

MANAGEMENT 414

SANS +S™

TRAINING PROGRAM

FOR THE CISSP®

CERTIFICATION EXAM

414.4

Security Architecture and Models, and Operations Security

Copyright © 2006, The SANS Institute. All rights reserved. The entire contents of this publication are the property of the SANS Institute. User may not copy, reproduce, distribute, display, modify or create derivative works based upon all or any portion of this publication in any medium whether printed, electronic or otherwise, without the express written consent of the SANS Institute. Without limiting the foregoing, user may not reproduce, distribute, re-publish, display, modify, or create derivative works based upon all or any portion of this publication for purposes of teaching any computer or electronic security courses to any third party without the express written consent of the SANS Institute.

Preface

The cardinal rule for SANS training is that after you take a course you should be able to apply what you learned directly the day you get back into the workplace. My journey into writing about the 10 Domains started when Stephen Northcutt asked that I lead the development of adding the ISC2 10 Domains of Knowledge into SANS Security Essentials. That is SANS most popular training course and when taught bootcamp style it does an amazing job of helping students become capable of using hands-on techie tools. However, there had been a split in the community whether Security Essentials, which favored technical and pragmatic material, or the ISC2 10 Domains, which favors theory, should be the baseline standard for an information security professional. We were discussing this hotly debated issue in the SANS faculty speaker room over lunch one day and it suddenly dawned on us, why not add the 10 Domains into Security Essentials? Tony Cole, CISSP, was assigned to evaluate Security Essentials and determined that about 60% of the 10 Domains material was already covered in Security Essentials. Clement Dupuis and I were the leads on the project. This was a very successful edit and a number of students have passed their CISSP exams after going through SANS Security Essentials with the ISC2 10 Domains. However, when we added the additional material there was no longer time to cover the application of the material to the workplace; in addition, there are some students who prefer the more formal 10 domain structure.

To best meet the needs of the students, SANS authorized the creation of Management 414, SANS CISSP® 10 Domains +S™, which covers the 10 Domains of Knowledge in a formal 10 domain structure. In the meantime, Clement Dupuis, Stephen Northcutt, Marcus Sachs, Bill Stearns and Joshua Wright are removing some of the 10 Domains material from SANS Security Essentials and returning it back to the original vision for that track, to fully cover the essentials of technical security. Moreover, SANS has insisted that the course teach the application of the 10 Domains in the workplace - something no other 10 Domains course, including ISC2's does. This course meets the SANS promise: what you learn in the course you will be able to apply in the work place. One of the most important things I have learned from Alan Paller, Director of Research, in the years I have been involved with SANS is the importance of community consensus. In order to provide the highest quality training we have recruited experts to review the material and come to consensus on the course content and the application of the information. With the help of Zoe Dias, SANS Faculty Director, we enlisted a total of 68 reviewers from 10 countries. All but two of the reviewers are active CISSP's. The main author for the course, Eric Cole, has been a CISSP for almost 10 years.

SANS enthusiastically applauds the expert work of our technical reviewers/editors:

Monica Anklam, CISSP No. 31995, USA
Alex Arndt, CISSP No. 52343, Canada
Hank Askin, CISSP No. 40792, USA
Anjali Atanacio, CISSP No. 27039, USA
Jason Bevis, CISSP No. 35285, USA
Ron Black, CISSP No. 24245, USA
Anton Bojanec, CISSP No. 24560, Slovenia
Olufremi Bolanle, CISSP No. 51582, Nigeria
Jeff Bontsas, CISSP No. 39135, USA
Derek Browne, CISSP No. 26099, Canada
Sherry Callahan, CISSP No. 21760, USA
Ed Capizzi, CISSP No. 35909, USA
Jim Cate, CISSP No. 37031, USA
Patrick Chan, CISSP No. 40222, Canada
Jerry Chen, CISSP No. 47413, Canada
Daniel Cline, CISSP No. 31366, USA
Chris Cook, CISSP No. 38254, UK

Edwin Covert, CISSP No. 3597, USA
Phil Curran, CISSP No. 31708, USA
Edgar Danielyan, CISSP No. 42834,
UK/Armenia
David Dann, CISSP No. 51571, USA
Gary Delaney, CISSP No. 37636, USA
Sandeep Dhameja, CISSP No. 33585, USA
Joe Dial, CISSP No. 25358, USA
Heinz Durr, CISSP No. 42160, Switzerland
Darin Dutcher, CISSP No. 41299, USA
Rene Evers, CISSP No. 29057, USA
Chris Farrow, CISSP No. 45570, USA
Kenneth Fox, CISSP No. 42293, USA
Roger Fradenburgh, CISSP No. 28099, USA
Brian Freedman, CISSP No. 49504, USA
Donald Glass, CISSP No. 42244, USA
Mark Heinrich, CISSP No. 36190, USA

Lorna Hutcheson, USA
Lawrence Johnson, CISSP No. 25456, USA
Chaiw Kok Kee, CISSP No. 31589, Malaysia
Darrin Lau, CISSP No. 29948, USA
Eliot Leibowitz, CISSP No. 43782, USA
Steven Leong, CISSP No. 30313, Singapore
Chip Meadows, CISSP No. 10070, USA
Sean Mitchell, CISSP No. 36817, USA
Michael Morrell, CISSP No. 36227, USA
Pamela Nottage, CISSP No. 3758, USA
Sanjay Pandit, CISSP No. 44786, USA
John Pao, CISSP No. 29876, USA
Ariya Parsamanesh, CISSP No. 36074, AUS
Stephen Patton, CISSP No. 49746, USA
Robert Pfau, CISSP No. 21572, USA
Gabriel Proulx, CISSP No. 34018, Canada

Jim Purcell, CISSP No. 34519, USA
Andrew Salzman, CISSP No. 25162, USA
Amarottam Shrestha, CISSP No. 41671, AUS
Michael Solomon, CISSP No. 26517, USA
Robert Sorensen, CISSP No. 48304, USA
George Starcher, CISSP No. 34689, USA
Bruce Swartz, CISSP No. 46522, USA
David Taylor, CISSP No. 55890, USA
Brad Towers, CISSP No. 27957, USA
Jill Treu, CISSP No. 43196, USA
Tim Weil, CISSP No. 44250, USA
Deborah Weinstein, CISSP No. 44411, USA
Melody Wilson, CISSP No. 4130, USA
Steven Winterfield, CISSP No. 38096, USA
Kelli Wolfe, USA
Wayde York, CISSP No. 30404, USA

I have had the privilege of the best seat in the house and have really enjoyed working with the CISSP team. I sincerely hope that you benefit greatly from the information in these books and am very interested in your feedback. Please feel free to send me suggestions, corrections or questions to [mgt.414\(a\)sans.org](mailto:mgt.414(a)sans.org).

Eric Cole, Senior Instructor and Research Fellow
The SANS Institute

6. Security Architecture and Design

10 Domains of Knowledge

This section covers Domain 6, the Security Architecture and Models domain.

Domain 6 Overview

- Hardware
- Firmware
- Trusted computing base
- Assurance models

Computer Architecture Fundamentals

To design a secure system, you must first have an understanding of how computers are designed. A modern, general-purpose computer requires several types of components including the following:

- Memory
- Mass storage
- Input device(s)
- Output device(s)
- Central Processing Unit (CPU)
- Software and Operating System

These components work closely together. Indeed, they are sometimes so well integrated that the distinctions between them seem to blur. Still, they are separate components, and the more you know about how they work, the better you can use that knowledge to protect your systems. We probably don't have to tell you much about hard drives and I/O devices because you interact with them directly every time you use a computer. There are, however, some concepts that relate to memory, CPUs, and operating systems that we'd like to discuss. In this domain, we give you an overview of each of these components and explain a few terms and ideas that are important.

Memory Overview

- Memory types
 - Cache memory
 - Cache memory is a relatively small amount (when compared to primary memory) of very high-speed RAM.
 - It holds the instructions and data from primary memory that have a high probability of being accessed during the portion of a program this being executed.
 - Random access memory (RAM)
 - RAM can be directly addressed.
 - Data that is stored can be altered.
 - RAM is volatile due to the fact that the data is lost if power is removed from the system.

Cache logic attempts to predict which instructions and data in main memory will be used by a currently executing program. It then moves these items to the higher-speed cache in anticipation of the CPU requiring these programs and associated data.

Caches can significantly reduce the apparent main memory access time, and thus, increase the speed of program execution.

Dynamic RAM (DRAM) stores the information on parasitic capacitance that decays over time. The data must be periodically refreshed. Refreshing is accomplished by reading and rewriting each bit every few milliseconds. Conversely, static RAM (SRAM) uses latches to store the bits and does not need to be refreshed. Both types of RAM, however, are volatile.

Memory Overview (2)

- Read-only memory (ROM)
 - ROM is used to hold programs and data that should normally not be changed or are changed infrequently.
 - Programs stored on these types of devices are referred to as firmware
- Real or primary memory
 - Primary memory is directly addressable by the CPU.
 - It is used for storage of instructions and data associated with the program that is being executed.
 - It is usually high-speed RAM.

Read only memory (ROM) is non-volatile storage where locations can be directly addressed in basic implementation. Data cannot be altered dynamically. Non-volatile storage retains its information even when the computer loses power. Some ROM is implemented with one-way fusible links and their contents cannot be altered. Other implementations can be altered by various means, but only at a relatively slow rate when compared to normal system reads and writes.

Memory Overview (3)

- Secondary memory
 - Slower memory (such as magnetic disks) that provides non-volatile storage
- Sequential memory
 - Memory from which information must be obtained by sequentially searching from the beginning rather than directly accessing the location

Secondary memory can store large amounts of data at relatively low cost per bit. The cost of such storage is in the speed with which you can access it.

The typical memory hierarchy proceeds from highest speed and highest cost per bit to the lowest speed and lowest cost per bit.

Memory Overview (4)

- Virtual memory
 - Uses secondary memory in conjunction with primary memory to present a CPU with a larger, apparent address space of real memory locations
- Memory protection
 - Means to prevent one program from accessing and modifying the memory space contents that belong to another program

By swapping memory from lower disk-based memory to higher-speed memory, the CPU can be made to "see" a larger amount of primary memory than physically exists. This virtual memory is implemented by the operating system or by hardware mechanisms.

Memory protection is where you protect one part of memory that is associated with Program A from being accessed and written to by Program B.

RAM

Random Access Memory (RAM):

- Volatile memory
- Data lost when power is lost
- Dynamic versus static
- Main memory

Memory

Computers run programs and operate on data. To do that, they need to have some place to store information. That place is the system's *memory*. It provides temporary storage for programs and the data they need to run. Modern computer systems use *Random Access Memory* (RAM), meaning that the system can directly read or write a byte stored in memory without affecting any of the other bytes. RAM is volatile, and the chips need continual power to preserve their contents. When the computer loses power, the data stored in RAM is lost. In common use, RAM usually refers to the system's *main memory*, that is, the memory that holds the running operating system, applications, and data. There are other types of RAM in a computer, though, which we will talk about later.

Most computers also possess a certain amount of *Read Only Memory* (ROM). ROM is similar to RAM in that the bytes can be read individually, but there are two important differences. First, ROM cannot be modified, only read from. Second, unlike RAM, ROM does not require continual power to preserve its contents. If you turn off the power to your computer, the ROM still retains its data. That's why ROM is typically used to store critical programs, such as the one that starts the boot process when the system powers up. Most systems contain a large amount of RAM and just a small amount of ROM so unless we indicate otherwise, you can assume that when we say memory, we mean RAM.

Dynamic Versus Static

Most computers use two different types of RAM. The main memory is usually *dynamic RAM* (DRAM). DRAM is considered dynamic because the system needs to continually refresh the data stored or it is lost. The data is rewritten thousands of times each second, otherwise it would decay and become unusable. This sounds like a useless piece of technology, but in reality, it is quite workable. After the CPU stores data in DRAM, the system's supporting electronics automatically take care of refreshing it, freeing the rest of the system to go on to do other things. The continual refresh process makes accessing the memory a little slower because the access can occur only between refresh cycles. However, DRAM is inexpensive, which more than makes up for its other faults. Because typical computers are equipped with hundreds (or sometimes thousands) of megabytes of main memory, inexpensive DRAM keeps the price down to a manageable level.

Types of RAM

- DRAM - Dynamic RAM
 - Refreshed on a regular basis
 - Cheapest and most common
- SRAM - Static RAM
 - Very fast
 - Less amount
 - Cache
 - Used to hold a copy of program instructions and data that are likely to be needed next by the processor in main storage.
- EDO RAM - Extended Data Output RAM
 - Improves read time
- SDRAM - Synchronous DRAM
 - Synchronized with clock speed
 - Increases the number of instructions that can be performed

Main memory is not the only place your computer uses RAM. Sometimes DRAM is just *too* slow for the task at hand. Most computers also include a small amount of *static RAM (SRAM)*. As long as it is supplied with electricity, SRAM keeps its contents safe without requiring constant refresh cycles. This means it is much faster because the system can immediately retrieve data stored in SRAM without waiting for a refresh cycle to complete. Unfortunately, SRAM is a lot more expensive than DRAM, which makes it unsuitable for use as main memory. SRAM is typically used as *cache memory*, a special fast storage buffer that holds copies of data or instructions likely to be requested soon by the CPU. Memory caching improves performance because many programs loop over the same data or program instructions several times. If this information is kept in cache, the computer can access it much more quickly than if it had to fetch it from the comparatively slow main memory.

Memory Addressing

- Memory isolation
- TOC/TOU protection
- CPUs can address memory in various ways:
 - By directly specifying the address (direct or absolute addressing)
 - By addressing the registers within a CPU (register direct addressing)
 - By addressing the register for the data's address in main memory (register indirect addressing)
 - By using an index register (indexed addressing)
 - By addressing the desired location of the program in memory (indirect addressing)

Memory Addressing

The theoretical ability to store and retrieve data in memory is useless without the ability to tell the memory system *where* to store or fetch the data. Each byte in memory is assigned a unique *address* that distinguishes it from the other bytes. There are several ways for the system to specify the address, but in the end, they all refer to the same location. These include:

- **Direct addressing:** This is the simplest form of addressing. The system knows the exact location of the data in memory and requests the data by passing the actual address to the memory subsystem. Direct addressing is sometimes referred to as *absolute addressing*.
- **Register direct addressing:** The CPU contains tiny memory areas known as *registers*. Registers are temporary storage for the task the CPU works on at that instant. To operate on values from main memory, the values must first be loaded into a register. Register direct addressing is slightly different from the other types of addressing in that it never refers to main memory. It simply refers to a specific register that already contains the required data.
- **Register indirect addressing:** In this addressing mode, the system looks in the specified register for the data's address in main memory.

Virtual Memory

- Virtual addresses are used by some operating systems to store data and instructions.
- The OS *maps the* virtual addresses to real addresses
 - *Paging occurs* when the OS copies from virtual memory to main memory.
 - *Page fault* is a request that results in paging.
- Locked memory
 - Prevents data from being written to virtual memory

Virtual Memory (Swap)

Up to this point, the different types of memory we have discussed correspond to physical hardware present in the system. In this section, we discuss something a bit different: virtual memory. *Virtual memory* (VM) is a set of memory addresses managed by the operating system that don't correspond directly to physical memory. To the CPU, virtual memory looks like physical memory. It can hold both programs and data, but using virtual memory gives the operating system the choice of where to store the data.

With physical memory, an address corresponds directly to a piece of hardware. If the physical address is specified, this is where the system will place the data. Physical addressing is very straightforward. Usually the operating system manages this sort of thing. Using virtual memory, it maps the virtual address space into the physical address space. When the system needs to access a memory address, the OS can translate the virtual address into a physical one and fetch the data from the correct location. Because virtual memory hides the actual storage location from the hardware, the OS is free to store the data wherever it likes, including a mass storage device, such as a hard drive. This lets the system address a larger amount of memory than it actually contains. For example, even if the system physically contains only 256 MB of main memory, virtual memory would allow it to hold a theoretically unlimited amount of data in memory.

The operating system uses the system's main memory as a cache to hold the most recently or most frequently accessed data, whereas the rest of the data is stored on the hard drive. When the CPU issues a request for more data, the OS first checks to see if the data is already stored in main memory. If it is, it notifies the system of the data's physical address and allows it to be read. If the data is not present in main memory, the OS fetches the data from the hard drive and copies it into main memory (it might first flush older data from main memory back to the disk to make room). After the data is in main memory again, the OS notifies the system of its physical address and processing continues. When the CPU wants to write to virtual memory, it writes to a physical address specified by the OS, which might then either keep the data in RAM for some time or flush it to the disk. This process of moving data to and from the hard disk is known as *paging*, and a request that results in paging is known as a *page fault*.

ROM

- Read Only Memory (ROM)
 - Can only be read
 - Allows system to be booted
- Firmware
 - Flash ROM
 - Can be updated with a flash program to correct bugs in the controller software or to improve performance of a device

Firmware

Firmware is a type of program somewhere between hardware and software. Firmware does not allow modification (writing or deleting) of data. Firmware is generally the controlling software for a device that is placed in a special type of ROM, which can be updated as new releases become available. We told you earlier that ROM means *Read Only Memory*, and as such you normally cannot update information stored in ROM.

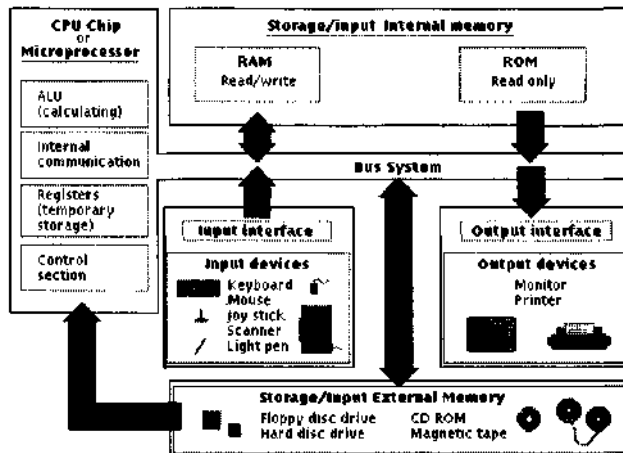
Types of ROM

- PROM - Programmable ROM
 - Modifiable once
 - Firmware
- EPROM - Erasable and Programmable ROM
 - Can be erased and reprogrammed
 - Not the norm
- EEPROM - Electrically Erasable ROM
 - Flash memory
 - Can be written
- Programmable Logic Devices (PLD)
 - Integrated circuit that can be modified programmatically
 - General technology for all EPROM

Firmware is the exception, because it is stored in special ROM chips called *Programmable Read Only Memory* (PROM). A PROM is like regular ROM, except the contents are blank when it is manufactured. It's meant for a system designer to program later. After it is written, a standard PROM is immutable, which makes it useless for firmware, but there is another type called an *Electrically Erasable PROM* (EEPROM), sometimes also called *flash memory*. EEPROMs can be rewritten, although it is a slow process. Most computer BIOS chips are actually EEPROMs, so they can be updated as the manufacturer corrects defects or adds new BIOS features.

All types of PROMs, including EEPROMs, are actually special cases of a more general sort of technology, the *Programmable Logic Device* (PLD). Although PROMs are simply a type of memory, other PLD devices offer fully programmable logic circuits, making them ideal for prototyping new chip designs. Other common types of PLDs include the *Programmable Logic Array* (PLA), *Programmable Array Logic* (PAL), and the *Generic Array Logic* (GAL). GAL is sometimes also called *Gate Array Logic*. PAL and GAL devices are especially well-suited for low-complexity, low-cost applications. GAL chips are particularly popular because, unlike PALs, they are reprogrammable. PLA devices, on the other hand, offer the highest level of flexibility, which comes at a higher cost.

Computer Architecture (Map of Targets)



This diagram shows you the complete layout of a system. The key areas of the system are shown along with the interaction each component has with the others. When looking at this architecture, there are different ways to analyze it. First, you can look at it to better understand how the system operates and why certain things occur. The more you understand how something works, the better chance you have to fix it. However, to fix a problem, you have to know what the problem is. The second way you can look at the architecture is to determine the attack vectors that one can use to compromise the system. If you understand the various points of attack, you can use that knowledge to secure the various components.

The CPU

The CPU contains:

- Arithmetic/logic unit—data transfer operations, arithmetic operations, data editing, and decision making
- Control unit—coordinates system activities during execution of code
- Primary storage memory unit—stores instructions and data for current programs in use

The CPU

You already know that the CPU is considered the brains of the computer. Just as a person's brain contains specialized regions that perform different functions, virtually every CPU is composed of at least two parts: the *control unit* and the *Arithmetic Logic Unit* (ALU). The control unit has a simple job: It manages the flow of execution in a program. It decides which instruction to process next, fetching them from memory, executing them, and storing the results. During execution, the control unit calls upon the ALU to perform whatever arithmetic and logical operations the program calls for.

Computer Architecture

Instruction cycle:

- Consists of two phases: fetch and execute
 - Fetch phase
 - CPU presents the address of the instruction to memory.
 - CPU retrieves the instruction located at that address.
 - Execute phase
 - The instruction is decoded and executed.
- The cycle is controlled by and synchronized with the CPU clock signals.

Multiple clock signals, known as multi-phase clock signals, are used in order to refresh dynamic RAM.

Some instructions might require more than one machine cycle to execute.

CPU Terms

- Pipelining
 - Combines the steps of different instructions
- Complex-Instruction-Set-Computer (CISC)
 - Performs many operations per instruction
- Reduced-Instruction-Set-Computer (RISC)
 - Simpler instructions using fewer cycles
- Interrupt
 - Allows for interruption of CPU execution

Speaking of instructions, this is probably a good time to mention that there are two basic types of *instruction sets*. An instruction set is just what it sounds like: a set of the low-level instructions a CPU knows how to execute. Most personal computers are *Complex Instruction Set Computers* (CISC), which means the CPUs include a wide variety of instructions, some of which are general-purpose and some of which are for specialized use, such as Intel's MMX and AMD's 3DNow multimedia technologies. A CISC CPU offers programmers a lot of flexibility with relatively little effort.

Reduced Instruction Set Computers (RISC), on the other hand, attempt to pare things down to their basics. RISC designers concentrate on making a small instruction set as efficient as possible. This boosts performance, but places more burden on the programmer. Of course, this burden is usually borne by the compiler writers and operating system manufacturers, not by the application programmers or end users. RISC workstations are popular among scientific and technical users.

Of course, a modern CPU architecture is vastly more complex than this description. In fact, the subject can easily consume several volumes. For our purposes, it is enough to know the difference between the control unit and the ALU, and to know that not all CPU instruction sets are created equal.

Scalar processor

- Executes one instruction at a time

Superscalar processor

- Enables concurrent execution of multiple instructions

Pipelining

- Combines the steps of different instructions

Timing and Instruction Execution

The control unit operates in units of time known as *clock cycles*. The system's supporting electronics generate a timing signal that synchronizes the system, including the CPU. Older processors would typically execute only one instruction per clock cycle, the definition of a *scalar* processor. Newer CPUs are capable of executing several instructions per cycle, and are known as *superscalar* processors. This performance boost is accomplished by adding additional ALUs, each of which can operate in parallel with the others. For example, a single additional ALU allows the CPU to simultaneously perform an additional operation and a logical comparison instruction.

Another common way to boost performance in the CPU is through *pipelining*. Pipelining takes advantage of the fact that there are several stages involved in executing a single instruction. In general terms, the instruction (and the data on which it operates) must be fetched from main memory, decoded to see what it is, and then executed. Finally, the result must be stored somewhere, usually either in a register or in main memory. This is a simple example of an execution path that consists of four stages; the actual stages vary somewhat from processor to processor.

Pipelining involves keeping pieces of the CPU busy processing instructions in the various stages. While the CPU is busy in one stage, resources for the other stages would normally have to wait. With pipelining, however, each stage is kept occupied with a different instruction. During execution, Instruction 1 passes through the fetch stage and progresses into the decode stage. While it is being decoded, Instruction 2 can enter the fetch stage. When Instruction 1 finishes the decode stage and enters the execution stage, Instruction 2 can proceed to the decode stage and Instruction 3 can enter the pipeline at the fetch stage. In this way, the CPU takes maximum advantage of its different modules. If we assume an average completion time for each stage of 10 milliseconds (ms), a pipelined processor can complete one instruction every 10 ms after its pipeline is full, whereas a non-pipelined processor takes 40 ms to process the same instruction. What a difference!

Most CPUs these days combine both superscalar processing and pipelining. In these CPUs, a pipeline is set up as we have just described, with the execute phase making use of multiple ALUs to provide the superscalar processing.

CPU Terms (3)

- **Multitasking**
 - Executes multiple tasks at the same time on one CPU
- **Multiprocessing**
 - Executes multiple programs at the same time on multiple processors
- **Multithreading**
 - Allows more than one user to utilize the system at the same time
- **Multiprogramming**
 - Interweaves execution of more than one program

Multitasking and Multiprocessing

Many people confuse the terms multitasking and multiprocessing. When the CPU can process more than one user program at the same time (or virtually the same time) it is called *multitasking*. If the computer has more than one CPU and it can execute instructions in parallel, it is called *multiprocessing*.

Some operating systems, such as Microsoft Windows NT/2000/XP, are *Symmetrical Multiprocessing Systems* (SMP). This means that they support more than one processor. SMP systems also have an interesting feature: The CPUs share the processing of system processes and application processes equally. In an SMP environment with two processors, system tasks and application tasks are divided equally between both CPUs.

Asymmetrical Multiprocessing Systems (AMP), on the other hand, operate differently. In an AMP system, one processor will take care of the system processes and the other processor(s) will run the applications.

Storage Devices

Types of storage devices:

- Primary
- Secondary
- Virtual
- Write once read many (WORM)
- Volatile
- Non-volatile

The general types of storage devices are:

- Primary
- Secondary
- Virtual
- Write once read *rt*r*** (WORM)
- Volatile
- Non-volatile

Types of Systems

- Client/server system
- Open system
- Closed system (proprietary)
- Shared system
- Standalone system

In order to secure a system you have to understand the general type of system and its requirements.

The key systems differentiators are:

- Client server system
- Open system
- Closed system (proprietary)
- Shared system
- Standalone system

Operating System

- The operating system (OS) is the heart of the computer and is loaded by a boot program.
- Mainframe is referred to as initial program load (IPL).
- GUI - graphical user interface
- OS services include program execution, system access, error detection, and accounting.
- Process states include run, wait, sleep, and interrupts.
- Control computer operations and resources
 - Memory management
 - Process management
 - File management
 - I/O management

The operating system (OS) is the heart of the computer. It is loaded by a boot program and controls everything that happens with the hardware and brings the hardware to life.

It controls computer operations and resources through the following calls:

- Memory management
- Process management
- File management
- I/O management

OS States

- User
 - Layer in the operating system where user applications run
- Privileged
 - Protected (or kernel) area of the operating system responsible for memory, process, disk, and task management

Operating System Structure

The *kernel* is the essential nucleus of an operating system, the core that provides basic services for all other parts. A kernel can be contrasted with a *shell*, the outermost part of an operating system that interacts with user commands. Typically, a kernel includes an interrupt handler that handles all requests that compete for the kernel's services, a scheduler that determines which programs share the kernel's processing time in what order, a virtual memory manager, and a supervisor that gives use of the computer to each process when it is scheduled.

Applications can request kernel services by using a set of program interfaces known as *system calls*. When the kernel is executing on a CPU, the system is operating in *privileged mode*. That is, it can interface directly with other parts of the OS and view all the internal data structures. On the other hand, user applications run in *user mode* and must rely on the system call interface to request services from the kernel.

Because the code that makes up the kernel is needed continuously, it is loaded into protected memory so that it will not be overlaid with other less frequently used parts of the operating system. In a virtual memory system, for example, the kernel would never be swapped out to the disk, but would remain in physical RAM at all times.

OS Protection Mechanisms

- **Layering**
- **Abstraction**
- **Process isolation**
- **Hardware segmentation**

OS Protection Mechanisms

One of the most important concepts in the design of secure systems is the concept of defense-in-depth. This is no different when it comes to the OS protection mechanisms. Attackers will try to attack the core of the computer system—the operating system — so protection of this important component is critical. Common OS protection mechanisms include *layering*, *abstraction*, *process isolation*, and *hardware segmentation*.

Layering is the organization of functions into separate components, each of which interacts with the others in a sequential way. Each layer will interface only with the layer above it and the layer below it and should work independently. If one layer in the system fails, it should not affect the other layers. *Abstraction* is the process of finding commonality in different objects, and then exploiting it to make the objects simpler to manage. The ultimate goal is to reduce complexity and to hide the inner workings of the system. A good example of this is a system call named *kill()*, whose purpose is to stop processes from running. All processes on the system share a common meta-information structure that tells the kernel what state the process is in and where its code lies. There is a lot of detail the programmer shouldn't need to worry about, so the OS abstracts the notion of a process into a *process ID*, which is an integer that uniquely identifies a particular process. Passing this process ID to the *killQ* system call is enough to cause the system to kill the process without bothering the programmer with a lot of needless detail or allowing him to look into the inner workings of the OS.

Process isolation is the mechanism for making sure that one process does not adversely affect another process or the operating system itself. In UNIX or Windows NT/2000/XP, for instance, each process runs in its own protected address space. This way, if something prevents the application from operating, it should not affect other processes running on the machine. Perhaps you remember the old DOS days, when one application could lock up the entire system, forcing the dreaded Ctrl-Alt-Del or a hard reboot. Process isolation ensures that one rogue process cannot drag the entire system down. It's also a necessary precursor to multi-user security. Without it, user John's process could look into the address space of Jane's e-mail application and pick out the password she just entered.

Ring Layer Protection

- A common protection scheme is the use of protection rings:
 - Ring 3: Applications and programs
 - Ring 2: I/O drivers and utilities
 - Ring 1: Operating system components that are not part of the kernel
 - Ring 0: Operating system kernel
- The Security Reference Monitor enforces the access controls on objects in the ring.

Ring Layer Protection

Some operating systems model their security framework on the concept of *rings*. A ring is a group of processes that share common security characteristics because they usually perform similar functions for the OS. These systems provide strict boundaries and definitions of what processes should work within each ring. The trusted, and therefore critical, components of the system increase as you travel from the outside to the innermost ring.

Operating system functions, memory access functions, and device drivers usually operate in the inner ring because they need to directly access hardware components. Applications usually operate in the outer ring. Following is a simple example of a ring protection scheme. Remember that this list is ordered so that the least trusted components are in the outermost (higher-numbered) rings:

- Ring 3: Applications and programs
- Ring 2: I/O drivers and utilities
- Ring 1: Operating system components that are not part of the kernel
- Ring 0: Operating system kernel

Windows NT/2000/XP uses the notion of a *Security Reference Monitor* (SRM) that resides in ring 0. Its job is to look at the access token of user objects trying to access resources. It compares the tokens to the resource's access control list and decides whether or not to grant the request. The access token is comprised of the user's *Security ID* (SID) and the SIDs of all of the groups for which the user belongs.

Programming Languages

Programming languages:

-Types of languages

- Machine
 - Executed directly by the computer
 - Relatively difficult to write
- Assembly
 - Mnemonics that have a one-to-one correspondence to machine language instructions
- High level
 - Easy-to-understand
 - One-to-many translation to assembly language

The general types of programming languages are machine, assembly, and high level.

A few definitions might be helpful:

- An *assembler* generates *object code*; original assembly code is called the *source code*.
- A *disassembler* translates machine language into assembly language.

Machine language comprises the 1 and 0 patterns that the CPU recognizes. For example, if we had a simple 4-bit computer, the pattern 1001 might represent the "add" instruction. Therefore, the programmer would have to write "1001" for an add instruction in the program. Now, using a simple translating program called an *assembler*, the programmer could write "add" and the assembler would generate the pattern 1001, which would be recognized by the computer. This translation would be a one-to-one translation, wherein, one assembly language instruction yields one, 4-bit pattern. With a high level language, the programmer can write a complex instruction such as "square root" and a program called a *compiler* will generate multiple assembly language instructions to accomplish the square root command.

Assemblers, Compilers, and Interpreters

- An **assembler** translates an assembly language program into a machine language program.
- A **disassembler** translates machine language into an assembly language program.
- A **compiler** translates a high-level program into machine language.
- An **interpreter** translates program commands one at a time.

An *assembler* translates an assembly language program into a machine language program. A *disassembler* translates machine language into an assembly language program. A *compiler* translates a high-level program into machine language. An *interpreter* translates program commands one at a time. Interpreted code is more secure than compiled code because a compiler generates large amounts of machine code that is difficult to examine for vulnerabilities.

Database

Database language types and functions:

- Permit external access to database management systems (dbms)
- Data definition language (ddl)
 - Defines database schema
- Data manipulation language (dml)
 - Examines and manipulates contents of a database

One can view a database and the language that controls it as a way to program some operations.

Two key terms to be familiar with are *data definition language* (ddl), which defines database schema, and *data manipulation language* (dml), which examines and manipulates contents of a database.

Certification and Accreditation

- Certification
 - Comprehensive evaluation of the technical and non-technical security features of an information system and other safeguards
 - Establishes the extent to which a particular design and implementation meets the set of specified security requirements
- Accreditation
 - Formal declaration by a Designated Approving Authority (DAA) where an information system is approved to operate in a particular security mode using a prescribed set of safeguards at an acceptable level of risk
 - Accreditors responsible for:
 - Evaluating certification evidence
 - Deciding on acceptability of application security safeguards
 - Approving corrective actions
 - Insuring corrective actions are accomplished
 - Issuing accreditation statement

In its second draft, NIST special publication 800 - 37 makes the distinction between certification and accreditation clear when it states, "Security certification is the comprehensive evaluation of the management, operational, and technical security controls in an information system."

"Security accreditation is the official management decision to authorize operation of an information system... By accrediting an information system, the agency official is not only responsible for the security of the system, but is also accountable for adverse impacts to the agency if a breach of security occurs."

At the core of 800-37, NIST offers U.S. government agencies and their senior officials a guideline for testing and accepting federal computer systems.

Certification and Accreditation (2)

- Two U.S. Defense and government certification and accreditation standards for the evaluation of critical information systems
 - Defense Information Technology Security Certification and Accreditation Process (DITSCAP)
 - The National Information Assurance Certification and Accreditation Process (NIACAP)

There are two U.S. Defense and government certification and accreditation standards for the evaluation of critical information systems: one for DoD (Department of Defense) and one for non-DoD government agencies.

DITSCAP is being replaced by DIACAP.

NIST (National Institute of Standards and Technology) has recently developed certification and accreditation guidelines for non-government agencies.

DITSCAP

- DoD directive 5200.40 "DoD information technology security certification and accreditation process (DITSCAP)"
 - Established DITSCAP as the standard C&A process for the Department of Defense
- Objective
 - Establish a DoD standard infrastructure-centric approach that protects and secures the entities comprising the defense information infrastructure (DII)
- Activities presented in DITSCAP standardize the C&A process for single IT entities that leads to more secure system operations and a more secure DII.
- DITSCAP applies to:
 - The Office of the Secretary of Defense (OSD)
 - The military departments
 - The Chairman of the Joint Chiefs of Staff
 - The combatant commands
 - The Inspector General of the Department of Defense (IG, DoD)
 - The defense agencies
 - The DoD field activities, their contractors, and agents

DITSCAP establishes a standard process, a set of activities, general task descriptions, and a management structure to certify and accredit the IT systems that will maintain the required security posture.

DITSCAP applies to:

- The Office of the Secretary of Defense (OSD)
- The military departments
- The Chairman of the Joint Chiefs of Staff
- The combatant commands
- The Inspector General of the Department of Defense (IG, DoD)
- The defense agencies
- The DoD field activities, their contractors, and agents

DITSCAP Phases

- Phase 1: Definition
- Phase 2: Verification
- Phase 3: Validation
- Phase 4: Post-accreditation



There are 4 phases to DITSCAP that will be discussed over the next several slides:

- Phase 1: Definition
- Phase 2: Verification
- Phase 3: Validation
- Phase 4: Post-accreditation

DITSCAP Phase 1

Phase 1: Definition

- Focuses on understanding the mission, the environment, and the architecture in order to determine the security requirements and level of effort necessary to achieve accreditation
- Tasks:
 - Document mission need
 - Registration
 - Negotiation
- Product is final Phase 1 SSAA

Phase 1: Definition

The first phase focuses on understanding the mission, the environment, and the architecture in order to determine the security requirements and level of effort necessary to achieve accreditation.

Key tasks are:

- Document mission need
- Registration
- Negotiation

The outcome of Phase 1 is the SSAA or System Security Authorization Agreement.

System Security Authorization Agreement (SSAA)

- The SSAA is a formal agreement among the DAA(s), certifier, user representative, and program manager.
- The objective is to establish an evolving, yet binding agreement on the level of security required before the system development begins or changes are made to a system.
- Used to:
 - Guide actions
 - Document decisions
 - Specify IA requirements
 - Document certification tailoring and level of effort
 - Identify possible solutions
 - Maintain operational systems security after accreditation
- The SSAA becomes the baseline security configuration document.

The SSAA is a formal agreement among the DAA(s), certifier, user representative, and program manager and becomes the baseline security configuration document.

The objective is to establish an evolving, yet binding agreement on the level of security required before the system development begins or changes to a system are made. This is then used to:

- Guide actions
- Document decisions
- Specify IA requirements
- Document certification tailoring and level of effort
- Identify possible solutions
- Maintain operational systems security after accreditation

System Security Authorization Agreement (SSAA) (2)

- Describes the operating environment and threat
- Describes the system security architecture
- Establishes the certification and accreditation boundary of the system to be accredited
- Documents the formal agreement among the DAA(s), certifier, program manager, and user representative

The SSAA documents all requirements necessary for accreditation. It minimizes documentation requirements by consolidating applicable information into the SSAA (security policy, concept of operations, architecture description, test procedures, etc). It also documents test plans and procedures, certification results, and residual risk.

DITSCAP Phase 2

Phase 2: Verification

- Verifies the evolving or modified system's compliance with the information agreed on in the system security authorization agreement (SSAA)
- Tasks:
 - Continuing refinement of the SSAA
 - System development or modification
 - Certification analysis
 - Analysis of the certification results

Phase 2: Verification

During this phase the evolving or modified system's compliance with the information agreed on in the system security authorization agreement (SSAA) is verified. The SSAA establishes an evolving, yet binding agreement on the level of security required before system development begins or changes are made to a system.

DITSCAP Phase 3

Phase 3: Validation

-Validates the compliance of a fully integrated system with the information stated in the SSAA

-Tasks:

- Review of the SSAA
- Evaluation of the integrated IT system
- Certification and accreditation

Phase 3: Validation

The focus during this phase is to validate the compliance of a fully integrated system with the information stated in the SSAA. It is composed of the following tasks:

- Review of the SSAA
- Evaluation of the integrated IT system
- Certification and accreditation

DITSCAP Phase 4

Phase 4: Post Accreditation

- Includes the activities that are necessary for the continuing operation of an accredited IT system in its computing environment and for addressing the changing threats that a system faces throughout its life cycle
- Tasks
 - Ongoing maintenance of the SSAA
 - System operations
 - Change management
 - Compliance validation

Phase 4: Post Accreditation

This phase includes the activities that are necessary for the continuing operation of an accredited IT system in its computing environment and for addressing the changing threats that a system faces throughout its life cycle. The following are the key tasks:

- Ongoing maintenance of the SSAA
- System operations
- Change management
- Compliance validation

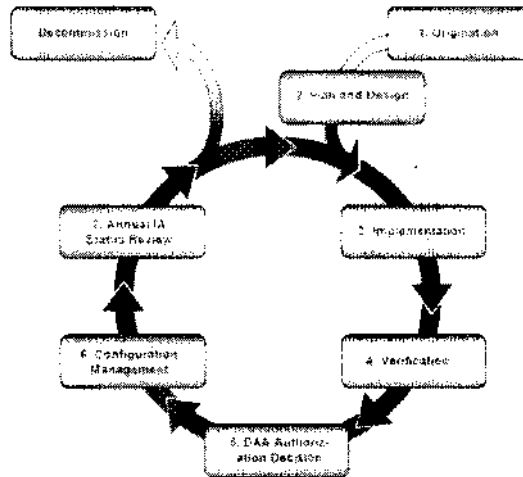
DIACAP

- DoD Information Assurance C&A Process (DIACAP)
- DITSCAP being replaced by DIACAP
 - With increasing levels of connectivity among DoD, international partners, NATO, and state governments, as well as the rise of the global information grid (GIG), a more time- and cost-effective method of C&A was needed.
- Builds upon DoD 8500.1, 8500.2 guidelines
- Utilizes new automated procedures that lower average time of developing a C&A package from six months down to only a matter of days
- Process comprised of seven basic steps

DITSCAP is being replaced by DoD Information Assurance C&A Process (DIACAP). With increasing levels of connectivity among DoD, international partners, NATO, and state governments, as well as the rise of the global information grid (GIG), a more time- and cost-effective method of C&A was needed.

DIACAP builds upon DOD 8500.1, 8500.2 guidelines. It utilizes new automated procedures that lower the average time of developing a C&A package from six months down to only a matter of days. The process is comprised of seven basic steps.

DIACAP (2)



This Slide shows the seven basic steps that DIACAP is composed of:

1. Origination
2. Plan and Design
3. Implementation
4. Verification
5. DAA Authorization Decision
6. Configuration management
7. Annual IA Status Review

NIACAP Certification and Accreditation

- Establishes the minimum national standards for certifying and accrediting national security systems
- Provides a standard set of activities, general tasks, and a management structure to certify and accredit systems that maintain the information assurance and the security posture of a system or site
- Designed to certify that:
 - The information system meets the documented accreditation requirements.
 - The system will continue to maintain the accredited security posture throughout the system's life cycle.

NIACAP certification and Accreditation standards are defined by national security telecommunications and information systems security instruction (NSTISSI) no. 1000

NIACAP provides guidance on how to implement the NSTISSP no. 6 policy , which establishes the requirement for federal departments and agencies to implement a C&A process to ensure that national security systems at all U.S. government executive branch departments, agencies, and their contractors and consultants meet the requirements of the NSTISSI no. 6.

The process is started when the concept design of a new information system or modification of an existing system is begun in response to an identified business case, operational requirement, or mission need.

Any security relevant changes should initiate the NIACAP for any existing or legacy information system.

NIACAP Roles

NIACAP Roles:

- Program manager
- Designated approving authority (DM)
- Certification agent (certifier)
- User representative



The following are the 4 core roles in NIACAP:

- Program manager
- Designated approving authority (DAA)
- Certification agent (certifier)
- User representative

The DAA is also referred to as the accreditor.

Program Manager

Represents the interests of the system in areas such as:

- Acquisition
- Life cycle schedules
- Funding responsibility
- System operation
- System performance
- Maintenance

The program manager is responsible for the interests of the system throughout the life cycle (cost, schedule, and performance of the system development).

This person ensures that the security requirements are integrated in a way that will result in an acceptable level of risk to the operational infrastructure as documented in the System Security Authorization Agreement (SSAA).

The program manager keeps all NIACAP participants informed of life cycle actions, security requirements, and documented user needs.

DM

- Primary government official responsible for implementing system security
- An executive with the authority and ability to balance the needs of the system with the security risks
- Determines the acceptable level of residual risk for a system and must have the authority to oversee the budget and IT business operations of systems under his/her purview

Based on the information available in the SSAA, the DAA can grant the accreditation, an interim approval to operate (IATO), or may determine that the system's risks are not at an acceptable level and the IT system is not ready to be operational.

Certification Agent/Certifier

- Provides the technical expertise to conduct the certification throughout the system's life cycle based on the security requirements documented in the SSAA
- Determines the existing level of residual risk and makes an accreditation recommendation to the DAA
- Determines whether a system is ready for certification and conducts the certification process
- At completion of the certification effort, the certifier
 - Reports the status of certification
 - Recommends to the DAA whether or not to accredit the system based on documented residual risk

The certification agent or certifier is the technical expert that documents tradeoffs between security requirements, cost, availability, and schedule to manage security risk.

To avoid conflicts of interest, the certifier should be independent from the organization responsible for the system development or operation. This ensures that the certifier can provide the most objective information possible for the DAA to make accreditation decisions.

User Representative

- Concerned with system availability, access, integrity, functionality, performance, and confidentiality as they relate to the mission environment
- Responsible for the identification of operational requirements
- Responsible for the secure operation of a certified and accredited is
- Represents the user community
- Assists in the C&A process:
 - Development, modification, integration, acquisition, and deployment

The operational interests of system users are vested in the user representative. The user representative defines the system's operational and functional requirements and is responsible for ensuring that the user's operational interests are maintained throughout system. This person functions as the liaison for the user community throughout the life cycle of the system.

NIACAP Accreditation

Three types of NIACAP accreditation:

- Site accreditation
 - Evaluates the applications and systems at a specific, self-contained location
- Type accreditation
 - Evaluates an application or system that is distributed to a number of different locations
- System accreditation
 - Evaluates a major application or general support system

NIACAP accreditation is used for a major system, where a major application could be SAP or some other large application.

Three types of NIACAP accreditation are:

- **Site accreditation**
Evaluates the applications and systems at a specific, self-contained location
- **Type accreditation**
Evaluates an application or system that is distributed to a number of different locations
- **System accreditation**
Evaluates a major application or general support system

Information Security Models

The access matrix:

- Provides access rights to subjects for objects
- Access rights can be read, write, and execute
- Subject is an active entity that is seeking rights to a resource or object
- A subject can be a person, a program, or a process
- An object is a passive entity such as a file or a storage resource
- In some cases, an item can be a subject in one context and an object in another
- Columns of the access matrix are called access control lists (ACLs)
- Rows are called capability lists

The access matrix supports discretionary access control (DAC). The entries in the matrix are at the discretion of the individual(s) who have the authorization authority over the table.

In the access control matrix, a subject's capability can be defined by the triple (object, rights, random #). The triple defines the rights a subject has to an object along with a random number used to prevent a replay or spoofing of the triple's source.

Enterprise Architecture

- Zachman Framework
 - Created by John Zachman
 - Six perspectives on how to view an enterprise
 - How the entity operates
 - What the entity uses to operate
 - Where the entity operates
 - Who operates the entity
 - When entity operations occur
 - Why the entity operates

In order to design and secure an enterprise architecture, John Zachman created an architecture called the Zachman Framework. It covers six perspectives on how to view an enterprise:

- How the entity operates
- What the entity uses to operate
- Where the entity operates
- Who operates the entity
- When entity operations occur
- Why the entity operates

Enterprise Architecture (2)

- Federal Enterprise Architecture Framework (FEAF)
 - Five models
 - Business Reference Model
 - Performance Reference Model
 - Data and Information Reference Model
 - Service Component Reference Model
 - Technical Reference Model

The Federal Enterprise Architecture Framework (FEAF) also created an enterprise architecture based on five models

- Business Reference Model
- Performance Reference Model
- Data and Information Reference Model
- Service Component Reference Model
- Technical Reference Model

Trusted Computing Base (TCB)

- Security-relevant parts
- Access control mechanisms
- Reference monitor
- Kernel
- Protective mechanisms
- Monitors
 - Process activation
 - Process execution domain switching
 - Memory protection
 - I/O operations

The Trusted Computing Base (TCB) consists of the security-relevant parts of a system that include: access control mechanisms, Reference Monitor, the kernel, and protective mechanisms. Qualifying and quantifying the TCB is the domain of TCSEC, ITS EC, the Common Criteria, and ISO 17799.

There is an important piece of information to remember when dealing with the TCB. For terms of analysis, you assume that the components are properly implemented and secure. This is a big argument that people make in this domain. If you do not assume that the TCB is secure, then there is no way the system can be secure. If you do assume that the TCB is secure, and it is not, the assumptions you make will be incorrect.

Enforcement

Enforcement:

- Security kernel
 - The security kernel is the central part of a computer system (software and hardware) that implements the fundamental security procedures for controlling access to system resources. It is a most trusted portion of a system that enforces a fundamental property and on which other portions of the system depend.
- Reference monitor concept
 - The reference monitor concept is an access control concept that refers to an abstract machine that mediates all accesses to objects by subjects.
- Reference monitor
 - A reference monitor is a system component that enforces access controls on an object (files or programs). It is a design concept for an operating system to assure secrecy and integrity. The reference monitor should always be invoked; it is not capable of being bypassed and is capable of being evaluated.

From an enforcement perspective the following are key:

- **Security kernel**

The security kernel is the central part of a computer system (software and hardware) that implements the fundamental security procedures for controlling access to system resources. It is a most trusted portion of a system that enforces a fundamental property and on which other portions of the system depend.
- **Reference monitor concept**

The reference monitor concept is an access control concept that refers to an abstract machine that mediates all accesses to objects by subjects.
- **Reference monitor**

A reference monitor is a system component that enforces access controls on an object (files or programs). It is a design concept for an operating system to assure secrecy and integrity. The reference monitor should always be invoked; it is not capable of being bypassed and is capable of being evaluated.

Domain Separation

Domain separation:

- Protects objects in the system
- Domain: set of objects that a subject is able to access
- Domain separation may be implemented by:
 - Execution rings
 - Base address registers
 - Segmentation descriptors

Domains can be defined in a number of ways. For example, components of a domain are under the same management and operate under the same policy. The human resources department of an organization might be considered a domain. Or, domains might be divided into different levels of security or privileges.

Rings are one abstract way to separate security or privilege domains.

Recovery Procedures

- Fault-tolerant system:
 - Computer or network continues to function when components fail
 - System must be capable of detecting that a fault has occurred.
 - System must then have the ability to correct the fault or operate around it.

Again, the two important characteristics of a fault tolerant system are:

It has to be able to detect a fault.

It has to be able to correct the fault.

Recovery Procedures (2)

- Fail safe system
 - Program execution is terminated.
 - The system is protected from being compromised when a hardware or software failure occurs and is detected.
- Fail soft or resilient system
 - Selected, non-critical processing is terminated when a hardware or software failure occurs and is detected.
 - The computer or network then continues to function in a degraded mode
- Failover
 - Failover refers to switching to a duplicate "hot" backup component in real-time when a hardware or software failure occurs.
 - It enables the system to continue processing.
- System cold start
 - A system cold start is used when an unexpected TCB or media failure takes place or when recovery procedures cannot bring system to a consistent state.
 - TCB and user objects may remain in an inconsistent state following automatic recovery attempts.

In a fail safe system, the systems continues processing even if a fault occurs. Usually, systems can tolerate one fault. They may not, however, be able to handle multiple faults.

In a fail soft system, IT sheds non-critical processes and tries to keep the critical processing running.

In real-time systems, such as failover, a map of all the sensor inputs has to be kept in a redundant database in order to keep the system functioning properly.

Maintenance mode is used to bring system to consistent state.

TCSEC Classes

- Trusted Computer Security Evaluation Criteria
 - Orange Book
 - Part of the Rainbow series
- Cover operating systems
- Key principles
 - Functionality
 - Effectiveness
 - Assurance
 - Operational assurance
 - Covert channel analysis
 - Trusted facility management
 - Trusted recovery
 - Life cycle assurance

The U.S. Department of Defense developed the Trusted Computer Systems Evaluation Criteria (TCSEC), also known as "The Orange Book." Operating systems, applications, and computer-related products are classified into one of four categories to describe their functionality, effectiveness, and assurance. These categories loosely fit the needs of various levels of security within the United States Department of Defense. They are:

- A: Verified protected
- B: Mandatory protected
- C: Discretionary protected
- D: Minimal security

Orange Book

- Four classes
 - A: Verified protected
 - B: Mandatory protected
 - C: Discretionary protected
 - D: Minimal security
- Based on:
 - Security policy
 - Object marking
 - Subject identification
 - Accountability
 - Assurance
 - Documentation
 - Continuous protection

The Orange book has 4 classes:

- A: Verified protected
- B: Mandatory protected
- C: Discretionary protected
- D: Minimal security

Based on:

- Security policy
- Object marking
- Subject identification
- Accountability
- Assurance
- Documentation
- Continuous protection

D and C Classes

- D
 - Base level
 - Every system by default is D
- C1: Discretionary protection
 - Separation of users and data
 - Cooperating users processing data at same level
 - Test documentation
- C2: Discretionary protection
 - More finely-grained than C1
 - Auditing
 - Resource isolation
 - Addresses object reuse

The following are the descriptions for D and C classes:

- D
 - Base level
 - Every system by default is D
- C1 Discretionary protection
 - Separation of users and data
 - Cooperating users processing data at same level
 - Test documentation
- C2 Discretionary protection
 - More finely-grained than C1
 - Auditing
 - Resource isolation
 - Addresses object reuse

B Class

- B1: Mandatory protection
 - Sensitivity labels enforce mandatory access controls
 - Informal statement of security policy model
- B2: Mandatory protection
 - Formal security policy
 - Strengthened authentication
 - System admin and operator functions
 - Configuration management
 - Covert channel analysis
- B3: Mandatory protection
 - Satisfy all reference monitor requirements
 - Security domains
 - Security administrator
 - Recovery procedures
 - Greater granularity than B2
 - Audit trail

The following are the descriptions for the three levels of the B class:

- B1: Mandatory protection
 - Sensitivity labels enforce mandatory access controls
 - Informal statement of security policy model
- B2: Mandatory protection
 - Formal security policy
 - Strengthened authentication
 - System admin and operator functions
 - Configuration management
 - Covert channel analysis
- B3: Mandatory protection
 - Satisfy all reference monitor requirements
 - Security domains
 - Security administrator
 - Recovery procedures
 - Greater granularity than B2
 - Audit trail

A Class

AI: Verified design:

- Analysis derived from formal design specifications and verification techniques
- Formal model of security policy

AI is very difficult too achieve and there are not many in this class. The following is the description for AI:

- AI: Verified design
 - Analysis is derived from formal design specifications and verification techniques.
 - It is a formal model of security policy.

Trusted Facility Management

- Typical system administrator or enhanced operator functions:
 - Installing system software
 - Starting up a system
 - Shutting down a system
 - Adding and removing users
 - Performing backups and recovery
 - Handling printers and managing print queues
- Typical security administrator functions:
 - Setting user clearances, initial passwords, and other security characteristics for new users
 - Changing security profiles for existing users
 - Setting or changing file sensitivity labels
 - Setting the security characteristics of devices and communications channels
 - Reviewing audit data

User account setup and maintenance should not be performed by same individual who initiates and authorizes the creation of the account.

- Typical system administrator or enhanced operator functions include:
 - Installing system software
 - Starting up a system
 - Shutting down a system
 - Adding and removing users
 - Performing backups and recovery
 - Handling printers and managing print queues
- Typical security administrator functions include:
 - Setting user clearances, initial passwords, and other security characteristics for new users
 - Changing security profiles for existing users
 - Setting or changing file sensitivity labels
 - Setting the security characteristics of devices and communications channels
 - Reviewing audit data

Trusted Recovery

- The objective is to ensure that the system remains secure in the event of system failure.
- The system must be able to be restarted without being vulnerable to attack.
- You can execute recovery without being compromised after the failure.
- Trusted recovery is required only for B3 and A1 level systems.

When a system crashes or is undergoing recovery, security mechanisms might be disabled, making the system vulnerable to attack.

Preparing for failure:

- Perform backups on scheduled basis.
- Provide for recovery in a safe manner.

System Recovery

- Bring the system up in a single user mode.
- Recover all file systems.
- Restore damaged or lost files.
- Restore file security designations.
- Make sure critical security files are in tact, such as password files.

Manual recovery requires system administrator involvement. This type of intervention is needed to resolve multiple failures.

Automated recovery is used for a single failure event and recovers systems to a secure state (without system administrator involvement). Manual intervention is required to resolve multiple failures.

Automated recovery without undue loss is similar to automated recovery.
A higher level of recovery prevents the undue loss of protected objects.

TNI

- Trusted Network Interpretation-
the Red Book
- Key features
 - Network protection**
 - CIA**
 - Labels**
 - Service protection**

The Trusted Network Interpretation is part of the Rainbow series. It is covered in the "Red Book.

The key features of TNI are:

- Network protection
- CIA
- Labels
- Service protection

ITSEC Classes

- European
 - First common standard
 - Main attributes
 - **Functionality (F)**
 - **Assurance (E)**
 - Target evaluation
- F1 + E1
 - F2 + E2
 - F3 + E3
 - F4 + E4
 - F5 + E5
 - F5 + E6

The ITSEC was the first attempt by European countries to establish a common standard for evaluation of computer security. Its main goal was to delineate between functionality and assurance. The Orange Book was thought to be too rigid and focused on assurance at the expense of functionality. The ITSEC segregated these main components, allowing for systems to have a more accurate definition.

ITSEC: Functionality

- F1-F5
 - Mirror functionality of orange book
- F6
 - High integrity requirements
 - Databases
- F7
 - High availability
- F8
 - High integrity for communication
- F9
 - High confidentiality
- F10
 - High confidentiality and integrity for data networks

The following are the ITSEC functionality levels:

- | | |
|-------|---|
| F1-F5 | Mirror functionality of orange book |
| F6 | High integrity requirements <ul style="list-style-type: none">• Databases |
| F7 | High availability |
| F8 | High integrity for communication |
| F9 | High confidentiality |
| F10 | High confidentiality and integrity for data networks |

ITSEC: Assurance

- EO
 - Inadequate assurance
- E1
 - General description
- E2
 - Configuration and process control
- E3
 - Source code analysis
- E4
 - Formal model of security policy
- E5
 - Vulnerability analysis
- E6
 - Formal specifications

The following are the ITSEC Assurance levels:

EO	Inadequate assurance
E1	General description
E2	Configuration and process control
E3	Source code analysis
E4	Formal model of security policy
E5	Vulnerability analysis
E6	Formal specifications

Common Criteria

- ISO
- International 2nd attempt
- Evaluation Assurance Level (EAL)
- The Evaluation Assurance Level (EAL) is applied to a product rather than a system. The rating system is as follows:
 - EAL 1: Functionally tested
 - EAL 2: Structurally tested
 - EAL 3: Methodically tested and checked
 - EAL 4: Methodically designed, tested, and checked
 - EAL 5: Semi-formally designed and tested
 - EAL 6: Semi-formally verified, designed, and tested
 - EAL 7: Formally verified, designed, and tested

ITSEC was thought to be a failure due to its confusing mix-and-match approach when applied to live business situations. From it was born the Common Criteria. The support for this document expanded to include Canadian and U.S. input. This document in all its parts is a monster work maintained by the International Organization for Standardization.

The Evaluation Assurance Level (EAL) is applied to a product rather than a system. The rating system is as follows:

- EAL 1: Functionally tested
- EAL 2: Structurally tested
- EAL 3: Methodically tested and checked
- EAL 4: Methodically designed, tested, and checked
- EAL 5: Semi-formally designed and tested
- EAL 6: Semi-formally verified, designed, and tested
- EAL 7: Formally verified, designed, and tested

ISO 17799

- 10 sections
 - Security policy
 - Security organization
 - Assets classification and control
 - Personnel security
 - Physical and environmental security
 - Computer and network management
 - Systems access control
 - Systems development and maintenance
 - Business continuity planning
 - Compliance
- Risk-based
- Holistic approach

Originally, the ISO 17799 was a British Standard (BS 7799) designed to offer another alternative to the Common Criteria. It is a risk-based approach for evaluation and assessment that is comprised of ten sections:

1. Security policy
2. Security organization
3. Assets classification and control
4. Personnel security
5. Physical and environmental security
6. Computer and network management
7. Systems access control
8. Systems development and maintenance
9. Business continuity planning
10. Compliance

The last part on compliance has yet to be ratified.

Summary

- Hardware represents a critical security risk.
- You must understand every component that processes data.
- You need to know who is responsible for validating and signing off on the security of a system.

This section discussed the Security Architecture and Design domain. This domain looks at the hardware components that make up a system. This includes the various ways programs run in memory and how computers typically use memory to process an application. It also addressed some of the critical risks that hardware vulnerabilities represent. Although this domain is not on the radar of what most users need to understand, it is critical for system designers to know. It is important that you strictly enforce need to know principles across all data.

This page intentionally left blank.

7. Operations Security

10 Domains of Knowledge

This section covers Domain 7, the Operations Security domain.

Domain 7 Overview

- Legal
- Administrative
- Operations controls
- Monitoring
- Auditing
- Reporting
- Roles and responsibilities

Operations security is all of those things that interact with technical security components to make them function properly. Some people refer to operations security as the softer side of security because it is not as technical as other areas. Although it is not technical, operations security is still important and is complicated in its own right. One key area of operations security involves the legal implications associated with the actions a company or individual performs. If you study the law, you might note that it is anything but straight-forward or easy. Therefore, it is important to interface with the correct parties who can provide the proper level of expertise.

One of the key things to keep in mind with operations security is to know your job. You might need to know, at a high level, that there is a legal implication to what you do, but you are not supposed to have knowledge of the details. You should know just enough so that you can ask the correct people for expert advice.

Operations Security

- The act of understanding threats and vulnerabilities in order to routinely support operational activities that enable computer systems to function correctly
- Addresses
 - Threats in an operating environment
 - External attackers
 - Internal malicious intruders
 - Operators inappropriately accessing resources

Operations security also involves security mechanisms for processing transactions, support operations, and system administration functions.

It also involves ensuring that auditing and monitoring are in place and functioning properly.

Controls should also protect relevant resources, including hardware, magnetic media, access tokens, software, backup files, and audit trails.

Operations Security

Operations security:

- Threat
 - An event that could cause harm through violation of operations security
- Vulnerability
 - A weakness in a system that could be exploited to cause harm to a system
- Asset
 - Computer resources, hardware, software, information, personnel, and so on

As in all other domains, the intent of operations security is to protect the CIA of information in an operations environment.

Threats in operations could involve system administrators with malicious intent, compromise of backup files, unauthorized monitoring of printout, etc.

Management Legal Requirements

- Legal and regulatory obligations
- Retention of records including e-mail
 - How long transactions and other records should be retained
 - Management issues
 - Legal issues
 - Audit issues
 - Tax compliance issues
- Privacy issues
- Due care versus due diligence

As senior managers, we can use operations security as a tool we to ensure the continuous maintenance of the organization and the environment in which the organization operates. Therefore, the legal and regulatory side of security must also be considered.

There are various legal requirements from your country and state that give specific guidelines to follow. Because requirements vary widely, we use examples from within the United States. Management must identify the laws their company or organization falls under so that they are aware of their legal responsibilities, particularly what their obligation is toward their customers, shareholders, employees, and the community overall.

In most industrialized nations, the law requires consumers and businesses to respect licenses and copyrights. In accordance with these laws, management must maintain accountability by utilizing proper controls to ensure that actual software usage corresponds with software license agreements. Although many companies now build software copyright protections into their licenses, unethical people will continue to attempt to overcome these protections and will often be successful in their attempts. It is our job to prevent this activity within our organizations. This will keep us out of legal trouble and allow us to develop a good relationship with our software vendors.

Record retention is simply a process for backing up and saving information in accordance with laws or corporate procedure. In our personal lives, we would never throw out a will or a birth certificate; likewise, many organizations have requirements to save all documents that are created or pass through their organization.

Media Security

- Controlling access to media
- Proper disposal of media
- Sanitizing media
 - Removing data
 - Overwriting - overwriting media multiple times with specific patterns of ones and zeroes
 - Degaussing - applying a large magnetic field to erase magnetic media
 - Destruction

Complete destruction is considered the safest way of sanitizing media. Overwriting—strict configuration controls must be in place both on the operating system and the software itself.

To purge the media, the DOD requires overwriting with a pattern, then its complement, and finally with another pattern; e.g., overwrite first with 0011 0101, followed by 1100 1010, then 1001 0111. To satisfy the DOD clearing requirement, you are required to write a character to all data locations in the disk. The number of times an overwrite must be accomplished depends on the storage media, sometimes on its sensitivity, and sometimes on differing DOD component requirements, but seven times is most frequently recommended.

Degaussing is often recommended as the best method for purging most magnetic media. Degaussing is a process whereby the magnetic media is erased, i.e., Returned to its initial virgin state. Erasure via degaussing may be accomplished in two ways: In AC erasure, the media is degaussed by applying an alternating field that is reduced in amplitude over time from an initial high value (i.e., AC-powered). In DC erasure, the media is saturated by applying a unidirectional field (i.e., DC-powered or by employing a permanent magnet). Another point about degaussing: degaussed magnetic hard drives will generally require restoration of factory-installed timing tracks, so data purging is recommended.

Data remnants refers to the data left on the media after the media has been erased. After erasure, there might be some physical traces left, which could enable the sensitive data to be reconstructed.

Documentation

- Documentation has to be controlled in an organization.
 - Have to know where documentation is
 - Accountability
 - Protection from unauthorized disclosure
- Documentation examples:
 - Security plans
 - Contingency plans
 - Risk analyses
 - Security policies and procedures

A security system needs documentation controls. Documentation can include several things: security plans, contingency plans, risk analyses, and security policies and procedures. Most of this documentation must be protected from unauthorized disclosure; for example, printer output must be sent to a secure location. Disaster Recovery documentation must also be readily available in the event of a disaster.

OPSEC Vulnerabilities Assessment

1. Identify critical information
2. Assess the threat
3. Assess vulnerabilities of critical information to the threat
4. Conduct risk versus benefit analysis
5. Implement appropriate countermeasures
6. Repeat

Five considerations allow you to apply operations security to an organization, a system, or a process. They are listed in the following:

- What is the information or resources that other people might consider important? What is the value of this information? Is the information important to you? Would the information be important to someone else?
- If someone else had access to this information, could this be a threat to your company or country? What are the threats to which the organization might be exposed?
- If you look carefully at what you have in place, what is the likelihood that some of these threats could become reality? What are some tricks and methods an adversary could use to get access to this information or resource or to modify it? What are the capabilities of a potential adversary? Could he possibly access the information? (If you are connected to a network that is connected to the Internet, this is a distinct possibility.) If your information is locked in an access-controlled vault and the adversary has to send in a Mission: Impossible team, this is less likely. How likely is it that someone can get at this information? How much work is involved to complete tricks?
- After you have identified potential threats and their impact on your organization, you must perform a cost/benefit analysis to identify what countermeasures are worth implementing and which ones make sense as far as business is concerned.
- Finally, you must fix problems and put in countermeasures. Make sure your employees and co-workers are aware of what you are doing to improve their job security. Weak or poorly implemented security can lead to the loss of a company's market share or undermine a country's very existence.

The following site lists additional information on OPSEC basics: http://www.nswc.navy.mil/ISSEC/Docs/Ref/GeneralInfo/opsec_basics.html.

Key Areas of OPSEC

- Resource protection
 - Protecting systems that contain critical data
- Privileged-entry control
 - Control and limit access
- Hardware control
 - Should allow for trusted recovery

The following are key areas of OPSEC:

- **Resource protection**
Protecting systems that contain critical data
- **Privileged-entry control**
Control and limit access
- **Hardware control**
Should allow for trusted recovery

Privacy and Protection

- Code of law
- Use and maintenance
- Protection of data
- Monitoring

Privacy and Protection

Privacy is one of the hottest topics today. Whether it is in a virtual world or in our daily life, our privacy is constantly challenged by new technology deployed to track what we are doing. There are video cameras at airports and in public places, photo radar systems, cameras to watch for traffic conditions, and the list goes on and on.

In recent years, the U.S. has seen the enactment of two privacy laws, the Gramm Leach Bliley Act (GLB) and, as mentioned before, the Health Insurance Portability and Accountability Act (HIPAA). The HIPAA law states that personal financial data and patient medical data must be protected from disclosure by the organization responsible for storing and using that data. Such laws are not only enacted in the U.S.; they are currently being introduced in many other countries as well.

The HIPAA privacy laws state that any personal medical information collected by a doctor, hospital, or insurance provider on a patient must be accurately maintained, used only for the purpose for which the information was collected in the first place, and not be transferred to another entity without the consent of the patient. Furthermore, the patient has a right to correct the information. When the information is to be transferred to another entity with patient approval, it must be ensured that the receiving entity can provide the same level of protection; if they can't provide that assurance, then the information cannot be transmitted.

As previously stated, the laws vary among jurisdictions, so ensure that you consult with local authorities to become familiar with what you are authorized to do and what protections you must put in place to protect such sources as your employees' data and your customers' data.

Illegal Activities

- **Fraud**
- **Theft**
- **Collusion**

Illegal Activities

A large part of operations security is the process of detecting any acts of fraud or theft that take place in the organization. Fraud is sometimes committed due to excessive rights being given to specific individuals in the organization. Sometimes it is also based on collusion that takes place between individuals who do not have enough privileges, by themselves, to commit the act. However, by acting as a group, they can access all of the functions needed to commit fraud.

For example, think of a system administrator who has access to most of the major servers in the company, including the HR servers, and who can create and delete accounts; he does not have access to the HR payroll servers. Throughout a one-year period he befriends the HR payroll database administrator and convinces him that together they can create a fictitious employee and collect his paychecks. This is collusion. By themselves they could not do this, but together they can.

Insiders or internal employees commit most fraud and theft cases. There are quite a few factors that can contribute to fraud taking place; some of the most common are pressure, opportunity, and rationalization. Audit trails and proper control are the most efficient deterrents against fraud. You have to ensure that fraud is not tolerated and that it is prosecuted when it does happen. In the upcoming pages, you will see some effective means of fighting fraud and collusion within an enterprise.

Administrative Management

- Job requirements
- Background checking
- Separation of duties
- Job rotation
- Vacation and leave
- Terminations

Administrative Management

As discussed in the introduction of this section, maintaining the security of your environment includes more than maintaining software and hardware. Proper security also includes the selection and screening of the personnel who work within your facilities. If a person has physical access to your servers, then there is nothing to stop her from abusing those privileges. Proper security screening and training of your personnel has to start before they are allowed into the work environment.

The first step is to ensure that your human resource department does a thorough verification of an applicant's past employment before hiring that person. Call previous employers and seek information from the human resource department. It is nice to have references on resumes; however, they are frequently the names of people who will talk highly of the person regardless of their history with a previous employer.

As seen in the previous section, fraud or theft is often the result of opportunities. To avoid presenting such opportunities, you must carefully select what privileges or access an individual is assigned. Take the time required to identify jobs that give too much access to one employee and look for ways to reallocate tasks across multiple positions. This can help cut down on opportunities to commit fraud by an individual, but it does not alleviate the chance of collusion between two or more employees. One of the ways to cut down on fraud is to regularly rotate positions in the IT department. This will allow employees to work with new counterparts and give them new skills (cross-training), while improving security by eliminating opportunities to commit fraud.

Many employees work hard all year and plan for a well-earned vacation; however, you might have some employees who do not wish to take an annual vacation. Although these employees might be some of your most dedicated staff, it could also indicate a hidden activity that would quickly come to light if another employee took over while the regular staffer was on vacation. This employee might be abusing resources, such as selling information or gathering intelligence for a competitor. It might not be in the best interest of your company to create a "big brother is watching" environment within your work place, but awareness of what can happen is important.

Employment Agreements

- General clauses
- Work hours and overtime
- Holidays, sick leave, and other leave
- Non-competition and non-solicitation
- Confidentiality
- Non-disclosure agreement (NDA)

Employment Agreements

Considering the extremely competitive market in which IT specialists perform, the lifelong commitment to a single company is less and less common. Most IT employees change jobs regularly as they seek new challenges and opportunities. For these reasons, it is important to have strong binding agreements within employee contracts. You do not want an employee to leave with your customer list, your Research and Development (R&D) information, and other valuable data that might give a competitor an unfair advantage.

Following are some important points related to operational security that can be included in a typical contract. These points are to protect the employer as well as the employee. Good negotiation and clear understanding of these issues from the beginning are key to a harmonious and long-term relationship.

General Information

- Position title
- Main duties and responsibilities
- Immediate supervisor

Special Provisions

- Employee's declarations and warranties
- Probation
- Exclusivity
- Care and diligence
- Loyalty and confidentiality
- Use and possession of the employer's property
- Intellectual property
- Medical examination
- Respect of company policies
- Undertaking of confidentiality
- Undertaking not to compete
- Notice of termination or resignation
- Reasons for dismissal
- Return of the employer's property

Individual Accountability

- Own responsibilities
- Act as a deterrent through audits
- Proper mechanisms must be in place
- Must be within the law

Individual Accountability

A company's auditing capabilities ensure that individuals are responsible for their own actions and that they comply with regulations, policies, and guidelines the company puts in place. Having an environment in which actions are tracked and logged offers an easy means of holding specific users accountable. It acts as a valid deterrent. If people know they are monitored, they will think twice before attempting an illegal action. Accountability can be applied toward internal users, but also against external users who may attempt to abuse the system and its network resources.

All monitoring and auditing should be done in accordance with local laws to avoid further prosecution or accusation of invasion of privacy on the employee's part. Usually employees can be subject to monitoring for security, performance, or other reasons, as long as they are informed of the mechanisms that are in place before the monitoring begins. It is a great idea to keep the policy in the HR hiring forms and to require new employees to read and sign the policy when they begin employment. From that point on, all you need to do is add it into the security awareness plan.

"Need to Know"

- More granular than least privilege
- Only when necessary
- Only to what is necessary
- Only where it is necessary
- A business requirement
- Usually combined with the "least privilege" principle

"Need to Know"

The "need to know" concept ensures that only people who have a need to access certain information or resources are authorized to do so. This access can be further restricted to specific hours, days, or a timeframe. Access is granted based on a business requirement and not simply because someone has a desire to see specific information.

Often we also talk about the least privilege principle, which is similar to need to know. The least privilege principle ensures that only the minimum required access is given at any time. The difference between the two is subtle; least privilege can mean someone gets only user rights and not administrator rights to their workstation, whereas with need to know, they might have access to all of the development data, but none of the HR data.

Sensitive Information

- Marking
- Handling
- Storage
- Destruction

Sensitive Information

All sensitive information should be clearly marked in accordance with the organization's security policy and security guidelines. Some companies mark information as proprietary or sensitive, but a far better course is to explicitly mark sensitive information by level. This should be done for all information, regardless of the media on which it is stored. Marking is used in systems that must enforce a mandatory security policy; information is marked with a classification or another sensitivity label. The two most common categories are management only for human resources information and proprietary for trade secrets. The system must ensure that these classifications or sensitivity labels are maintained while the information flows through the systems.

Electronic labels and paper labels should be used the same way. **Paper**, media, and any other items that contain sensitive information should be marked with the appropriate sensitivity label, and the label should remain clearly visible. This will ensure that you do not erase or destroy important information by error and also that you do not transfer classified information to unauthorized persons.

Control Types

- Directive controls
- Preventive controls
- Detective controls
- Deterrent controls
- Corrective controls
- Recovery controls

Control Classification

Gene Kim, the CTO of Tripwire, did a study of hundreds of organizations in late 2002 and early 2003. He found that many organizations were struggling with patch management and with system administrator-to-server ratios of one administrator to five or six servers. Other organizations were humming along with ratios that had one administrator to a hundred servers. These organizations also had strong security. The difference between the strong organizations and the ones that were simply struggling to survive was controls. There are different ways of classifying controls. Some organizations like to classify them by their nature, such as administrative controls (policies and HR actions), technical controls (IDS and single sign-on), and physical controls (mantraps and fences). Another way of classifying controls is based on actions, such as directive, preventive, detective, corrective, and recovery.

Directive controls are the equivalent of administrative controls. This category includes items such as policies, standards, guidelines, personnel screening, and security awareness training. Directive controls are important and form the foundation for enterprise security.

Preventive controls are the equivalent of technical controls. These contain the methods, tools, practices, and the techniques used to ensure that systems remain secure and highly available. They also include logical access control, encryption, security devices, identification, authentication, firewalls, antivirus, separation of duties, access rights, data classifications, physical access controls, and many more.

Detective controls are used to validate that the preventive controls and the directive controls perform adequately. This is how computer abuse, fraud, or crime is detected with both automated and manual tools. This type of control area contains log review, surveillance, auditing, and integrity checkers, to name a few.

Corrective controls provide information, procedures, and instructions for correcting detected shortcomings. These shortcomings can include attacks that have been detected, errors, or system misuse. In this category are procedures, instruction manuals, audit trails, and many more.

Operation Controls

- Resource protection
- Privileged-entity controls
- Hardware controls
- Input/output controls
- Media controls
- Administrative controls

Operation Controls

Operation control means controlling the access to a computer facility, the entire site housing the computer facility, and the movement of data within networks. You implement operation controls to protect resources that the organization has designated as sensitive or business-critical. These resources must be protected to keep an organization operating. Auditing is one way to determine if resources have been tampered with.

Operation controls deal with protecting anything that processes or touches critical data. Some of these key controls are: hardware, input/output, media, and administrative.

Controls and Protections

Other types of operational controls:

- Transaction controls
 - A processing control verifies that a transaction is processed correctly.
 - An output control ensures the confidentiality of output information and verifying the integrity of the output data.
 - Test controls are implemented during tests to prevent unauthorized access to sensitive data.
 - Change controls are used to protect the integrity of data when changes to the system configuration are made.

Processing controls also ensure that incorrect entries are reprocessed accurately.

Output controls can include restricted access to printout locations, obtaining receipts and ID when releasing critical information, and appropriate banners at the beginning and end of printouts.

Test controls are designed to preserve the confidentiality and integrity of data.

Monitoring

Steps in monitoring

- Review
- Watch
- Take action

Types of monitoring

- Real-time
- Ad hoc
- Passive

Auditing

- Monitoring and looking for change

Monitoring and Auditing

Monitoring is an all-inclusive term that can mean many things. In this context, it means to review, to watch, and to audit the network. The network is monitored to identify access attempts to the data or resources, and although this is really one of the operation controls previously discussed, it is important enough for us to go into more depth and look at the types of monitoring available today. It is also important to look at what can be done with the data collected by the monitoring devices and how configuration management can help network administrators identify when someone else has placed an unauthorized monitor on the network.

Auditing is closely related to monitoring; some smaller organizations that monitor network security devices also audit the network and security device data. The relationship between the two is simple. A team monitors the network for suspicious activity; if activity is detected and a possible compromise of a system is suspected, then an audit team is asked to review the system in question, physically and electronically. This would include the system logs and critical files for any anomalies.

Monitoring tools are used to ensure that security devices, resources, and usage are in accordance with the policies that are put in place. Monitoring is essential to ensure availability, security, and proper care of resources. There are different types and categories of monitoring. Three popular ones in widespread use are real-time, ad hoc, and passive monitoring. It is not accurate to call monitoring a real-time technique because activities cannot be reported until they actually happen or are identified. Monitoring is not truly real-time, but is as immediate as possible after the fact. Intrusion detection systems (IDS) are sometimes called real-time. This is because they have the capability to react to specific events identified on the network. This is based on exploit signatures that an IDS monitors for or a threshold that is crossed, such as a certain amount and type of traffic across the network.

The ad hoc monitoring technique is performed at regular intervals or whenever a need arises. This category contains tools such as vulnerability checkers, file integrity checkers, network sniffers, and log consolidation tools.

Some monitoring tools are passive on the network; they have no active role and cannot interfere with passing traffic, unlike a firewall that can be configured to stop traffic. Some IDS devices can be set up to "reset" traffic connections. This means that if a system connected to the network exhibits a sign of suspicious activity, the connection can be terminated with the IDS device. Many devices are not set up to automatically reset connections due to false positives or activity that looks suspicious but is actually normal traffic. It is not a good idea to automatically terminate what might be an important link. This is why trained analysts review the alarms and data.

Types of Monitoring

- Keystroke
 - Hardware vs. software monitoring
- Illegal software
- Traffic analysis
- Trend analysis

Types of Monitors

Keystroke monitoring is a technique that records all keys (in some cases, it records all mouse clicks and menu selections) while a user is at a computer. It is generally used with legal permission in a criminal investigation to monitor the activity of suspected criminals. In most countries, it is used only with the permission of a judge. This is an effective means of collecting information from a user. The monitoring tools used to record the keystrokes can be either software- or hardware-based.

Software tools are well-made and usually have a blind mode that will not show in a list of processes. You will not see them on an application bar or system tray, and in most cases, there is no simple way of knowing that they are active on a machine. These tools are sometimes distributed by a Trojan that entices you to look at a new software piece (a game or utility).

The hardware tools used to record keystrokes are in the form of a keyboard adaptor. In most cases, you install this device at the back of the computer and then connect your keyboard to it. The device will record and store keys pressed at the keyboard and later on, you simply pick the device up and have access to all passwords that were entered on the keyboard while the device was in place. This is a technique that was previously used in some criminal cases in which cryptography was used and the secret passphrase was needed by the authorities to break the encrypted information.

One of the more interesting things you can do with the data collected from your security monitors is perform an analysis of the traffic traversing your network. The fundamental basis of traffic analysis is in the detection of a message passed from A to B.

Monitoring Methods

Violation tracking, processing
and analysis:

- Establish clipping levels.
 - Baseline of user activity
 - Activity considered "normal" if it does not exceed clipping level
- When the clipping level is exceeded a violation record is produced.

For example, clipping level could be set at three for failed login attempts. Any login attempts below three would be considered operator error and not flagged as a possible attack.

Configuration Management

- Process of tracking and approving changes to a system
- Incorporates change control
- Involves identifying, controlling, and auditing all changes made to a system
- Primary security goal is to ensure that changes to the system do not unintentionally diminish security

You have a responsibility to make sure that any changes, additions, etc. to system configuration do not jeopardize the ability of the system to meet the requirements to be considered a trusted system. Configuration management is a high-level, formalized approach to managing changes to complex systems.

Configuration management is required for formal, trusted, systems because it ensures that changes occur in an identifiable and controlled environment.

Configuration Management (2)

- Configuration identification
 - Decompose system security verification components into identifiable, trackable, and manageable units known as **configuration items**.
- Configuration control
 - System changes are approved before implementation.

Configuration management is used to identify and document the physical and functional characteristics of information systems.

This includes managing, reporting, and recording change processing and implementation. The two key areas are configuration identification and configuration control. Configuration identification decomposes system security verification components into identifiable, trackable, and manageable units known as **configuration items**. Configuration control is where system changes are approved before implementation.

Configuration Management (3)

- Configuration accounting
 - Provides means for managers to trace changes to the system and establish a history of problems and issues
- Configuration audit
 - Verifies that policies are being followed
 - Determines state of completeness and consistency of related accounting information

Configuration accounting documents the status of configuration control actions, keeping track of changes as they progress through the process.

It determines the granularity of recorded information.

The accounting function should be able to identify and locate all versions of a configuration item and its corresponding changes.

Audit is the quality control portion of change control.

Change Control

- Primary functions of change control
 - Ensure that the change occurs in an orderly manner
 - Ensure that the user base is informed of the pending change
 - Analyze the effect of the change on the system
 - Reduce the negative impact the change may have had on the computing resources

Change control is the process of tracking and approving of changes to a system, including identifying, controlling, and auditing all system changes.

That process includes hardware, software, and networks.

Change control is also concerned with changes that might affect security.

It ensures that changes are reflected in current documentation.

Change Control (2)

- Generally accepted procedures to implement change control process
 - Applying to introduce a change
 - Cataloging the intended change
 - Scheduling the change
 - Implementing the change
 - Reporting the change to the appropriate parties

An application for a change is presented to the entity responsible for approving and administering changes.

A change approval involves a trade-off analysis of the change and the corresponding justifications.

Changes should be documented and updated and the change recorded in a change control log.

The change should be formally tested.

A full report must be submitted to management with a summary of the change.

Intrusion Detection

- Intrusion prevention = Before
- Intrusion detection = During
- Intrusion response = After

Intrusion prevention is about putting the controls in place to ensure intrusions are stopped before they start. It is not IPS or intrusion prevention systems. Intrusion prevention is correctly identifying and authenticating users. Are your systems in place using 2, 3, or 4 factor authentication? Is physical access to terminals controlled?

Intrusion detection is about knowing that an intruder is or has been in the system(s). Intrusion detection can take place in the following ways: actively, through trace evidence, by blatant communication, or by system anomalies. For example, you might discover the intrusion during the breach, such as catching the burglar in the act. That black beret and mask over the eyes is typically a dead giveaway. You might see the remnants of the intrusion, such as going to turn on the television, but it is not there. You might get a message from the intruder, such as a post card from Barbados that has a message thanking you for the cash. Or, you might notice strange activities after the intruder has left, such as the door never closing correctly again.

Intrusion response is the next step. Ensure you have clear evidence of an intrusion. What actions can be taken to restore the systems to working order? What steps can be taken to ensure that it doesn't happen again? Computer Security Incident Response Team (CSIRT) can help define the plan of action for the organization.

Types of Intrusion Detection

- Integrity checking
- Anomaly identification
- Attack signature identification

Both network-based intrusion detection systems (NIDS) and host-based intrusion detection systems (HIDS) focus on one or more types of intrusion detection.

Integrity checking products verify the current system by hashing values of all the files, yielding a checksum, which gives an exact snapshot of what files should be on the systems and what values those files have at an instant in time. As long as the initial system is trustworthy, changes are reported.

Anomaly identification measures the baseline of activity over time and highlights exceptions. The weakness is that users and network activity do change over time for many reasons, causing many false positives.

Attack signature identification (signatures) is a database of known attacks (events) programmed into a NIDS or HIDS. Each event has a unique signature that causes an alarm. The weakness with attack signature identification is that slight variations cannot be programmed into the database.

Auditing

- Compliance checks
- Internal and external
- Frequency of review
- Standard of due care

Auditing

Auditing is a function that will verify the security of systems and resources and whether or not a system has been compromised or misused. Auditing also tests the effectiveness of the operation controls implemented throughout the network. It can also help determine where more controls might be needed. This is an important step in the accountability process. If you don't audit your systems, then it is extremely difficult to make your users responsible for their actions. Auditing is a broad topic. In this instance, it covers the review of data gathered from security devices (log reviews) and the security assessment of devices connected to the network with an assessment tool. There are many types of assessment tools — all of them gather similar data. They generally gather information on a specified list (built into the tool) of vulnerabilities or weaknesses in the operating system or application. This gives a list of what must be fixed to keep someone from compromising the vulnerability. The tools are run either automatically at pre-determined times or manually by the auditing team members.

There are normally two types of audits: internal audits and external audits. Internal employees perform internal audits and external audits are performed by a third party or outside trusted firm. It is important to regularly use both audit types, even though this can be costly. It is not a good idea to use your network administrator to audit your network. This creates a situation in which a person is both the accused and the judge. It's likely the administrator would find himself innocent. Audits are normally conducted through the use of a detailed checklist. This allows for consistency. Moreover, results of audits can be compared.

The frequency at which audits take place depends on multiple factors, such as the complexity of your environment, regulations, and policies. In some cases, an audit might occur after major changes have been implemented or an incident has taken place. Although audits are usually conducted at regular intervals, it is also a good idea to conduct surprise audits occasionally. This ensures that security is maintained throughout the year and not only when there is an upcoming audit announced.

The whole audit process is a verification to ensure that due care is carried out in accordance with best practices of the industry. This is what is sometimes referred to as the Prudent Man rule.

Audit Trails

- Must be reviewed
 - Periodic manual review to make sure tools are working
- Must be part of a routine
- Ease task with use of tools
- Ensure tool works properly
- Records the history of transactions on the system
- Can be used to flag indications of abnormal behavior by attackers
- Provides accountability through the ability to reconstruct past events and identify users associated with those events

Audit Trails

It is fine to collect audit information, but it is useless unless you review this information regularly. Audit trails must be reviewed on a schedule set by policy. This is a very important step in protecting your environment. Often logs will give the first indication that something suspicious is going on with your systems and that abuse might be taking place. From a legal standpoint, it is important to be able to demonstrate that audit trails are conducted on a regular basis. This may be the only way that your evidence might be admissible in court.

It is also understood that manual review of audit logs can be cumbersome and quickly become a full-time job. It is strongly recommended that you use a log-reduction tool to avoid looking at hundreds of megabytes of information. Such a tool could be used in conjunction with an anomaly detection tool that would notice unusual trends in traffic patterns. A good example is CodeRed. It is abnormal for your Web server to start browsing the Web on port 80 outbound. This type of traffic would be flagged by an anomaly detector tool or by a firewall that is properly configured to only allow the authorized traffic outbound.

Just a word of warning: ensure that the tools that you use are working properly and that they do not give you a false sense of security. This is why it is important to manually process part of the logs to ensure that your tools are working.

Audit Trails (2)

- Transaction information logged
 - Individual conducting transaction
 - Date
 - Time
 - Location (workstation) used to process the transaction
- Audit information should be protected at the highest level of security in the system.
- Audit information should be retained and protected when stored off-site.
- The integrity of the audit information must be protected.
- Audit data has to be available even during a security breach.

The following is transactional information that can be logged:

- Individual conducting transaction
- Date
- Time
- Location (workstation) used to process the transaction

Audit trails show:

- Transaction's date and time
- Who processed the transaction
- At which terminal the transaction was processed
- Various security events relating to the transaction
- Production job reruns
- Computer operator practices
- All commands directly initiated by the user
- All identification and authentication attempts
- Files and resources accessed

Audit Log Backup

- No log, no audit.
- Central logging
 - Prevent attackers from covering their tracks
- Make sure you use a NTP server.

Audit Log Backup

Maintaining a centralized backup copy of your logs is critically important to your monitoring. You must have a means to ensure that the logs were not modified, deleted, altered, or changed in order to consider them a reliable source of information. Many attackers know that the last thing to do before exiting a system is to erase all traces of malicious activity by removing logs, shell history, and a few other files that may leave evidence of their visit. This is why you should implement a centralized logging host where a copy of all logs will be sent. This centralized server has to be very secure, as it will contain important information that you will need if something ever goes wrong. By default, syslog does not provide any integrity features that can confirm the authenticity of the logs, but other third party utilities such as Syslog-NG have such features. It is also important to regularly back up your centralized syslog server in order to protect all of the logs that it has stored. In highly critical environments, a copy of logs can be sent to multiple servers at once. This will greatly increase the chances of having a reliable copy of the logs somewhere.

Just as a side note, you must ensure that all of your systems are using a reliable and accurate time source. This will ease log correlation. If the time is erroneous, it will be very difficult to reconstruct the events that took place.

Note

Imagine finding out your systems were hacked and you know the person in your company that did it. Wouldn't it be nice to be able to prosecute them? You won't be able to do it without the logs!

Reconstruction of Events

- Console messages
- Logs
- Correlation from multiple sources
- Extract data from system
- Not an easy task
- Seek help, and again, seek help

Reconstruction of Events

Reconstruction of events is sometimes necessary to determine how and why an incident occurred. It might allow you to see if this was the result of an abuse or simply an application or subsystem that failed to perform as expected.

The first thing to look for are error messages on the console that were generated by the system or its applications. This, of course, takes for granted that the application or system is built to provide error or warning messages.

The next location to look for information for the reconstruction of events is in the system logs. If for any reason you believe that your logs might not be reliable, use the copy on your central logging server. Logs from systems and firewalls can sometimes be gigabytes in size. You will need log-parsing tools to help you in your quest. After you put the puzzle together, you will probably need to correlate information from multiple sources, such as firewalls, routers, IDS (Intrusion Detection System), and other monitoring tools.

If the logs do not indicate what happened, you can dig through some of the system files or temporary space areas to see if they contain information that can help you.

The reconstruction effort is a puzzle. You need dedication and patience to succeed, but it can be done. You might need assistance from a forensic expert. You will be amazed to see how much information forensic experts can dig out of a system.

Protection against Alteration

- Use of integrity controls
- MAC
- Digital signatures

Protection against Alteration

Protection against alterations to your operating systems, applications, and other information is usually achieved by the use of integrity controls. These alterations could be accidental or malicious. The controls that are deployed will give you a high level of assurance that the information has not been modified.

Two cryptographic techniques are usually used to ensure the integrity of information. The first technique is the Message Authentication Code, usually referred to as MAC; the second one is through the use of digital signatures that prove the origin of the data, as well as its authenticity. Both techniques use a hashing algorithm that produces a digest of the information. If any modifications are made to the data, the digest will no longer match and changes will be detected. Some of the most popular hashing algorithms are MD5 and SHA1.

Today, the market has a whole series of tools that can help you detect unauthorized alteration. Products, such as Tripwire and Veracity, are commonly used for detecting changes to files. These products have centralized consoles where alterations are reported. They are fantastic for catching an intruder if malicious activity is not blocked or detected by your network perimeter defenses. Because they are host-based, they are also great for catching the potential insider who might have physical access to the system, but who is not performing malicious activity across the network itself.

Protection against Unavailability

- Single point of failure
- Redundancy
- Fail-over
- Load sharing
- Alternate site

Protection against Unavailability

Availability is one of the three tenets on which security is based. It is part of the CIA triangle (confidentiality, integrity, and availability) and the opposite of what we continually guard against, DAD (destruction, alteration, and disclosure).

To properly protect yourself against unavailability, you must go through the exercise of identifying the single points of failure in your environment. These single points of failure could be hardware, software, or human resources. If you have someone who keeps all information in his head to protect his job (at least he thinks!), then you might be in a lot of trouble if he gets hit by a car and is unable to come to work, or worse, if he is not able to continue his employment with your company. These are the weakest links you have to identify in your environment.

How much redundancy is needed will depend on the criticality of the service or resource and on how long the acceptable down time is for your company. What is the impact of not having Internet access or e-mail, for example? If you have an e-commerce server, the impact can be very high.

After the critical elements are clearly identified, you should ensure that you reduce the risks by adding redundancy into your environment. If you conduct online commerce and it is your main revenue source, then your connectivity to the Internet is crucial and you need redundant links and servers for the e-commerce service.

Fail-over mechanisms refer to having an alternative component take over immediately when the main component fails. These mechanisms are usually automated and do not require human intervention. There is a heartbeat monitoring the machines so when the primary mechanism becomes unavailable, the secondary one automatically takes over. Of course, the downside of such a mechanism is that you need a second resource that will not be used until a disaster occurs. Another method that makes better use of resources is load sharing.

Load sharing is the use of multiple servers to ensure continuous service. In this case, there is more than one server providing services and they are all used at the same time. Some firewalls, such as Checkpoint and StoneGate, allow for large clustering that can respond to the most demanding environment. Clustering is an easy way to manage a large number of servers because it treats them as one with third-party software.

Roles and Responsibilities (IS/IT)

- Policy
- Risk management
- Life cycle planning
- Auditing and monitoring
- Recovery strategies
- Incident handling
- Awareness

Roles and Responsibilities (IS/IT)

The roles and responsibilities regarding data security in an organization are important; these need to be defined and placed in the security policy to ensure that employees are aware of their security duties. The day-to-day functions in an IT department are beyond the scope of this section. Essentially it is critical that all of the key roles, including policy, risk management, auditing, and incident response, are covered either by the IT or the IS department.

IS/IT Functions

- Audit
- Physical security
- Disaster recovery
- Monitoring
- Incident response
- Training and awareness

There are some basic functions performed by most IS/IT departments, regardless of size. These are:

- Audit functions. This function ensures that the controls put in place offer the proper protection in accordance with policies.
- Physical security. This is sometimes delegated to the building owner or a third-party security service. In such a case, you should ensure that they provide proper levels of protection and that access control is properly enforced.
- Disaster recovery. This has taken on a life of its own since September 11, 2001. Disaster recovery used to be overlooked in most environments, but it is now one of the greatest worries for IT/IS managers.
- Monitoring. As previously discussed in this section, monitoring is done for performance, abuse, or other reasons. Ensure that it is done within the limits of the law and that it is effective.
- Incident response. This is another key point in the numerous IS/IT functions. Without a response capability, deploying controls is not of much value. If you monitor your network, review the data and take appropriate actions.

Finally, you have to train your user population at all levels—train them to recognize threats to their systems and information. Increase their awareness to make them more receptive to your protection plan and to understand why countermeasures are deployed. Also, trained employees will likely cause fewer incidents.

Roles and Responsibilities (Manager/Custodian)

- Usually the immediate supervisor of an employee
- Responsible for user IDs
- Responsible for contractors
- Looks after termination procedures
- Looks after passwords

Roles and Responsibilities (Manager)

A manager is usually a person who has an employee who works directly for him. Managers have the ultimate responsibility for information assets owned by the employees working for him. The manager also has to manage the assets of temporary employees, such as contractors, consultants, and interns. Here are a few guidelines for managers:

- Ensure that employees are aware of security policies, directives, guidelines, procedures, and standards.
- Receive the initial password for employees under your care; this ensures that an employee will not be granted access without the approval of his manager.
- Notify human resources and the IT security staff immediately when an employee terminates employment or is terminated. This ensures that accounts and access rights are revoked or suspended in a timely manner.
- Inform the security administrators of changes in a person's role to ensure that only required access is granted in accordance with the role that the employee plays.

Roles and Responsibilities (Owner)

- Final say towards security
- Decides what is appropriate
- Ultimately responsible
- Determines what backup to use
- Determines who can access

Roles and Responsibilities (Owner)

Companies often have complicated structures in which key information resources have assigned owners. These owners are responsible for defining the appropriate protection for the information under their care. Owners make the ultimate decisions, and they are accountable if a compromise, loss, or abuse occurs. Some of their responsibilities include:

- Assigning a classification to the information under their care
- Ensuring that proper security controls are in place to protect the information for which they are accountable
- Regularly reviewing who has access to the information under their care
- Determining what backups are needed
- Serving as the main point of contact to approve access to data or information under their care
- Naming someone else to replace them in case of absence

Roles and Responsibilities (User)

- Security involves all personnel.
- End users play a critical role.
- Users must be aware of their role,
- Awareness is key.
- Ensure proper training.

Roles and Responsibilities (User)

Security is a matter that should concern everyone in a company. End users must understand what their role with respect to security is and how they can contribute to maintaining proper security for the corporation. Users often notice strange problems on their computers. If they are trained well and are conscious of security, they will quickly learn to identify problems that might be related to security.

An end user is anyone in a company, including contractors, vendors, and partners, who use the company information resources as part of their daily tasks. Some user responsibilities include the following:

- They must not share user IDs and passwords with others.
- They must follow proper procedures to protect information under their care.
- They must use company assets only for company-related activities.
- They must be conversant with the policies, procedures, guidelines, and standards that they must follow.
- They have a responsibility to report security incidents that they are aware of.

Employee Sabotage

- Destruction of hardware
- Destruction of facility
- Planting bombs
- Deleting or modifying data
- Holding systems hostage
- To avoid it, be fair and honest.

Employee Sabotage

Sabotage happens regularly; however, because companies are afraid of negative publicity, you seldom hear anything about it. There are well-documented cases of employees who have been fired who took great care in causing the maximum amount of disruption before or after they left by using a logic bomb. There are employees who simply used a fire axe to cut and destroy data center equipment. Sabotage is real and it can happen to any company.

Protecting against sabotage is not always easy. You have to exercise great care in the way you handle promotions, terminations, or any other matter in your daily dealing with employees. If employees feel that they are cheated, harassed, or endangered, they might want revenge. This is when sabotage occurs. Most sabotage attacks cost a lot of money due to the privileged access employees have and their intimate knowledge of systems **and** networks.

Loss of Infrastructure

- Power failures
- Spike and brownouts
- Loss of communications
- Water outage or leaks
- Lack of transportation
- Fire, flood, civil unrest, and strike

Loss of Infrastructure

Here are some possible examples of loss of infrastructure:

- Power failures
- Spike and brownouts
- Loss of communications
- Water outage or leaks
- Lack of transportation
- Fire, flood, civil unrest, and strike

These losses often create a situation in which down time might be experienced. There is little you can do to prevent most of them because they are not easily predictable and/or are beyond our control. For example, a strike in public transportation is not directly related to your IT systems, but if your employees cannot get to work, it might become a problem. The same applies for cities that are accessible only through the use of bridges from one shore to the other; if the bridge is closed, you cannot get to work.

Violations and Reporting

- **Policies**
- **Compliance**
- **Audit role**
- **Procedures for reporting**
- **Part of security awareness training**
- **Disciplinary and administrative actions**

Violations and Reporting

Although proper screening takes place and proper controls are in place, there will always be people who intentionally or unintentionally attempt to trespass your network. Your policies and acceptable usage agreement should clearly state what employees can and cannot do, what is considered acceptable activity, and what is not. Ensure that the policy is disseminated to all new employees and that they have agreed to and signed the company's acceptable use agreement before they are granted access.

Violations are often discovered while conducting audits or through monitoring tools. There are cases in which a user might notice a violation. Users should be made aware, in their security awareness training, of what procedures they must follow if they discover a violation. There should be a documented procedure that they can follow or refer to if they are in doubt. All violations should be investigated and a motive established. In some cases, it might just be a lack of training or a mistake from a user, but in other cases it might be obvious that it was unauthorized activity.

Reporting Concepts

- Content
- Format
- Structure
- Hierarchy
- Escalation
- Frequency

Reporting Concepts

Audits are of no value unless the information discovered is put together in a structured and organized fashion. The report must be valuable to the high-level executive, the system administrator responsible for implementing the fixes, and all levels of personnel in between. A high-level executive will not want the details; however, a well-thought out summary of the security status of the network is vital to a CIO. Most organizations have fairly similar auditing reports. Each report should state, at a minimum, the purpose of the audit, the scope, and what results were found. Audit reports also typically include the auditor's name, the date and time, what systems were evaluated, the location of the audit, and other pertinent information that may help the auditors.

The reports should be structured in such a way that they are easy to read and understand. Reports are usually distributed to the people that are directly responsible for taking corrective measures. These people should always ensure that the report is given proper consideration. Usually a summary of the report is given only to upper management who might not be interested in the technical jargon of the detailed report.

Audit reports are usually produced right after an audit has been completed; however, in case of large audits there might be periodic reports produced to demonstrate the progression of the audit. In case a serious weakness is discovered that might endanger security perimeters, a special report might be prepared immediately to address this shortcoming as quickly as possible.

Penetration Testing (Operations Evaluation)

- War dialing
- Sniffing
- Eavesdropping
- Radiation monitoring
- Dumpster diving
- Social engineering

After the operational security plan is in place, it must be tested. *Penetration testing* is the process of examining the limitations of the security measures in place. Some tests include :

- *War dialing* — attempts to attack the systems via dialing all the phone numbers in an exchange
- *Sniffing* — passively monitors network traffic for network knowledge, such as passwords
- *Eavesdropping* — involves listening to phone conversations
- *Radiation monitoring* — the process of receiving images, data, or audio from an unprotected source by listening to radiation signals
- *Dumpster diving* — obtains passwords and corporate directories by searching through discarded media
- *Social engineering* — a euphemism for non-technical or low-technology means of breaching security, such as lies, impersonation, tricks, bribes, blackmail, and threats. These are used to attack information systems.

Avoiding Threats to Operation Security

- Errors and omissions
- Fraud and theft
- Employee sabotage
- Malicious attackers
- Malicious code

Threats to operation security come in many forms. The following is not an exhaustive list:

- Errors and omissions can be avoided with proper audit.
- Fraud and theft can be avoided with proper administrative controls.
- Employee sabotage by giving people too much access should be avoided.
- Malicious hackers and crackers can be thwarted with proper intrusion prevention and detection.
- Malicious code can be avoided with proper code review and auditing.

Summary

- Maintaining a proper security stance across an organization is critical.
- Understanding legal requirements is important to properly protect the confidentiality of critical information.
- It is critical that proper operations security gets mapped to critical risks.

We cannot discuss all the issues you might run across in operations security during your tenure as a security professional. You should now have a glimpse into what comprises a good security environment and the information you need to build and maintain a good secure stance.

We talked about legal requirements for intellectual property, due diligence, due care, and why we must care about these. We also discussed how you can interact with an HR department to make sure you do not hire the wrong people. We also talked about how correct policies must be in place in the event a bad apple slips through the screening process. We also talked about how to mitigate the threat from a potential bad apple with job rotation, mandatory vacations, least privilege, and the need to know for specific information.

If the bad apple is successful in doing something malicious to the network, you can identify illegal activity and control it through the use of different types of operational controls including preventive controls, detective controls, corrective controls, and recovery controls. These controls are what allow you to do the job. They consist of many things including policies, IDSs, firewalls, and recovery tools.