NIST Special Publication 800-30

Risk Management Guide

NIST
National Institute of Standards and Technology
Technology Administration
U.S. Department of Commerce

COMPUTER SECURITY

1st Public Exposure DRAFT
June 2001

1901-2001
NIST CENTENNIAL

NIST Special Publication 800-30

Risk Management Guide

Recommendations of the
National Institute of Standards and Technology

COMPUTER SECURITY

1st Public Exposure DRAFT – June 2001

U.S. DEPARTMENT OF COMMERCE

Donald L. Evans, Secretary

TECHNOLOGY ADMINISTRATION

Karen H. Brown, Acting Under Secretary of Commerce for Technology

NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY

Karen H. Brown, Acting Director

Reports on Computer Systems Technology

The Information Technology Laboratory (ITL) at the National Institute of Standards and Technology (NIST) promotes the U.S. economy and public welfare by providing technical leadership for the Nation's measurement and standards infrastructure. ITL develops tests, test methods, reference data, proof of concept implementations, and technical analyses to advance the development and productive use of information technology. ITL's responsibilities include the development of technical, physical, administrative, and management standards and guidelines for the cost-effective security and privacy of sensitive unclassified information in Federal computer systems. This Special Publication 800-series reports on ITL's research, guidance, and outreach efforts in computer security, and its collaborative activities with industry, government, and academic organizations.

National Institute of Standards and Technology Special Publication 800-27
Natl. Inst. Stand. Technol. Spec. Publ. 800-27, 30 pages (June 2001)
CODEN: NSPUE2

U.S. GOVERNMENT PRINTING OFFICE
WASHINGTON: 2001

For sale by the Superintendent of Documents, U.S. Government Printing Office
Internet: bookstore.gpo.gov — Phone: (202) 512-1800 — Fax: (202) 512-2250
Mail: Stop SSOP, Washington, DC 20402-0001

TABLE OF CONTENTS

TABLE OF FIGURES

TABLE OF TABLES

1.0 INTRODUCTION

As a result of the move to a digital economy, information and information technology (IT) have become valuable mission assets that need to be protected. With this development has come the recognition that fulfilling these basic functions requires comprehensive, well-designed, and reliable IT security programs.

One important component in an IT security program is an effective risk management process. The most fundamental principle on which to base the organization's risk management process is that the goal is to protect the *organization*, not simply its IT assets. Therefore, risk management should not be treated primarily as a technical function of the IT experts, but as an essential management function of the mission owner.

1.1 PURPOSE

This guide provides both definitional and practical guidance regarding the concept and practice of managing IT-related risks. Risk - the net impact of an adverse IT-related event - is a function of the likelihood of a given threat-source exercising a particular vulnerability, and the resulting impact.

In following the guidance in this document, IT personnel will be able to isolate a wide variety of risks, many of which are subjective, determine the extent of a compromise, and identify potential mitigation options. The approach that follows will help in identifying risks based on potential threats and the consequences of those threats as well as the associated risk mitigation techniques. In addition the document will provide information on the selection of security controls based on cost and the degree of risk reduction. This knowledge is a means for management to make well-informed risk management decisions and to justify the expenditures that are part of an information security budget.

The intent of this document is to provide a common, thorough foundation for use in the development of detailed risk management guidance and procedures.

1.2 DOCUMENT RELATIONSHIPS

This guide uses the general concepts presented in the National Institute of Standards and Technology (NIST)'s Special Publication 800-27, *Engineering Principles for IT Security* along with the principles and practices in NIST Special Publication 800-14, *Generally Accepted Principles and Practices for Securing Information Technology Systems*. In addition, it is consistent with the policies in the Office of Management and Budget (OMB) Circular A-130, Appendix III, *Security of Federal Automated Information Resources.*

1.3 DOCUMENT STRUCTURE

This guide is organized by the three phases of an ongoing risk management process: performing a risk assessment; addressing the mitigation of that risk; and evaluating the results. The guide also contains two appendixes. Appendix A is a glossary of terms used frequently in this document, and Appendix B provides a sample outline to use in documenting results.

2.0 RISK MANAGEMENT OVERVIEW

2.1 IMPORTANCE OF RISK MANAGEMENT

Risk management is the process that allows managers to balance operational and economic costs of protective measures with the resulting gain in mission effectiveness. This process is not unique to the IT environment; indeed it pervades our decision-making on a daily basis. Take the case of home security. Many people decide to have security systems installed and pay a monthly fee to a service provider to have these systems monitored. Presumably, the homeowners calculated the cost of installation and monitoring against the value of their household goals and their family's safety, a fundamental “mission” need.

Mission owners and managers must also weigh the cost of protective measures for their IT systems that store, process, and transmit their mission information against the risks to their mission. Few organizations have unlimited resources to spend on IT security and therefore IT security spending must be considered just as thoroughly as other mission decisions.

2.2 INTEGRATION OF RISK MANAGEMENT IN LIFECYCLE PLANNING

Minimizing risk to the mission is the fundamental reason why organizations implement information security measures. In fact, all security-related activities are a part of the risk management process. As a result, risk management spans the entire system development life cycle (SDLC). Throughout this document reference will be made to the SDLC phases in which the particular activity is relevant. As described in several NIST Special Publications, the SDLC has five phases. They are–

- **Initiation Phase**: During the initiation phase, the need for a system is expressed and the purpose of the system is documented.
- **Development/Acquisition Phase**: During this phase, the system is designed, purchased, programmed, developed, or otherwise constructed. This phase often consists of other defined cycles, such as the system development cycle or the acquisition cycle.
- **Implementation Phase**: During implementation, the system is tested and installed or fielded.
- **Operation/Maintenance Phase**: During this phase, the system performs its work. Typically the system is being modified on an ongoing basis by the addition of hardware and software and by changes to mission, policy and procedures.
- **Disposal Phase**: The disposal phase of the IT system life cycle involves the sanitizing of information, hardware, and software.

2.3 KEY ROLES

This section describes the key roles in the risk management process

Agency Senior Management

Management's role is primarily associated with decisions about spending levels and acceptable amounts residual risk. Senior decision makers therefore should play a significant role in the risk management process. They must see the security costs being worth the benefit.

Mission Process Owner

The mission owner is responsible for accomplishing the mission. As such, they must determine what constitutes acceptable risk and what tradeoffs are appropriate to maximize mission effectiveness within the given constraints. Thus, it is crucial that they understand the risk management process.

IT System Owner

The system owners are typically responsible for changes made to the system. As such, they usually have to sign off on work before it can be done. The system owner must therefore understand the role risk management plays in overall system effectiveness.

Information System Security Officers

Information system security officers are responsible for IT security (hardware, software, and data) of the system. In this role, these individuals will select the appropriate controls and meet with vendors to discuss the capabilities of the controls. They will also arrange training for personnel in order to familiarize them with security requirements and rules of behavior necessary to protect the system and its data.

System Administrators

As new components are added to a system, or the existing infrastructure is reconfigured for optimal performance, the system administrator will need to adjust lifecycle activities accordingly. This justification is necessary because changes to the system will often impact system and data security. Therefore, system administrators need to be aware of how changes will affect the security posture of the organization.

End Users

The organization's staff is the users of the system. Their use of the system according to appropriate guidelines and rules of behavior are critical in protecting an organization's mission. It is essential that they understand the potential risks and their role in the risk management process.

3.0 RISK ASSESSMENT

The cornerstone of risk management is the risk assessment, a process for organizations to use in determining the level of risk associated with a given system throughout its SDLC. The output of this process is the residual risk and a determination whether this is at an acceptable level or whether additional security controls should be implemented to further reduce risk.

Risk is a function of the likelihood of a security event and the impact that event would have on the organization's mission. To determine likelihood, threats to the system are analyzed in conjunction with the vulnerabilities present. Impact is determined by considering the criticality of the system in supporting the organizational mission. This methodology is represented in Figure 1, below; the individual components are described in sections 3.1 through 3.4.

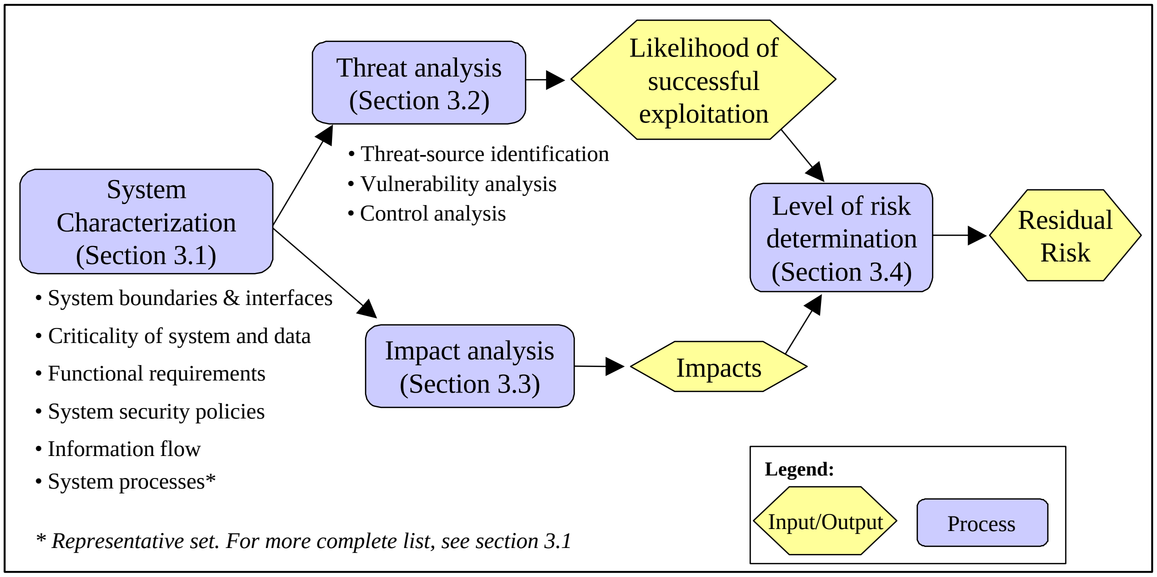
Threat analysis (Section 3.2)
Likelihood of successful exploitation
• Threat-source identification
• Vulnerability analysis
• Control analysis
System Characterization (Section 3.1)
Level of risk determination (Section 3.4)
Residual Risk
• System boundaries & interfaces
• Criticality of system and data
• Functional requirements
• System security policies
• Information flow
• System processes*
Impact analysis (Section 3.3)
Impacts
Legend:
Input/Output
Process
* Representative set. For more complete list, see section 3.1

Figure 1 - Risk Assessment Methodology

3.1. SYSTEM CHARACTERIZATION

Characterizing the system establishes the scope of the risk management effort and provides information essential to defining the risk. This step is necessary so that everyone involved can understand the organization's mission and system operations and the nature of the potential mission impact arising from the IT. In this step, boundaries of the system are identified, along with the resources and information that constitute it. These assets are usually classified as follows:

- Information infrastructure
- Hardware
- Data and information
- People
- System interfaces and connectivity.

Additional information to be collected about the system and its data includes:

- The organization's mission
- The processes performed by the system
- The functional requirements of the system
- Users of the system
- All applicable system security policies governing the system (agency policies, federal requirements, law)
- System security architecture
- The operating environment of the system
- The facilities where the system is contained
- The information storage requirements of the system
- The flow of information pertaining to the system.

For an operational system, the data is collected about the system as it exists, regardless of whether the information is formally documented. For a system under development, analysis needs to be performed to define key security rules and attributes of the future system.

The system description should also include any assumptions made as well as all sources of information used to develop the description. Assumptions may be necessary if the documentation is silent on a given topic or if the discussion is incomplete. They might include assumptions about security provided by the underlying infrastructure or about future plans for the system.

3.2 THREAT ANALYSIS

Threat is expressed as a function of the likelihood that a given threat-source will successfully exploit a given vulnerability. A vulnerability is a weakness that can be accidentally triggered or intentionally exploited. Without a vulnerability that can be exercised, a threat-source does not present a risk. In determining likelihood, one must consider threat-sources, vulnerabilities, and existing controls, as described in sections 3.2.1 through 3.2.3.

Threat: The potential for a "threat-source" to exercise (accidentally trigger or intentionally exploit) a specific vulnerability.

3.2.1 Threat-Source Identification

The goal in this step is to identify and develop a list of potential threat-sources: natural, human, and environmental. A threat-source is defined as any circumstance or event with the potential to cause harm to an information system.

Threat-source: Either (1) intent and method targeted at the intentional exploitation of a vulnerability or (2) the situation and method that may accidentally trigger a vulnerability.

In assessing threat-sources, it is important to ensure appropriate natural and environmental threats to the system are considered. Many times these can be overlooked but can cause as much, or more, damage as manmade threats. Of course, these threats are highly dependent on the location of the system. Systems located in a desert may not have “flood” listed in their list of threat-sources.

Manmade threat-sources can either be intentional - a deliberate attack - or unintentional. A deliberate attack can be either (1) a malicious attempt to gain unauthorized access to an information system to compromise its integrity, availability, or confidentiality or (2) a benign but nonetheless purposeful attempt to circumvent security (for example bypassing controls to “get the job done). One example of an attack is an attempt to get a valid user’s password to gain access to personal information.

Common Threat-Sources

- Natural Threats—floods, earthquakes, tornadoes, landslides, avalanches, electrical storms, and other events.
- Human Threats—events that are either enabled by or caused by human beings such as unintentional acts (inadvertent data entry) or deliberate actions (network based attacks, malicious software upload, unauthorized access to confidential information).
- Environmental Threats—long-term power failure, pollution, chemicals, liquid leakage.

In order for a human to be a valid threat-source, motivation and the resources to carry out the attack must be present. Table 1 below presents an overview of the types of attackers, what their motivations might be, and the means by which they might carry out the attack. Whether these agents may be interested in the system will depend on many factors. Using the information from the system characterization, identify which might apply in each case. Once a list of potential threat agents has been identified, one should develop a reasonable estimate of the resources and capabilities that may be required to carry out an attack. These range from an external connection into the system using automated tools to perform the attack all the way to insider knowledge of weaknesses in the system not generally known.

Table 1 - Human Threats

Threat-source	Motivation	Means
Hacker, cracker	Ego Challenge Rebellion	System Intrusion Unauthorized system access
Criminal	Illegal Disclosure Alteration Monetary Gain	Crime/Intrusion Fraudulent act
Terrorist	Blackmail Exploitation Destruction	System Attack/intrusion
Foreign Interests	Classified Information Other government interests	Intrusion/penetration
Insider (disgruntled, negligent, or dishonest employee)	Intelligence Revenge Ego Monetary Gain	Intrusion, computer abuse, unauthorized system access

Known threats can be obtained from many government and private sector organizations. Intrusion detection systems are becoming more prevalent and government and industry organizations continue collecting more data on security events, thereby improving ability to realistically assess threats. These sources include—

- Intelligence agencies (for example, the FBI's National Infrastructure Protection Center, NIPC)
- Federal Computer Incident Response Center (FedCIRC)
- Mass media, particularly web-based resources such as SecurityFocus.com, SecurityWatch.com, SecurityPortal.com, and SANS.org.

In general, information on natural threats, e.g., floods, earthquakes, etc., should be readily available. In the absence of hard data it may be necessary to estimate the threat, but this still has value. In either case, the potential threat-sources should be tailored to the individual organization. The output of this step is a threat statement that lists potential threat-sources that are applicable to the system being evaluated.

3.2.2 Vulnerability Analysis

The goal in this step is to develop a list of the system flaws or weaknesses that could be exercised by the potential threat-sources. This step systematically evaluates the technical and non-technical weaknesses associated with the system. This information is collected via site surveys, interviews with personnel responsible for the system, network-scanning tools, and available system and organizational documentation. Other industry sources (e.g., vendor Web pages

Vulnerability: A flaw or weakness in system security procedures, design, implementation, internal controls, etc., that could be exercised (accidentally triggered or intentionally exploited) and result in a violation of the system's security policy.

that identify system bugs) are helpful in identifying vulnerabilities that may be applicable to specific systems. The specific types of vulnerabilities, and the methodology needed to determine whether they are present, will usually vary depending on the nature of the system and whether the system is in the design phase or has already been implemented.

If the system has not yet been designed, then the search for vulnerabilities focuses on security policies, procedures, and system requirement definitions. If the system is being implemented, the vulnerability identification would expand to include more specific information such as design documentation. If the system is operational, then the vulnerability identification methodology would include determining and analyzing whether the security features implemented or security controls used to mitigate the risk are applicable and effective.

Some proactive methods used to collect vulnerability information are—

- Automated vulnerability scanning
- Network mapping
- Security testing and evaluation
- Penetration testing.[1]

The Internet is a source of information. Known vulnerabilities are commonly posted by vendors, along with hot fixes, service packs, patches, or other remedial measures that may be applied to eliminate or mitigate vulnerabilities. Therefore, to perform a thorough vulnerability analysis, documented vulnerability sources that should be considered include:

- Previous risk assessment
- Audit reports, security review reports, and system test and evaluation reports
- Vulnerabilities lists such as the NIST I-CAT vulnerability database (http://icat.nist.gov/icat.taf)
- Security advisories such as Federal Computer Incident Response Capability (FedCIRC) and Department of Energy’s Computer Incident Advisory Capability (CIAC) bulletins
- Vendor advisories
- System software security analyses
- System anomaly reports.

The organization should research and analyze the available resources to support system vulnerability analysis and associate the identified vulnerabilities with specific system or information elements within the construct of the threat environment. Vulnerability analysis attempts to uncover all flaws and weaknesses, indicating those that may be exercised and those that probably will not be exercised. A flaw is unlikely to be exercised due to a low-level of threat-source interest or capability, effective security controls, or both.

[1] *Penetration testing is a portion of security testing in which assessors attempt to circumvent security features of the system to test the system from the point of view of a threat agent and identify potential failures in information system protection schemes. Most organizations have strict rules governing penetration testing activities.*

3.2.3 Control Analysis

During this step, the organization determines whether the security requirements collected during system characterization are being met by existing or planned security controls. Typically, the system security requirements are presented in matrix form where an explanation can be included that describes how the system's design or implementation does or does not satisfy the specific security control requirement. Security controls for the system can be extrapolated from the following sources:

- Security policies and guidelines
- System operating procedures
- System security specifications
- Industry standards and good practices

Many organizations find it helpful to group the controls into three categories as shown in the examples below. Each of these control categories is described in more detail in Section 4.3, Implementing Controls. In each case, some controls serve to *prevent* a security event, others to *detect* a security event.

Technical controls are those safeguards incorporated in computer hardware, software or firmware. Table 2 lists some of the technical controls used to mitigate risk.

Table 2 - Example Technical Controls

Prevent	Detect
▪ Access control mechanisms ▪ Antivirus software ▪ Identification & Authentication mechanisms ▪ Firewalls ▪ Encryption	▪ Audit trails ▪ Intrusion detection systems

Operational Controls are those operational procedures and personnel and physical security measures established to provide an acceptable level of protection for computing resources. Table 3 lists some of the operational controls used to mitigate risk.

Table 3 - Examples of Operational Controls

Prevent	Detect
▪ Security awareness and training ▪ Disaster recovery, contingency, and emergency plans ▪ Background investigations	▪ Security reviews and audits

Management Controls are those security measures that focus on the management of the system and the management of risk. By their nature they all fall into the “prevent” category. They include security reviews and assessments, risk assessments, and rules of behavior.

3.2.4 Likelihood Determination

The final step in the threat assessment is to derive an overall likelihood rating. Factors that govern the threat likelihood include threat-source motivation and capability, the nature of the vulnerability, and the effectiveness of current countermeasures. A simple way to describe the likelihood that a vulnerability will be exercised by a given threat-source is high, moderate, or low. The table below describes these three likelihood levels.

Table 4 - Likelihood Definitions

Likelihood	Description
High	The threat-source is highly motivated and sufficiently capable and countermeasures to prevent the vulnerability from being exercised are ineffective.
Moderate	The threat-source is motivated and sufficiently capable but countermeasures are in place that will impede successful exercise of the vulnerability. or The threat-source lacks specific motivation to exercise this vulnerability or is only marginally capable of doing so.
Low	The threat-source lacks motivation or capability or controls are in place to prevent, or at least significantly impede, the vulnerability from being exercised.

3.3 IMPACT ANALYSIS

The next major step in the risk assessment process is to determine the mission impact resulting from the threats (exercise of a vulnerability by a threat-source). The impact of a security event can be described in terms of mission impacts attributed to loss or degradation of the five security goals – integrity, availability, confidentiality, accountability, and assurance. Below is a brief description of each security goal and the related consequence if they are not met:

- **Loss of Integrity.** Integrity is lost if unauthorized changes are made to the data or system, whether these changes are intentional or accidental. Loss of system or data integrity may cause impacts similar to those due to the loss of availability. Additionally, if the loss of system or data integrity is not discovered, continued use of the corrupted system or data could cause future problems. Also, violation of integrity (1) may be a first step toward achieving a successful attack against availability or confidentiality and (2) reduces the assurance of the system.

- **Loss of Availability.** If a system becomes partially or completely unavailable to its authorized users, mission accomplishment may suffer. Loss of functionality and operational effectiveness may, for example, result in loss of public confidence or lost productive time. Additionally, unauthorized use of system resources may result in additional loss of confidence and it various forms of liability.

- **Loss of Confidentiality.** Confidentiality refers to the protection of data (both user and system) from unauthorized disclosure. The impact of unauthorized disclosure can range from jeopardizing national security to embarrassment.

- **Loss of Accountability.** Accountability refers to the ability to trace the actions of an individual user. The accountability supports non-repudiation, deterrence, fault isolation, intrusion detection and prevention, and after-action recovery and legal action. Loss of accountability impacts the ability to perform these functions. Additionally, reducing the system's accountability capability is a frequent part of a achieving other ends such as violating integrity, confidentiality, or availability.

- **Loss of Assurance.** Assurance is the grounds for confidence that the other four security goals (integrity, availability, confidentiality, and accountability) have been adequately met by a specific implementation. Loss of assurance implies that there is not sufficient protection against unintentional user and/or software errors or the existence of adequate resistance to intentional penetration or by-pass. A successful exercise of a vulnerability results in a reduction in the grounds for confidence in the system.

Some impacts can be measured quantitatively in lost revenue, costs of repairing the systems, or costs in terms of levels of effort required to correct problems caused by a successful exploitation. Other intangible impacts (e.g., loss in public confidence, credibility) cannot be measured in specific units, but can be qualified in terms of high, moderate, and low. Because of the generic nature of this discussion, this guide uses the qualitative categories – critical, high, moderate, and low - as defined in Table 5 below.

Table 5 - Magnitude of Impact Definitions

Impact	Description
Critical Impact	Threat results in unavailability, modification, disclosure, or destruction of valued data or other system assets or loss of system services that is unacceptable due to the resulting disastrous national impact or likely deaths.
High Impact	Threat results in unavailability, modification, disclosure, or destruction of valued data or other system assets or loss of system services that is unacceptable due to the resulting significant degradation of mission or possible injury to persons.
Moderate Impact	Threat results in discernible but recoverable unavailability, modification, disclosure, or destruction of data or other system assets or loss of system services, resulting in transitory, yet important mission impact but no injury to persons.
Low Impact	Threat results in unavailability, modification, disclosure, or destruction of data or degradation of system services that does not cause a significant mission impact nor injury to persons.

Quantitative verses Qualitative

Consideration should be given to the advantages and disadvantages of a quantitative versus qualitative assessment. The advantage of the qualitative impact analysis is that it provides a relative prioritization of the risks and identifies immediate areas for improvement against the vulnerabilities. The disadvantage of the qualitative impact analysis is that it does not provide specific quantifiable measurements of the magnitudes of impact, therefore making the cost-benefit analysis of any recommended controls difficult at best.

On the other hand, the advantage of a quantitative impact analysis is that it provides a measurement of the magnitude that can be used in the cost-benefit analysis of recommended controls. The disadvantage is that depending on the units in which the measurement is expressed, the meaning of a quantitative impact analysis may be unclear, requiring that the result be interpreted in a qualitative manner. Additionally, if the quantitative values are the result of subjective judgments (frequently the case), then the use of quantitative methods may just hide the fact that the results are actually qualitative. Factors that assist in quantifying the magnitude of impact may include, but are not limited to—

- An estimation of the frequency of the threat-source exercising the vulnerability over a specified time period (e.g., 1 year)
- An approximate cost for each occurrence of the threat-source exercising the vulnerability

A weighted factor based on a subjective analysis of the relative priority of a specific threat exploiting a specific vulnerability.

3.4 LEVEL OF RISK DETERMINATION

As mentioned at the outset of this section, the final determination of risk to the system and data is derived by combining the two ratings generated in the previous two sections—threat and impact. Table 6 below provides an example of how overall risk rating might be determined based on an input from each threat likelihood and impact categories.

Table 6 - Level of Risk Determination

	Likelihood of Threat Occurrence		
Impact	***High***	***Moderate***	***Low***
Critical Impact	Critical	High	Moderate
High impact	High	Moderate	Low
Moderate impact	Moderate	Moderate	Low
Low impact	Low	Low	Low

3.5 RISK ASSESSMENT IN THE SDLC

Risk assessment activities are relevant in the first four SDLC phases. As described in NIST Special Publication 800-18, in the initiation phase, a criticality assessment should be conducted. As described above, in this step owners define how the system relates to mission accomplishment. In addition, potential threats and vulnerabilities can be identified. These activities will be useful in the next two phases, development and acquisition and implementation.

In the development and acquisition phase, a risk assessment may be conducted to ensure that the overall system design and architecture, including controls, provides a security capability commensurate with the acceptable risk levels. As detailed in the process description above, the mission(s) that will be supported by the system, and how a security breach of the system might impact the mission(s) are important considerations. Addressing risk while the system is in the design phase allows performance and cost trades for the security features of the system to be made deliberately and with fewer constraints than is typical in subsequent phases.

In the implementation phase, a risk assessment is conducted when new controls or system components have been added. For example, an analysis might be conducted after the addition of a remote access terminal. Performing a risk assessment then would allow administrators to determine how the support of external connections might impact the mission.

Finally, risk analyses are considered essential during the operations and maintenance phase of the SDLC - in anticipation of the occurrence of an event, or even after the occurrence of an event, to analyze vulnerabilities and recommend controls. A risk assessment would be appropriate if, for example, a server inside the firewall boundary were to experience a penetration. The analysis would seek to retrace the events leading up to the penetration to determine the penetration technique and its effects on the server. Once the vulnerabilities and threats have been identified, controls can be recommended to sufficiently reduce the risk in the future.

4. RISK MITIGATION

During this step of the process, additional controls are identified to sufficiently mitigate the identified risks to the organization's operations based on the results of the risk assessment. The goal in selecting controls is to reduce the level of risk to the mission to an acceptable level, with minimum decrease in other system capabilities. The elimination of all risk is typically impractical or impossible. Consequently, the goal is to protect a system with cost-effective and feasible security controls that are applicable to the system environment and supportive of mission accomplishment.

For each control that is proposed, the cost versus the benefit should be taken into account. This section will identify the various points at which mitigation of risk can be performed. Afterwards, the necessity of conducting a cost-benefit analysis for each proposed control will be covered.

4.1 RISK MITIGATION APPROACHES

To mitigate risk, organizations usually consider implementing a blend of the following three approaches:

- **Prevent**: Eliminate the threat by removing the flaw or weakness or the ability to exercise it.
- **Limit**: Implement controls that constrain the impact of a threat without the need to take additional actions.
- **Detect and Respond**: Implementing measures to detect the exercise of a vulnerability and take action to mitigate adverse outcomes.

In implementing technical and administrative solutions for each approach, it is important to keep in mind the goals and mission of the organization when deciding upon solutions. Simply because a threat can be addressed, does not necessarily mean that it makes sense to do so. Threats that would result in little impact to the mission should be a low priority to mitigate. Threats that result in the potential for significant mission impact should be given priority for mitigation.

The devices and applications used to implement controls may be from many different sources. The "best of breed" approach brings together various components from different vendors, along with administrative measures, with the idea of each contributing in its own, specialized way. These components, both software and hardware, come together to form the security architecture.

4.2 RISK MITIGATION OVERVIEW

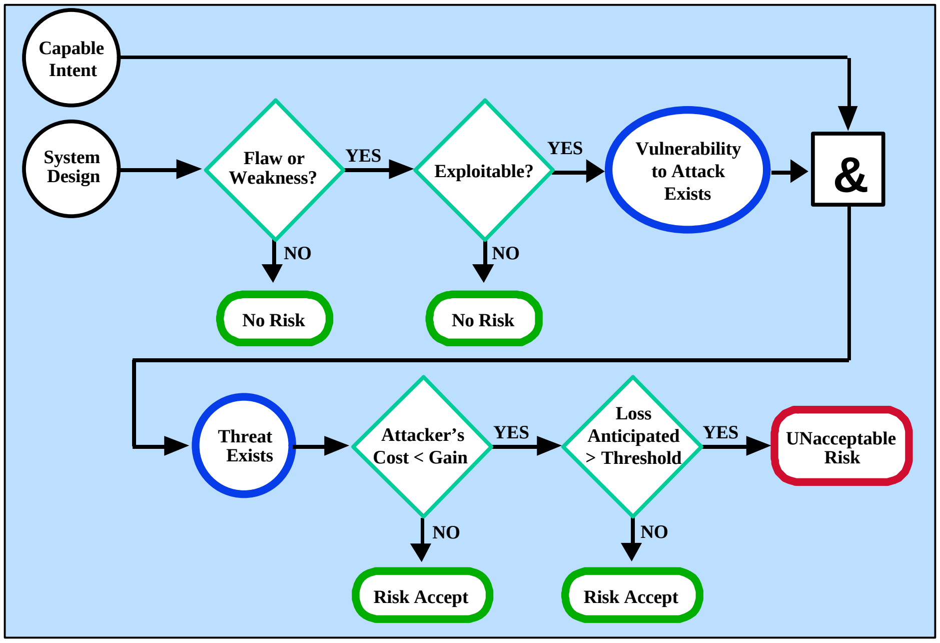

Figure 2 - Basics of Risk Mitigation

Figure 2 shows that the mitigation of risk can be accomplished at the following points:

- **Flaw exists—**implement assurance techniques to reduce the likelihood of a flaw
- **Flaw is exploitable—**apply layered protections, architectural designs, and administrative controls to prevent exploitability
- **Attacker's cost is less than gain—**apply protections to increase attacker's cost (note that administrative choices such as limiting what is processed can significantly reduce attacker's gain)
- **Loss too great—**apply design principles, architectural designs, and administrative protections to limit extent of attack, thereby reducing loss. (Again, note that administrative choices such as limiting what is processed may provide the most effective risk mitigation.)

Figure 2 also applies to the mitigation of risks arising from system and user errors. For these situations the mitigation of risk is very similar; the only point that does not apply is the third, since there is no "attacker."

4.3 IMPLEMENTING CONTROLS

In implementing controls, the organization should consider both technical, operational, and management security controls. In making a decision about which controls to choose, target the greatest risks and strive for sufficient risk mitigation at the lowest cost with minimal impact on other mission capabilities. Security controls seek to prevent, limit, or deter threat-sources from inflicting damage to the organization's mission, but are rarely a "bulletproof" solution. Recognizing which risks each security control is targeting and the degree of protection that can be reasonably expected are essential to cost-effective protection of the organizational mission.

The following section provides a high level overview of some of the types of controls to be considered as well as a discussion on how to conduct a cost-benefit analysis. More detailed guidance can be found in NIST Special Publication 800-18, *Guide for Developing Security Plans for Information Technology Systems* and NIST Special Publication 800-12, *An Introduction to Computer Security: The NIST Handbook.*

4.3.1 Technical Security Controls

Technical means of risk mitigation can be tailored to protect against given types of threats. These may range from simple to complex measures, and typically involve system architectures, engineering disciplines, and security packages involving a mix of hardware and software – all working together towards securing critical data and vital functions. Technical controls can be grouped into one of the following three major categories, according to primary purpose:

- Support. These controls are generic and underlie most information technology security capabilities.
- Prevent. These controls focus on preventing a security breach from occurring
- Detect and Recover. The controls in this category focus on the detection and recovery from a security breach.

Figure 3 provides a pictorial representation of the primary technical controls and an indication of the relationships between these controls.

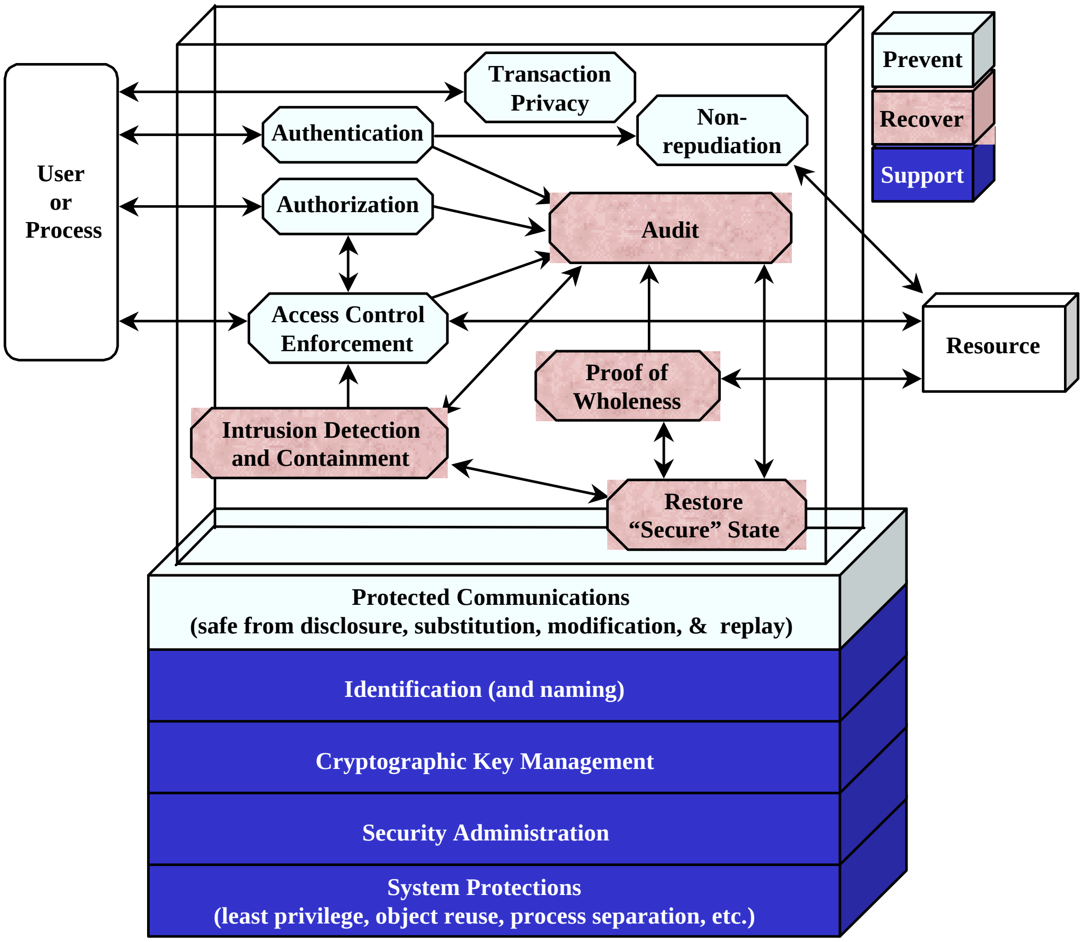
Prevent
Recover
Support
User or Process
Transaction Privacy
Authentication
Non-repudiation
Authorization
Audit
Access Control Enforcement
Resource
Proof of Wholeness
Intrusion Detection and Containment
Restore "Secure" State
Protected Communications
(safe from disclosure, substitution, modification, & replay)
Identification (and naming)
Cryptographic Key Management
Security Administration
System Protections
(least privilege, object reuse, process separation, etc.)

Figure 3 - Technical Security Controls

Supporting:

Supporting controls are, by their very nature, pervasive and inter-related with many other controls. The supporting controls are:

- Identification (and naming) In order to implement many of the other controls, it is essential that both subjects and objects be identifiable. This control provides the capability to uniquely identify users, processes, and information resources.
- Cryptographic key management Cryptographic keys must be securely managed when cryptographic functions are implemented in various other controls. That ability is provided by this control.
- Security administration The security features of the system need to be administered in order to meet the needs of a specific installation and to account for changes in the operational environment. This control provides this needed administration.
- System protections Underlying the various security functional capabilities is a base of confidence in the technical implementation. This represents the quality of the implementation from both the perspective of the design processes used and the manner in which the implementation was accomplished. Some examples of system protections are: residual information protection (also known as object reuse), least privilege, process separation, modularity, layering, and minimization of what needs to be trusted.

Prevention:

These controls can prevent the security breach from ever happening.

- Protected communications In a distributed system, the ability to accomplish security objectives is highly dependent on trustworthy communications. The protected communications control ensures the integrity, availability, and confidentiality of information while in transit. It is the rare situation where all three elements are not essential requirements, with confidentiality being needed at least for authentication information.
- Authentication It is often extremely important to ensure that a claimed identity is valid. The authentication control provides the means to verify the identity of a subject.
- Authorization The authorization control enables specification and subsequent management of the allowed actions for a given system.
- Access control enforcement When the subject requesting access has been validated for access to particular processes, it is still necessary to enforce the defined security policy. The access control enforcement control provides this enforcement, and frequently the enforcement mechanisms are distributed throughout the system. It is not only the correctness of the access control decision, but also the strength of the access control enforcement that determines the level of security obtained.
- Non-repudiation System accountability depends upon the ability to ensure that senders cannot deny sending information and that receivers cannot deny receiving it. Non-repudiation is a control that spans prevention and detection. This control has been placed into the prevention category because the mechanisms implemented prevent the ability to successfully repudiate an action. As a result, this control is typically performed at the point of transmission or reception, rather than later.
- Transaction privacy Both government and private systems are increasingly required to maintain the privacy of individuals using these systems. The transaction privacy control protects against loss of privacy with respect to transactions being performed by an individual.

Detection and Recovery:

Because no set of prevention measures is perfect, it is necessary to both detect security breaches and to take actions to reduce their impact.

- Audit The auditing of security relevant events is a key element for after-the-fact detection of and recovery from security breaches.
- Intrusion detection and containment It is essential to detect insecure situations in order to respond in a timely manner. Also, it is of little use to detect a security breach if no effective response can be initiated. The intrusion detection and containment control provides these two capabilities.
- Proof of Wholeness In order to determine that integrity has been compromised, the ability must exist to detect when information or system state is potentially corrupted. The proof of wholeness control provides this ability.
- Restore 'secure' state When a security breach occurs, the system must be able to return to a state that is known to be secure. That is the purpose for this service.

4.3.2 Management Security Controls

System Security Plan

OMB Circular A-130 Appendix III requires a System Security Plan (see NIST Special Publication 800-18) for each general support system and major application. The plan should document system identification, system type, sensitivity levels and security controls. Recommendations from the risk assessment performed as a part of this guide should be included in the system security plan.

Procedural security controls

Security procedures ensure that the processes vital to the successful fulfillment of organizational goals and missions are performed in line with a base set of requirements. This both provides for an environment that reduces the chance of a security breach and ensures that in the event of an incident, there will be a means to trace the chain of occurrences back to the origin of the problem for future reference. Some examples of procedural controls are:

> The procedures should regulate user access to data by defining levels of authorized access and associating each user with one or more of these levels. For example, regular users will typically be authorized significantly less access than supervisors or system administrators. This procedural control is typically enforced by a combination of technical and non-technical mechanisms. The user's level of access will also typically be documented in authorization memorandums used to verify their identity and the authenticity of their user profile.
>
> For systems relying on passwords for access control, procedures should be in place setting standards for password generation, control, and use.
>
> SDLC activities are also carried out under the guidelines of the security procedures. These include system tests, design reviews and proposed changes to security-related code. Configuration management is also completed under the procedures outlined in the security controls documentation.

If standard procedures are commonly followed, it will be easier to re-trace the events leading to a security incident and possibly even the source of the trouble. Additionally, repeatable, procedural controls boost efficiency in general and help achieve a controlled environment that is less susceptible to threats.

4.3.3 Operational Controls

One of the most effective measures to improve an organization's information security posture is a security education, awareness, and training program. The level of training depends on the degree of responsibility and interaction the person has with the system. The more interaction and individual has with the security-related issues, the more comprehensive is the training need.

Virus protection is important to the enterprise due to the ability of malicious code to cripple everyday operations. While virus-detection mechanisms are technical controls, the effectiveness of such mechanisms is largely determined by the related operational controls. To be effective, virus detection software must be running on all machines and be applied to all the vectors for virus infection such as email, network downloads, and floppy disk transfers. It is important that upgrades of scanning software definition files be performs frequently, on all copies of the scanners to ensure that new viruses are detected.

External storage media should be considered in the security procedures of any organization. It can be a source for unauthorized information dissemination if not physically protected or properly disposed of when no longer needed. Also, careful records and storage requirements will aid in deterring theft and/or illegal duplication of the data contained on peripheral storage media.

The maintenance of any system is vital to its successful operation. However, this process may subject system components to tampering, damage, or theft if not carried out according to specific guidelines. Thus, system maintenance must be scheduled and documented and the personnel who carry it out must be monitored. Additionally, change control and configuration management are essential.

A contingency plan allows organizations to fulfill mission objectives regardless of problems that may occur. Such plans provide for enhanced system resilience via means such as pre-planned rerouting of critical functions, alternate facilities for system operations, and personnel replacement.

An incident response plan ensures that for a wide variety of potential security breaches, there is an applicable, known, and rehearsed procedure to follow in response to each type of breach.

Personnel security controls

Limiting the personnel authorized to use a system is also important. Background checks are one example of such controls and can be extremely helpful in determining if an individual might be a threat to the organization's mission.

Physical security controls

Physical security controls are the third important category of controls that should be employed to mitigate risk to the system. They are the measures put in place to protect information and information systems from compromise, theft, or damage by any of the potential threats.

One type of physical controls is physical barriers such guards, fences, locks, and segmented work areas. Another is the issuance and wearing of badges.

Physical controls such as protective equipment are especially important in mitigating against non-manmade threats, such as fire, flood, earthquakes, etc. Examples of protective equipment that can be installed are voltage regulating transformers, uninterruptible power supplies, and on-site power generators.

4.4. COST-BENEFIT ANALYSIS

It is not usually feasible to implement all possible controls. To determine which ones are required and appropriate for a specific organization, a cost-benefit analysis for each proposed control should be conducted. The cost-benefit analysis can be qualitative or quantitative. Its purpose is to demonstrate that the costs of implementing the controls can be justified by the reduction in the level of risk. In other words, the organization may not want to spend $10,000 on a control to reduce a $200 risk.

The first step in the cost-benefit analysis is to identify the benefits of the controls (increase in mission effectiveness) relative to the cost of implementation and operation of the control. Organizations can analyze the extent of reduction in the level of risk generated by the controls in terms of the two parameters that define the level of risk to the mission: likelihood and impact.

The impact parameter is mission-based and cannot usually be negotiated. However, if cost-effective controls do not exist, then risk mitigation is possible only by modifying the role the system plays in support of the mission. When cost-effective controls can be implemented their purpose is to reduce the likelihood of occurrence, the impact should a security breach occur, or both. This is accomplished by achieving one or more of the following:

- Eliminate some of the flaws and weakness, therefore reducing the number of possible threat-source/vulnerability pairs
- Add a targeted control. For example, a vulnerability that requires physical access to the system remains uncorrected, but administrative controls are implemented to make physical access harder to achieve.
- Reduce the magnitude of the impact of successful exercise of a vulnerability by either limiting the extent of a vulnerability or modifying the nature of the relationship between the system and the organization's mission.

This concept is graphically presented in Figure 4.

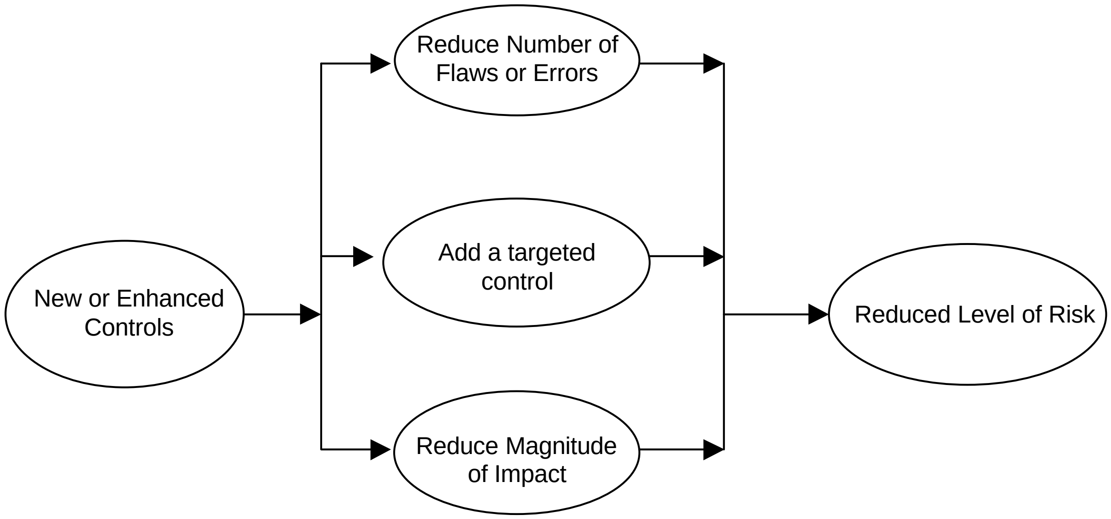

Figure 4 - Control Recommendation Process

The control recommendation process as covered will involve choosing between a combination of technical, management, and operational recommendations for making the organization's security posture more effective.

For example, consider the need to enforce correct entry of security parameters. A technical control may be more complex and expensive than a procedural one, but is likely to be more effective since the enforcement is automated. On the other hand, a procedural control might be accomplished via a simple memorandum to all concerned individuals and an amendment to the security guidelines for the organization. However, trying to ensure that users consistently follow this memorandum is a much tougher task.

After identifying appropriate controls and determining their benefits, the organization will need to determine the associated costs. The costs of implementing controls may include, but are not limited to, the following:

- Hardware and software purchases
- Reduced operational effectiveness if system performance or functionality may be reduced for increased security
- Costs of implementing additional policies and procedures
- Costs of hiring additional personnel to implement proposed policies, procedures, or services.
- Training costs

Finally, the organization will then need to assess the benefits of the controls to the organization in terms of maintaining an acceptable mission posture. Just as there is a cost to implement a needed control, there is a cost for not implementing it. Relating the result of not implementing the control to the organizational mission serves to determine whether it is feasible to forego its implementation. As the amount of acceptable risk is a management decision, the mission process owner must determine what constitutes an acceptable level of mission risk. Once a range of operationally feasible risk levels are created, the control's impact may then be assessed and either included or excluded, depending on the outcome. This range will vary from one organization to the next, but the calculation of the risk present based upon the controls employed will generally be made as follows:

If control would reduce risk more than needed, then see if a less expensive alternative exists.

If control would cost more than the risk reduction provided, then find something else.

If control does not reduce risk enough, then look for more controls or a different control.

If control provides enough risk reduction and is cost-effective, then use it.

It is important to note that the cost to implement a control is often more tangible than the cost of not implementing it. This makes the mission process owner even more critical in the decision whether to implement control measures. It is often only the process owner who can make a determination as to the relative measures of these two, often very different, 'costs'.

4.5 RESIDUAL RISK

Few, if any, systems will ever be completely risk free; every system will have some degree of residual risk. It is the process owner's responsibility to make the final decision about the degree of risk they are willing to accept. This decision should be based on the cost-benefit analysis described in section 4.4 as well as the risk assessment described in section 3. Within the federal government, the acceptance of risk is closely linked with system certification and accreditation.

This process results in a formal approval for the system to become operational or, if it is already on line, to remain so.

4.6 RISK MITIGATION IN THE SDLC

Risk mitigation activities generally begin in the second phase of the SDLC, development and acquisition, when technical, management, and operational security controls for the system are defined. As described above, the controls that are selected should address specific, identified vulnerabilities, or specific identified threat-sources, thereby reducing the overall threat faced by the system. In this phase technical controls are designed into the system. Industry and government alike agree that the beginning of the system life cycle is the best time to address security to ensure cost effective, interoperable solutions.

In the implementation phase, the controls are integrated into the existing system. In the operation and maintenance phase, the controls are put to the test, keeping unwanted and unauthorized incidents from occurring. New controls may be put into place during this phase in response to any number of changes - threats may increase, the criticality of the system to the organization may change, and new vulnerabilities might be discovered.

In the final phase, disposal, the network components are destroyed according to commonly accepted practices, and degaussing and other measures used to keep sensitive residual data from falling into the wrong hands.

5. EVALUATION AND ASSESSMENT

The results of a risk assessment are only the beginning of an ongoing process aimed at reducing the possibility of, or degree to which, the mission operations will be adversely affected by an information technology, security event. In most organizations, the network itself will continually be updated, its components changed, and its software applications replaced with newer versions. In addition, personnel changes will occur and adherence to security policies is likely to change over time. The existence of these variables means that new risks will periodically surface and risks previously mitigated will again become a concern. Thus, the risk management process is ongoing and evolving. There should be a specific schedule, but the process should also be flexible enough to allow changes where warranted. As a rule of thumb, the analysis is usually repeated within 24 months or less. However, certain instances will require an immediate analysis. These include the installation of new equipment, upgrading of software applications or a new system platform installation. From time to time, new employees or the assignment of employees from within the organization to new roles will warrant examination as well. In any event the requirement of OMB Circular A130, Appendix for a review at least every three (3) years must be met.

Periodic reassessments are necessary to maintain an accurate picture of the system's security posture. As these results are reported, changes in policy should be made to better address the weak points in the existing security program. Appendix B provides a suggested outline for documenting the results of a risk assessment in a report or briefing format. The results should be of sufficient detail to allow the organization's management to make informed decisions on the appropriate actions to take in response to the identified risks to their mission. One must try to make comments as specific as possible, especially when recommending controls and remedies. Also, diagrammatic explanations are useful in the implementation of the security controls for less familiar individuals. Above all, one must keep in mind the feasibility of recommendations given the resources available and the risks present.

APPENDIX A: GLOSSARY

TERM DEFINITION

Accountability

The security goal that generates the requirement for actions of an entity to be traced uniquely to that entity. This supports non-repudiation, deterrence, fault isolation, intrusion detection and prevention, and after-action recovery and legal action.

Assurance

Grounds for confidence that the other four security goals (integrity, availability, confidentiality, and accountability) have been adequately met by a specific implementation. “Adequately met” includes (1) functionality that performs correctly, (2) sufficient protection against unintentional errors (by users or software), and (3) sufficient resistance to intentional penetration or by-pass.

Availability

The security goal that generates the requirement for protection against:

- intentional or accidental attempts to (1) perform unauthorized deletion of data or (2) otherwise cause a denial of service or data and
- unauthorized use of system resources.

Confidentiality

The security goal that generates the requirement for protection from intentional or accidental attempts to perform unauthorized data reads. Confidentiality covers data in storage, during processing, and while in transit.

Denial of service

The prevention of authorized access to resources or the delaying of time-critical operations.

Integrity

The security goal that generates the requirement for protection against either intentional or accidental attempts to violate data integrity (the property that data has not been altered in an unauthorized manner) or system integrity (the quality that a system has when it performs its intended function in an unimpaired manner, free from unauthorized manipulation).

IT-related risk

The net mission impact considering (1) the probability that a particular threat-source will exercise (accidentally trigger or intentionally exploit) a particular information system vulnerability and (2) the resulting impact if this should occur. IT related-risks arise from legal liability or mission loss due to:

1. Unauthorized (malicious, non-malicious, or accidental) disclosure, modification, or destruction of information.
2. Non-malicious errors and omissions.
3. IT disruptions due to natural or man-made disasters.
4. Failure to exercise due care and diligence in the implementation and operation of the IT.

IT security goal
See “Security goal”.

Risk
Within this document, synonymous with “IT-related risk.”

Risk analysis
See risk assessment

Risk assessment
The process of identifying the risks to system security and determining the probability of occurrence, the resulting impact, and the additional safeguards that mitigate this impact. Part of risk management and synonymous with risk assessment.

Risk management
The total process of identifying, controlling, and mitigating information system related risks. It includes risk assessment; cost-benefit analysis; and the selection, implementation, test, and security evaluation of safeguards. This overall system security review considers both effectiveness and efficiency, including impact on the mission and constraints due to policy, regulations, and laws.

Security
Security is a system property. Security is much more that a set of functions and mechanisms. Information system security is a system characteristic as well as a set of mechanisms which span the system both logically and physically.

Security goals
The five security goals are integrity, availability, confidentiality, accountability, and assurance.

Threat
The potential for a “threat-source” (defined below) to exercise (accidentally trigger or intentionally exploit) a specific vulnerability.

Threat-source
Either (1) intent and method targeted at the intentional exploitation of a vulnerability or (2) the situation and method that may accidentally trigger a vulnerability.

Threat analysis
The examination of threat-sources against system vulnerabilities to determine the threats for a particular system in a particular operational environment.

Vulnerability
A flaw or weakness in system security procedures, design, implementation, internal controls, etc., that could be exercised (accidentally triggered or intentionally exploited) and result in a violation of the system’s security policy.

APPENDIX B: SAMPLE RISK ASSESSMENT REPORT OUTLINE

EXECUTIVE SUMMARY

Introduction

Begin with a brief description of the team and the analysis process. Reserve greater detail for scope statement below.

Purpose

To protect the accuracy, confidentiality and availability of data or functions and assure that safe and consistently correct procedures are being employed to conduct the work of the organization.

Scope

Describe the elements of the network, its architecture, the system components, users, field site locations (if any), and any other details about the system to be considered in the analysis. Use diagrams here, as they will assist others in understanding the scope of the project.

RISK ASSESSMENT APPROACH

Include in the description whether the organization's approach is to perform an analysis after an event has taken place or if the analysis is considering the likelihood of an event taking place in the future.

System characterization

This is where system resources and information that constitute the system and its boundaries are identified in order to provide the foundation for the remaining steps in the risk assessment process. One must use the system characterization statement to give readers a detailed view of the hardware, software and setup examined. This section should describe the relationship between IT components and the organization's mission processes.

Threat statement

Identifies and explains the existing threats (threat-source/vulnerability pairs) to the system and outlines them specifically in terms of potential problems.

FINDINGS

Each finding must include:
A discussion of the threat-source and vulnerability pair
Identification of existing mitigating security controls
Impact analysis discussion
Risk rating
Recommended controls
Appendices incl. System diagram, etc.

APPENDICES

There should a few descriptive sections to the end of the report. These should include a System Diagram, Glossary of Terms, and a List of Acronyms and Abbreviations. The diagram is particularly important, as it will provide staff and administration with an overall view of the architecture employed by the system, as well as the individual components mentioned in the report. Additionally, a list of key staff members is helpful. Individual contact information, including phone, fax, and E-mail should be included.

APPENDIX C: REFERENCES

NIST Special Publication 800-12, *An Introduction to Computer Security: The NIST Handbook*. October 1995.

NIST Special Publication 800-14, *Generally Accepted Principles and Practices for Securing Information Technology Systems*. September 1996. Co-authored with Barbara Guttman.

NIST Special Publication 800-18, *Guide For Developing Security Plans for Information Technology Systems*. December 1998. Co-authored with Federal Computer Security Managers' Forum Working Group.

NIST Special Publication 800-27, *Engineering Principles for IT Security*. June 2001.

OMB A-130, Office of Management and Budgeting, *Management of Federal Information Resources* Appendix 3, November 2000.