionality as soon as possible, and then restore
   normal operation in limited stages. Bear in mind that removing all
   access while an incident is in progress will obviously notify all
   users, including the alleged problem users, that the administrators
   are aware of a problem; this may have a deleterious effect on an
   investigation. However, allowing an incident to continue may also
   open the likelihood of greater damage, loss, aggravation, or
   liability (civil or criminal).

   If it is determined that the breach occurred due to a flaw in the
   systems' hardware or software, the vendor (or supplier) and the CERT
   should be notified as soon as possible. Including relevant telephone
   numbers (also electronic mail addresses and fax numbers) in the site
   security policy is strongly recommended. To aid prompt
   acknowledgment and understanding of the problem, the flaw should be
   described in as much detail as possible, including details about how
   to exploit the flaw.

As soon as the breach has occurred, the entire system and all its
components should be considered suspect. System software is the most
probable target. Preparation is key to recovering from a possibly
tainted system. This includes checksumming all tapes from the vendor
using a checksum algorithm which (hopefully) is resistant to
tampering [10]. (See sections 3.9.4.1, 3.9.4.2.) Assuming original
vendor distribution tapes 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 tapes to
recover from; consider that the incident may have continued for
months or years before discovery, and that 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. At worst-case, restoration from
the original manufactures' media and a re-installation of the systems
will be the most prudent solution.

Review the lessons learned from the incident and always update the
policy and procedures to reflect changes necessitated by the
incident.

6.2.1 Assessing Damage

   Before cleanup can begin, the actual system damage must be
   discerned. This can be quite time consuming, but should lead into
   some of the insight as to the nature of the incident, and aid
   investigation and prosecution. It is best to compare previous
   backups or original tapes when possible; advance preparation is
   the key. 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.

6.2.2 Cleanup

   Once the damage has been assessed, it is necessary to develop a
   plan for system cleanup. 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.

   It may be necessary to go back to the original distributed tapes
   and recustomize the system. To facilitate this worst case

scenario, a record of the original systems setup and each
customization change should be kept current with each change to
the system.

6.2.3 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. 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 section 3.9.8.2 (e.g., COPS) as a start. Remember,
these tools don't replace continual system monitoring and good
systems administration procedures.

6.2.4 Keep a Security Log

As discussed in section 5.6, a security log can be most valuable
during this phase of removing vulnerabilities. There are two
considerations here; the first is to keep logs of the procedures
that have been used to make the system secure again. This should
include command procedures (e.g., shell scripts) that can be run
on a periodic basis to recheck the security. Second, keep logs of
important system events. These can be referenced when trying to
determine the extent of the damage of a given incident.

6.3 Capturing Lessons Learned

6.3.1 Understand the Lesson

After an incident, it is prudent to write a report describing the
incident, method of discovery, correction procedure, monitoring
procedure, and a summary of lesson learned. This will aid in the
clear understanding of the problem. Remember, it is difficult to
learn from an incident if you don't understand the source.

6.3.2 Resources

6.3.2.1 Other Security Devices, Methods

Security is a dynamic, not static process. Sites are dependent
on the nature of security available at each site, and the array
of devices and methods that will help promote security.
Keeping up with the security area of the computer industry and
their methods will assure a security manager of taking
advantage of the latest technology.

6.3.2.2 Repository of Books, Lists, Information Sources

Keep an on site collection of books, lists, information
sources, etc., as guides and references for securing the
system. Keep this collection up to date. Remember, as systems
change, so do security methods and problems.

6.3.2.3 Form a Subgroup

Form a subgroup of system administration personnel that will be
the core security staff. This will allow discussions of
security problems and multiple views of the site's security
issues. This subgroup can also act to develop the site
security policy and make suggested changes as necessary to
ensure site security.

6.4 Upgrading Policies and Procedures

6.4.1 Establish Mechanisms for Updating Policies, Procedures, and Tools

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.

6.4.2 Problem Reporting Procedures

A problem reporting procedure should be implemented to describe,
in detail, the incident and the solutions to the incident. Each
incident should be reviewed by the site security subgroup to allow
understanding of the incident with possible suggestions to the
site policy and procedures.

7. References

[1] Quarterman, J., "The Matrix: Computer Networks and Conferencing
    Systems Worldwide", Pg. 278, Digital Press, Bedford, MA, 1990.

[2] Brand, R., "Coping with the Threat of Computer Security
    Incidents: A Primer from Prevention through Recovery", R. Brand,
    available on-line from: cert.sei.cmu.edu:/pub/info/primer, 8 June
    1990.

[3] Fites, M., Kratz, P. and A. Brebner, "Control and Security of

Computer Information Systems", Computer Science Press, 1989.

[4] Johnson, D., and J. Podesta, "Formulating a Company Policy on
    Access to and Use and Disclosure of Electronic Mail on Company
    Computer Systems", Available from: The Electronic Mail
    Association (EMA) 1555 Wilson Blvd, Suite 555, Arlington VA
    22209, (703) 522-7111, 22 October 1990.

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

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

[7] Linn, J., "Privacy Enhancement for Internet Electronic Mail: Part
    I -- Message Encipherment and Authentication Procedures", RFC
    1113, IAB Privacy Task Force, August 1989.

[8] Kent, S., and J. Linn, "Privacy Enhancement for Internet
    Electronic Mail: Part II -- Certificate-Based Key Management",
    RFC 1114, IAB Privacy Task Force, August 1989.

[9] Linn, J., "Privacy Enhancement for Internet Electronic Mail: Part
    III -- Algorithms, Modes, and Identifiers", RFC 1115, IAB Privacy
    Task Force, August 1989.

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

[11] Postel, J., "Internet Protocol - DARPA Internet Program Protocol
     Specification", RFC 791, DARPA, September 1981.

[12] Postel, J., "Transmission Control Protocol - DARPA Internet
     Program Protocol Specification", RFC 793, DARPA, September 1981.

[13] Postel, J., "User Datagram Protocol", RFC 768, USC/Information
     Sciences Institute, 28 August 1980.

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

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

[16] Pfleeger, C., "Security in Computing", Prentice-Hall, Englewood

Cliffs, N.J., 1989.

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

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

[19] Postel, J., and J. Reynolds, "Telnet Protocol Specification", RFC
     854, USC/Information Sciences Institute, May 1983.

[20] Postel, J., and J. Reynolds, "File Transfer Protocol", RFC 959,
     USC/Information Sciences Institute, October 1985.

[21] Postel, J., Editor, "IAB Official Protocol Standards", RFC 1200,
     IAB, April 1991.

[22] Internet Activities Board, "Ethics and the Internet", RFC 1087,
     Internet Activities Board, January 1989.

[23] Pethia, R., Crocker, S., and B. Fraser, "Policy Guidelines for
     the Secure Operation of the Internet", CERT, TIS, CERT, RFC in
     preparation.

[24] Computer Emergency Response Team (CERT/CC), "Unauthorized
     Password Change Requests", CERT Advisory CA-91:03, April 1991.

[25] Computer Emergency Response Team (CERT/CC), "TELNET Breakin
     Warning", CERT Advisory CA-89:03, August 1989.

[26] CCITT, Recommendation X.509, "The Directory: Authentication
     Framework", Annex C.

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

8. Annotated Bibliography

    The intent of this annotated bibliography is to offer a
    representative collection of resources of information that will help
    the user of this handbook. It is meant provide a starting point for
    further research in the security area. Included are references to
    other sources of information for those who wish to pursue issues of
    the computer security environment.

8.1 Computer Law

[ABA89]
          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.

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

          Kept up to date with supplements.
          Years covering 1978-1984 focuses on: Computer law,
          evidence and procedures. The years 1984 to the current
          focus on general computer law. Bibliographical
          references and index included.

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

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

          Court cases and decisions rendered by federal and state
          courts throughout the United States on federal and state
          computer law. Includes Case Table and Topical Index.

[CONLY]
          Conly, C., "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.

[FENWICK]
          Fenwick, W., 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.

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

    [HUBAND]
            Huband, F., 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.

    [MCEWEN]
            McEwen, J., "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.

    [PARKER]
            Parker, D., "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.

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

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

8.2 Computer Security

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

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

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

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

R. Brand, 8 June 1990.

As computer security becomes a more important issue in
modern society, it begins to warrant a systematic approach.
The vast majority of the computer security problems and the
costs associated with them can be prevented with simple
inexpensive measures. The most important and cost
effective of these measures are available in the prevention
and planning phases. These methods are presented in this
paper, followed by a simplified guide to incident
handling and recovery. Available on-line from:
cert.sei.cmu.edu:/pub/info/primer.

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

Brief abstract (slight paraphrase from the original
abstract): AT&T maintains a large internal Internet that
needs to be protected from outside attacks, while
providing useful services between the two.
This paper describes AT&T's Internet gateway. This
gateway passes mail and many of the common Internet
services between AT&T internal machines and the Internet.
This is accomplished without IP connectivity using a pair
of machines: a trusted internal machine and an untrusted
external gateway. These are connected by a private link.
The internal machine provides a few carefully-guarded
services to the external gateway. This configuration
helps protect the internal internet even if the external
machine is fully compromised.

This is a very useful and interesting design. Most
firewall gateway systems rely on a system that, if
compromised, could allow access to the machines behind
the firewall. Also, most firewall systems require users
who want access to Internet services to have accounts on
the firewall machine. AT&T's design allows AT&T internal
internet users access to the standard services of TELNET and
FTP from their own workstations without accounts on
the firewall machine. A very useful paper that shows
how to maintain some of the benefits of Internet
connectivity while still maintaining strong
security.

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

          This paper describes measures that you, as a system
          administrator can take to make your UNIX system(s) more
          secure. Oriented primarily at SunOS 4.x, most of the
          information covered applies equally well to any Berkeley
          UNIX system with or without NFS and/or Yellow Pages (NIS).
          Some of the information can also be applied to System V,
          although this is not a primary focus of the paper. A very
          useful reference, this is also available on the Internet in
          various locations, including the directory
          cert.sei.cmu.edu:/pub/info.

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

          This book serves as a good guide to the issues encountered
          in forming computer security policies and procedures. The
          book is designed as a textbook for an introductory course
          in information systems security.

          The book is divided into five sections: Risk Management (I),
          Safeguards: security and control measures, organizational
          and administrative (II), Safeguards: Security and Control
          Measures, Technical (III), Legal Environment and
          Professionalism (IV), and CICA Computer Control Guidelines
          (V).

          The book is particularly notable for its straight-forward
          approach to security, emphasizing that common sense is the
          first consideration in designing a security program. The
          authors note that there is a tendency to look to more
          technical solutions to security problems while overlooking
          organizational controls which are often cheaper and much
          more effective. 298 pages, including references and index.

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

          Approx 450 pages, $29.95. Orders: 1-800-338-6887
          (US & Canada), 1-707-829-0515 (Europe), email: nuts@ora.com

          This is one of the most useful books available on Unix

security. The first part of the book covers standard Unix
and Unix security basics, with particular emphasis on
passwords. The second section covers enforcing security on
the system. Of particular interest to the Internet user are
the sections on network security, which address many
of the common security problems that afflict Internet Unix
users. Four chapters deal with handling security incidents,
and the book concludes with discussions of encryption,
physical security, and useful checklists and lists of
resources. The book lives up to its name; it is filled with
specific references to possible security holes, files to
check, and things to do to improve security. This
book is an excellent complement to this handbook.

[GREENIA90]
        Greenia, M., "Computer Security Information Sourcebook",
        Lexikon Services, Sacramento, CA, 1989.

        A manager's guide to computer security. Contains a
        sourcebook of key reference materials including
        access control and computer crimes bibliographies.

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

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

        A white paper prepared for the EMA, written by two experts
        in privacy law. Gives background on the issues, and presents
        some policy options.

        Available from: The Electronic Mail Association (EMA)
        1555 Wilson Blvd, Suite 555, Arlington, VA, 22209.
        (703) 522-7111.

[KENT]
        Kent, Stephen, "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.

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

[LU1]
          Lu, W., 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.

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

          NSA's catalogue contains chapter on: Endorsed Cryptographic
          Products List; NSA Endorsed Data Encryption Standard (DES)
          Products List; Protected Services List; Evaluated Products
          List; Preferred Products List; and Endorsed Tools List.

          The catalogue is available from the Superintendent of
          Documents, U.S. Government Printing Office, Washington,
          D.C. One may place telephone orders by calling:
          (202) 783-3238.

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

          This report, prepared for congressional committee considering
          Federal policy on the protection of electronic information, is
          interesting because of the issues it raises regarding the
          impact of technology used to protect information. It also
          serves as a reasonable introduction to the various encryption
          and information protection mechanisms. 185 pages. Available
          from the U.S. Government Printing Office.

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

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

          A general textbook in computer security, this book provides an
          excellent and very readable introduction to classic computer

security problems and solutions, with a particular emphasis on
encryption. The encryption coverage serves as a good
introduction to the subject. Other topics covered include
building secure programs and systems, security of database,
personal computer security, network and communications
security, physical security, risk analysis and security
planning, and legal and ethical issues. 538 pages including
index and bibliography.

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

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

        This is a good general reference on computer viruses and
        related concerns. In addition to describing viruses in
        some detail, it also covers more general security issues,
        legal recourse in case of security problems, and includes
        lists of laws, journals focused on computers security,
        and other security-related resources.

        Available from: ADAPSO, 1300 N. 17th St, Suite 300,
        Arlington VA 22209. (703) 522-5055.

[STOLL88]
        Stoll, C., "Stalking the Wily Hacker", Communications
        of the ACM, Vol. 31, No. 5, Pgs. 484-497, ACM,
        New York, NY, May 1988.

        This article describes some of the technical means used
        to trace the intruder that was later chronicled in
        "Cuckoo's Egg" (see below).

[STOLL89]
        Stoll, C., "The Cuckoo's Egg", ISBN 00385-24946-2,
        Doubleday, 1989.

        Clifford Stoll, an astronomer turned UNIX System
        Administrator, recounts an exciting, true story of how he
        tracked a computer intruder through the maze of American
        military and research networks. This book is easy to
        understand and can serve as an interesting introduction to
        the world of networking. Jon Postel says in a book review,

"[this book] ... is absolutely essential reading for anyone
that uses or operates any computer connected to the Internet
or any other computer network."

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

8.3 Ethics

[CPSR89]
Computer Professionals for Social Responsibility, "CPSR
Statement on the Computer Virus", CPSR, Communications of the
ACM, Vol. 32, No. 6, Pg. 699, June 1989.

This memo is a statement on the Internet Computer Virus
by the Computer Professionals for Social Responsibility
(CPSR).

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

A collection of 40 pieces divided into six sections: the
emergence of worldwide computer networks, electronic breakins,
worms, viruses, counterculture (articles examining the world
of the "hacker"), and finally a section discussing social,
legal, and ethical considerations.

A thoughtful collection that addresses the phenomenon of
attacks on computers. This includes a number of previously
published articles and some new ones. The previously
published ones are well chosen, and include some references
that might be otherwise hard to obtain. This book is a key
reference to computer security threats that have generated
much of the concern over computer security in recent years.

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

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

From the preface: "The aim of this book is two-fold: (1) to
describe some of the problems created by society by computers,
and (2) to show how these problems present ethical dilemmas
for computers professionals and computer users.

The problems created by computers arise, in turn, from two
main sources: from hardware and software malfunctions and
from misuse by human beings. We argue that computer systems
by their very nature are insecure, unreliable, and
unpredictable -- and that society has yet to come to terms
with the consequences. We also seek to show how society
has become newly vulnerable to human misuse of computers in
the form of computer crime, software theft, hacking, the
creation of viruses, invasions of privacy, and so on."

The eight chapters include "Computer Crime", "Software
Theft", "Hacking and Viruses", "Unreliable Computers",
"The Invasion of Privacy", "AI and Expert Systems",
and "Computerizing the Workplace." Includes extensive
notes on sources and an index.

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

[IAB89]
        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.

        This memo is a statement of policy by the Internet
        Activities Board (IAB) concerning the proper use of
        the resources of the Internet. Available on-line on
        host ftp.nisc.sri.com, directory rfc, filename rfc1087.txt.
        Also available on host nis.nsf.net, directory RFC,
        filename RFC1087.TXT-1.

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

[MIT89]
        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.

This memo is a statement of policy by the Massachusetts
Institute of Technology (MIT) on the responsible use
of computers.

[NIST]

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

[NSF88]

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.

This memo is a statement of policy by the National Science
Foundation (NSF) concerning the ethical use of the Internet.

[PARKER90]
Parker, D., Swope, S., and B. Baker, "Ethical Conflicts:
Information and Computer Science, Technology and Business",
QED Information Sciences, Inc., Wellesley, MA. (245 pages).

Additional publications on Ethics:

The University of New Mexico (UNM)

The UNM has a collection of ethics documents. Included are
legislation from several states and policies from many
institutions.

Access is via FTP, IP address ariel.umn.edu. Look in the
directory /ethics.

8.4 The Internet Worm

[BROCK]

Brock, J., "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.

Testimonial statement of Jack L. Brock, Director, U. S.
Government Information before the Subcommittee on
Telecommunications and Finance, Committee on Energy and

Commerce, House of Representatives.

[EICHIN89]
          Eichin, M., and J. Rochlis, "With Microscope and Tweezers:
          An Analysis of the Internet Virus of November 1988",
          Massachusetts Institute of Technology, February 1989.

          Provides a detailed dissection of the worm program. The
          paper discusses the major points of the worm program then
          reviews strategies, chronology, lessons and open issues,
          Acknowledgments; also included are a detailed appendix
          on the worm program subroutine by subroutine, an
          appendix on the cast of characters, and a reference section.

[EISENBERG89]
          Eisenberg, T., D. Gries, J. Hartmanis, D. Holcomb,
          M. Lynn, and T. Santoro, "The Computer Worm", Cornell
          University, 6 February 1989.

          A Cornell University Report presented to the Provost of the
          University on 6 February 1989 on the Internet Worm.

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

          This 36 page report (GAO/IMTEC-89-57), by the U.S.
          Government Accounting Office, describes the Internet worm
          and its effects. It gives a good overview of the various
          U.S. agencies involved in the Internet today and their
          concerns vis-a-vis computer security and networking.

          Available on-line on host nnsc.nsf.net, directory
          pub, filename GAO_RPT; and on nis.nsf.net, directory nsfnet,
          filename GAO_RPT.TXT.

[REYNOLDS89]
          The Helminthiasis of the Internet, RFC 1135,
          USC/Information Sciences Institute, Marina del Rey,
          CA, December 1989.

          This report looks back at the helminthiasis (infestation
          with, or disease caused by parasitic worms) of the
          Internet that was unleashed the evening of 2 November 1988.
          This document provides a glimpse at the infection,its
          festering, and cure. The impact of the worm on the Internet
          community, ethics statements, the role of the news media,

crime in the computer world, and future prevention is
discussed. A documentation review presents four publications
that describe in detail this particular parasitic computer
program. Reference and bibliography sections are also
included. Available on-line on host ftp.nisc.sri.com
directory rfc, filename rfc1135.txt. Also available on
host nis.nsf.net, directory RFC, filename RFC1135.TXT-1.

[SEELEY89]
        Seeley, D., "A Tour of the Worm", Proceedings of 1989
        Winter USENIX Conference, Usenix Association, San Diego, CA,
        February 1989.

        Details are presented as a "walk thru" of this particular
        worm program. The paper opened with an abstract,
        introduction, detailed chronology of events upon the
        discovery of the worm, an overview, the internals of the
        worm, personal opinions, and conclusion.

[SPAFFORD88]
        Spafford, E., "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.

        Describes the infection of the Internet as a worm
        program that exploited flaws in utility programs in
        UNIX based systems. The report gives a detailed
        description of the components of the worm program:
        data and functions. Spafford focuses his study on two
        completely independent reverse-compilations of the
        worm and a version disassembled to VAX assembly language.

[SPAFFORD89]
        Spafford, G., "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.

8.5 National Computer Security Center (NCSC)

   All NCSC publications, approved for public release, are available
   from the NCSC Superintendent of Documents.

           NCSC = National Computer Security Center

              9800 Savage Road
              Ft Meade, MD 20755-6000

              CSC = Computer Security Center:
              an older name for the NCSC

              NTISS = National Telecommunications and
              Information Systems Security
              NTISS Committee, National Security Agency
              Ft Meade, MD 20755-6000

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

              The security provided by a password system depends on
              the passwords being kept secret at all times. Thus, a
              password is vulnerable to compromise whenever it is used,
              stored, or even known. In a password-based authentication
              mechanism implemented on an ADP system, passwords are
              vulnerable to compromise due to five essential aspects
              of the password system: 1) a password must be initially
              assigned to a user when enrolled on the ADP system;
              2) a user's password must be changed periodically;
              3) the ADP system must maintain a 'password
              database'; 4) users must remember their passwords; and
              5) users must enter their passwords into the ADP system at
              authentication time. This guideline prescribes steps to be
              taken to minimize the vulnerability of passwords in each of
              these circumstances.

      [NCSC1]
              NCSC, "A Guide to Understanding AUDIT in Trusted Systems",
              NCSC-TG-001, Version-2, 1 June 1988, 25 pages.

              Audit trails are used to detect and deter penetration of
              a computer system and to reveal usage that identifies
              misuse. At the discretion of the auditor, audit trails
              may be limited to specific events or may encompass all of
              the activities on a system. Although not required by
              the criteria, it should be possible for the target of the
              audit mechanism to be either a subject or an object. That
              is to say, the audit mechanism should be capable of
              monitoring every time John accessed the system as well as
              every time the nuclear reactor file was accessed; and
              likewise every time John accessed the nuclear reactor
              file.

[NCSC2]

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

Discretionary control is the most common type of access
control mechanism implemented in computer systems today.
The basis of this kind of security is that an individual
user, or program operating on the user's behalf, is
allowed to specify explicitly the types of access other
users (or programs executing on their behalf) may have to
information under the user's control. [...] Discretionary
controls are not a replacement for mandatory controls. In
any environment in which information is protected,
discretionary security provides for a finer granularity of
control within the overall constraints of the mandatory
policy.

[NCSC3]

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

Configuration management consists of four separate tasks:
identification, control, status accounting, and auditing.
For every change that is made to an automated data
processing (ADP) system, the design and requirements of the
changed version of the system should be identified. The
control task of configuration management is performed
by subjecting every change to documentation, hardware, and
software/firmware to review and approval by an authorized
authority. Configuration status accounting is responsible
for recording and reporting on the configuration of the
product throughout the change. Finally, though the process
of a configuration audit, the completed change can be
verified to be functionally correct, and for trusted
systems, consistent with the security policy of the system.

[NTISS]

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

This document provides guidance to users, managers, security
officers, and procurement officers of Office Automation
Systems. Areas addressed include: physical security,
personnel security, procedural security, hardware/software
security, emanations security (TEMPEST), and communications

security for stand-alone OA Systems, OA Systems
used as terminals connected to mainframe computer systems,
and OA Systems used as hosts in a Local Area Network (LAN).
Differentiation is made between those Office Automation
Systems equipped with removable storage media only (e.g.,
floppy disks, cassette tapes, removable hard disks) and
those Office Automation Systems equipped with fixed media
(e.g., Winchester disks).

Additional NCSC Publications:

   [NCSC4]
           National Computer Security Center, "Glossary of Computer
           Security Terms", NCSC-TG-004, NCSC, 21 October 1988.

   [NCSC5]
           National Computer Security Center, "Trusted
           Computer System Evaluation Criteria", DoD 5200.28-STD,
           CSC-STD-001-83, NCSC, December 1985.

   [NCSC7]
           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.

   [NCSC8]
           National Computer Security Center, "Technical Rationale
           Behind CSC-STD-003-85: Computer Security Requirements",
           CSC-STD-004-85, NCSC, 25 June 85.

   [NCSC9]
           National Computer Security Center, "Magnetic Remanence
           Security Guideline", CSC-STD-005-85, NCSC, 15 November 1985.

           This guideline is tagged as a "For Official Use Only"
           exemption under Section 6, Public Law 86-36 (50 U.S. Code
           402). Distribution authorized of U.S. Government agencies
           and their contractors to protect unclassified technical,
           operational, or administrative data relating to operations
           of the National Security Agency.

   [NCSC10]
           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.

[NCSC11]
        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.

[NCSC12]
        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.

[NCSC13]
        National Computer Security Center, "Trusted Network
        Interpretation", NCSC-TG-005, NCSC, 31 July 1987.

[NCSC14]
        Tinto, M., "Computer Viruses: Prevention, Detection, and
        Treatment", National Computer Security Center C1
        Technical Report C1-001-89, June 1989.

[NCSC15]
        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.

8.6 Security Checklists

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

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

8.7 Additional Publications

   Defense Data Network's Network Information Center (DDN NIC)

      The DDN NIC maintains DDN Security bulletins and DDN Management

bulletins online on the machine: NIC.DDN.MIL. They are available
via anonymous FTP. The DDN Security bulletins are in the
directory: SCC, and the DDN Management bulletins are in the
directory: DDN-NEWS.

For additional information, you may send a message to:
NIC@NIC.DDN.MIL, or call the DDN NIC at: 1-800-235-3155.

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

        A Defense Data Network Management Bulletin announcement
        on the 4.2bsd and 4.3bsd software fixes to the Internet
        worm.

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

    IEEE Proceedings

[IEEE]
        "Proceedings of the IEEE Symposium on Security
        and Privacy", published annually.

    IEEE Proceedings are available from:

            Computer Society of the IEEE
            P.O. Box 80452
            Worldway Postal Center
            Los Angeles, CA 90080

    Other Publications:

    Computer Law and Tax Report
    Computers and Security
    Security Management Magazine
    Journal of Information Systems Management
    Data Processing & Communications Security
    SIG Security, Audit & Control Review

9. Acknowledgments

    Thanks to the SSPHWG's illustrious "Outline Squad", who assembled at
    USC/Information Sciences Institute on 12-June-90: Ray Bates (ISI),
    Frank Byrum (DEC), Michael A. Contino (PSU), Dave Dalva (Trusted
    Information Systems, Inc.), Jim Duncan (Penn State Math Department),
    Bruce Hamilton (Xerox), Sean Kirkpatrick (Unisys), Tom Longstaff
    (CIAC/LLNL), Fred Ostapik (SRI/NIC), Keith Pilotti (SAIC), and Bjorn
    Satdeva (/sys/admin, inc.).

    Many thanks to Rich Pethia and the Computer Emergency Response Team
    (CERT); much of the work by Paul Holbrook was done while he was
    working for CERT. Rich also provided a very thorough review of this
    document. Thanks also to Jon Postel and USC/Information Sciences
    Institute for contributing facilities and moral support to this
    effort.

    Last, but NOT least, we would like to thank members of the SSPHWG and
    Friends for their additional contributions: Vint Cerf (CNRI),
    Dave Grisham (UNM), Nancy Lee Kirkpatrick (Typist Extraordinaire),
    Chris McDonald (WSMR), H. Craig McKee (Mitre), Gene Spafford (Purdue),
    and Aileen Yuan (Mitre).

10. Security Considerations

    If security considerations had not been so widely ignored in the
    Internet, this memo would not have been possible.

11. Authors' Addresses

    J. Paul Holbrook
    CICNet, Inc.
    2901 Hubbard
    Ann Arbor, MI 48105

    Phone: (313) 998-7680
    EMail: holbrook@cic.net

    Joyce K. Reynolds
    University of Southern California
    Information Sciences Institute
    4676 Admiralty Way
    Marina del Rey, CA 90292

    Phone: (213) 822-1511
    EMail: JKREY@ISI.EDU