CHAPTER 1

BRIEF HISTORY AND MISSION OF INFORMATION SYSTEM SECURITY

Seymour Bosworth and Robert V. Jacobson

1.1 INTRODUCTION TO INFORMATION SYSTEM SECURITY.

The growth of computers and of information technology has been explosive. Never before has an entirely new technology been propagated around the world with such speed and with so great a penetration of virtually every human activity. Computers have brought vast benefits to fields as diverse as human genome studies, space exploration, artificial intelligence, and a host of applications from the trivial to the most life-enhancing.

Unfortunately, there is also a dark side to computers: They are used to design and build weapons of mass destruction as well as military aircraft, nuclear submarines,

and reconnaissance space stations. The computer's role in formulating biologic and chemical weapons, and in simulating their deployment, is one of its least auspicious uses.

Of somewhat lesser concern, computers used in financial applications, such as facilitating the purchase and sales of everything from matchsticks to mansions, and transferring trillions of dollars each day in electronic funds, are irresistible to miscreants; many of them see these activities as open invitations to fraud and theft. Computer systems, and their interconnecting networks, are also prey to vandals, malicious egotists, terrorists, and an array of individuals, groups, companies, and governments intent on using them to further their own ends, with total disregard for the effects on innocent victims. Besides these intentional attacks on computer systems, there are innumerable ways in which inadvertent errors can damage or destroy a computer's ability to perform its intended functions.

Because of these security problems and because of a great many others described in this volume, the growth of information systems security has paralleled that of the computer field itself. Only by a detailed study of the potential problems, and implementation of the suggested solutions, can computers be expected to fulfill their promise, with few of the security lapses that plague less adequately protected systems. This chapter defines a few of the most important terms of information security and includes a very brief history of computers and information systems, as a prelude to the works that follow.

Security can be defined as the state of being free from danger and not exposed to damage from accidents or attack, or it can be defined as the process for achieving that desirable state. The objective of information system security[1] is to optimize the performance of an organization with respect to the risks to which it is exposed.

Risk is defined as the chance of injury, damage, or loss. Thus, risk has two elements: (1) chance—an element of uncertainty, and (2) potential loss or damage. Except for the possibility of restitution, information system security actions taken today work to reduce *future* risk losses. Because of the uncertainty about future risk losses, perfect security, which implies zero losses, would be infinitely expensive. For this reason, risk managers strive to optimize the allocation of resources by minimizing the total cost of information system security measures taken and the risk losses experienced. This optimization process is commonly referred to as risk management.

Risk management in this sense is a three-part process:

1. Identification of material risks
2. Selection and implementation of measures to mitigate the risks
3. Tracking and evaluating of risk losses experienced, in order to validate the first two parts of the process

The purpose of this *Handbook* is to describe information security system risks, the measures available to mitigate these risks, and techniques for managing security risks. (For a more detailed discussion of risk assessment and management, see Chapter 47 and Chapter 54.)

Risk management has been a part of business for centuries. Renaissance merchants often used several vessels simultaneously, each carrying a portion of the merchandise, so that the loss of a single ship would not result in loss of the entire lot. At almost the same time, the concept of insurance evolved, first to provide economic protection against the loss of cargo and later to provide protection against the loss of buildings by fire. Fire insurers and municipal authorities began to require adherence to standards

intended to reduce the risk of catastrophes like the Great Fire of London in 1666. The Insurance Institute was established in London one year later. With the emergence of corporations, as limited liability stock companies, corporate directors have been required to use prudence and due diligence in protecting shareholders' assets. Security risks are among the threats to corporate assets that directors have an obligation to address.

Double-entry bookkeeping, another Renaissance invention, proved to be an excellent tool for measuring and controlling corporate assets. One objective was to make insider fraud more difficult to conceal. The concept of separation of duties emerged, calling for the use of processing procedures that required more than one person to complete a transaction. As the books of account became increasingly important, accounting standards were developed, and they continue to evolve to this day. These standards served to make books of account comparable and to assure outsiders that an organization's books of account presented an accurate picture of its condition and assets. These developments led, in turn, to the requirement that an outside auditor perform an independent review of the books of account and operating procedures.

The transition to automated accounting systems introduced additional security requirements. Some early safeguards, such as the rule against erasures or changes in the books of account, no longer applied. Some computerized accounting systems lacked an audit trail, and others could have the audit trail subverted as easily as actual entries.

Finally, with the advent of the Information Age, intellectual property has become an increasingly important part of corporate and governmental assets. At the same time that intellectual property has grown in importance, threats to intellectual property have become more dangerous, because of information system (IS) technology itself. When sensitive information was stored on paper and other tangible documents, and rapid copying was limited to photography, protection was relatively straightforward. Nevertheless, document control systems, information classification procedures, and need-to-know access controls were not foolproof, and information compromises occurred with dismaying regularity. Evolution of IS technology has made information control several orders of magnitude more complex. The evolution and, more important, the implementation of control techniques have not kept pace.

The balance of this chapter describes how the evolution of information systems has caused a parallel evolution of IS security and at the same time has increased the importance of anticipating the impact of technical changes yet to come. This overview will clarify the factors leading to today's IS security risk environment and mitigation techniques and will serve as a warning to remain alert to the implication of technical innovations as they appear. The remaining chapters of this *Handbook* discuss IS security risks, threats, and vulnerabilities, their prevention and remediation, and many related topics in considerable detail.

1.2 EVOLUTION OF INFORMATION SYSTEMS.

The first electromechanical punched-card system for data processing, developed by Herman Hollerith at the end of the nineteenth century, was used to tabulate and total census field reports for the U.S. Bureau of the Census in 1890. The first digital, stored-program computers developed in the 1940s were used for military purposes, primarily cryptanalysis and the calculation and printing of artillery firing tables. At the same time, punched-card systems were already being used for accounting applications and were an obvious choice for data input to the new electronic computing machines.

1.2.1 1950s: Punched-Card Systems. In the 1950s, punched-card equipment dominated the commercial computer market.[2] These electromechanical devices could perform the full range of accounting and reporting functions. Because they were programmed by an intricate system of plugboards with a great many plug-in cables, and because care had to be exercised in handling and storing punched cards, only experienced persons were permitted near the equipment. Although any of these individuals could have set up the equipment for fraudulent use, or even engaged in sabotage, apparently few, if any, actually did so.

The punched-card accounting systems typically used four processing steps. As a preliminary, operators would be given a "batch" of documents, typically with an adding machine tape showing one or more "control totals." The operator keyed the data on each document into a punched card and then added an extra card, the batch control card, which stored the batch totals. Each card consisted of 80 columns, each containing, at most, one character. A complete record of an inventory item, for example, would be contained on a single card. The card was called a unit record, and the machines that processed the cards were called either unit record or punched-card machines. It was from the necessity to squeeze as much data as possible into an 80-character card that the later Y2K problem arose. Compressing the year into two characters was a universally used space-saving measure; its consequences 40 years later were not foreseen.

A group of punched cards, also called a "batch," were commonly held in a metal tray, Sometimes a batch would be rekeyed by a second operator, using a "verify-mode" rather than actually punching new holes in the cards, in order to detect keypunch errors before processing the card deck. Each batch of cards would be processed separately, so the processes were referred to as "batch jobs."

The first step would be to run the batch of cards through a simple program, which would calculate the control totals and compare them with the totals on the batch control card. If the batch totals did not reconcile, the batch was sent back to the keypunch area for rekeying. If the totals reconciled, the deck would be sort-merged with other batches of the same transaction type, for example, the current payroll. When this step was complete, the new batch consisted of a punched card for each employee in employee-number order. The payroll program accepted this input data card deck and processed the cards one by one. Each card was matched up with the corresponding employee's card in the payroll master deck to calculate the current net pay and itemized deductions and to punch a new payroll master card including year-to-date totals. The final step was to use the card decks to print payroll checks and management reports. These steps were identical with those used by early, small-scale electronic computers. The only difference was in the speed at which the actual calculations were made. A complete process was still known as a batch job.

With this process, the potential for abuse was great. The machine operator could control every step of the operation. Although the data were punched into cards and verified by others, there was always a keypunch machine nearby for use by the machine operator. Theoretically, that person could punch a new payroll card and a new batch total card to match the change before printing checks and again afterward. The low incidence of reported exploits was due to the controls that discouraged such abuse and possibly to the pride that machine operators experienced in their jobs.

1.2.2 Large-Scale Computers. While these electromechanical punched card machines were sold in large numbers, research laboratories and universities were working to design large-scale computers that would have a revolutionary effect on

the entire field. These computers, built around vacuum tubes, are known as the first generation. In March 1951, the first Universal Automatic Computer (UNIVAC) was accepted by the U.S. Census Bureau. Until then, every computer had been a one-off design, but UNIVAC was the first large-scale, mass-produced computer, with a total of 46 built. The word "universal" in its name indicated that UNIVAC was also the first computer designed for both scientific and business applications.[3]

UNIVAC contained 5,200 vacuum tubes, weighed 29,000 pounds, and consumed 125 kilowatts of electrical power. It dispensed with punched cards, receiving input from half-inch-wide metal tape recorded from keyboards, with output either to a similar tape or to a printer. Although not a model for future designs, its memory consisted of 1,000 72-bit words and was fabricated as a mercury delay line. Housed in a cabinet about six feet tall, two feet wide, and two feet deep was a mercury-filled coil running from top to bottom. A transducer at the top propagated slow-moving waves of energy down the coil to a receiving transducer at the bottom. There it was reconverted into electrical energy and passed on to the appropriate circuit, or recirculated if longer storage was required.

In 1956, IBM introduced the RAMAC (Random Access Method of Accounting and Control) magnetic disk system. It consisted of 50 magnetically coated metal disks, each 24 inches in diameter, and mounted on a common spindle. A servomotor, controlled by feedback from digital addresses read off each track of a disk, moved two coupled read/write heads to span each side of the disk, and then inward to any one of 100 tracks. In one revolution of the disks, any or all of the information on those two tracks could be read out, or recorded. The entire system was almost the size of a compact car and held what, for that time, was a tremendous amount of data—5 megabytes. The cost was $10,000 per megabyte, or $35,000 per year to lease. This compares with some of today's magnetic hard drives that measure about $3\frac{1}{2}$ inches wide by 1 inch high, store as much as 1,000 gigabytes, and cost less than $400, or about $0.0004 per megabyte.

Those early, massive computers were housed in large, climate-controlled rooms. Within the room, a few knowledgeable experts, looking highly professional in their white laboratory coats, attended to the operation and maintenance of their million-dollar charges. The concept of a "user" as someone outside the computer room who could interact directly with the actual machine did not exist.

Service interruptions, software errors, and hardware errors were usually not critical. If any of these caused a program to fail or abort, beginning again was a relatively simple matter. Consequently, the primary security concerns were physical protection of the scarce and expensive hardware, and measures to increase their reliability. Another issue, then as now, was human fallibility. Because the earliest computers were programmed in extremely difficult machine language, consisting solely of ones (1s) and zeroes (0s), the incidence of human error was high and the time to correct errors was excessively long. Only later were assembler and compiler languages developed to increase the number of people able to program the machines and to reduce the incidence of errors and the time to correct them.

Information system security for large-scale computers was not a significant issue then for two reasons. First, only a few programming experts were able to utilize and manipulate computers. Second, there were very few computers in use, each of which was extremely valuable, important to its owners, and, consequently, closely guarded.

1.2.3 Medium-Size Computers. In the 1950s, smaller computer systems were developed with a very simple configuration; punched-card master files were replaced by punched paper tape and, later, by magnetic tape, and disk storage systems.

1 · 6 BRIEF HISTORY AND MISSION OF INFORMATION SYSTEM SECURITY

The electromechanical calculator with its patchboard was replaced by a central processor unit (CPU) that had a small main memory, sometimes as little as 8 kilobytes,[4] and limited processing speed and power. One or two punched-card readers could read the data and instructions stored on that medium. Later, programs and data files were stored on magnetic tape. Output data were sent to cardpunches, for printing on unit record equipment and later to magnetic tape. There was still no wired connection to the outside world, and there were no online users because no one, besides electronic data processing (EDP) people within the computer room, could interact directly with the system. These systems had very simple operating systems and did not use multiprocessing; they could run only one program at a time.

The IBM Model 650, as an example, introduced in 1954, measured about 5 feet by 3 feet by 6 feet and weighed almost 2,000 pounds. Its power supply was mounted in a similarly sized cabinet, weighing almost 3,000 pounds. It had 2,000 (10-digit) words of magnetic drum primary memory, with a total price of $150,000 or a rental fee of $3,200 per month. For an additional $1,500 per month, a much faster core memory, of 60 words, could be added. Input and output both utilized read/write punch-card machines. The typical 1950s IS hardware was installed in a separate room, often with a viewing window so that visitors could admire the computer. In an early attempt at security, visitors actually within the computer room were often greeted by a printed sign saying:

> Achtung! Alles Lookenspeepers!
>
> Das computermachine ist nicht fur gefingerpoken und mittengrabben.
> Ist easy schnappen der springenwerk, blowenfusen, und poppencorken mit spitzensparken.
> Ist nicht fur gewerken bei das dumbkopfen. Das rubbernecken sightseeren keepen hans in das pockets muss. . .:
> Relaxen und watch das blinkenlichten.

Since there were still no online users, there were no user IDs and passwords. Programs processed batches of data, run at a regularly scheduled time—once a day, once a week, and so on, depending on the function. If the data for a program were not available at the scheduled run time, the operators might run some other job instead and wait for the missing data. As the printed output reports became available, they were delivered by hand to their end users. End users did not expect to get a continuous flow of data from the information processing system, and delays of even a day or more were not significant, except perhaps with paycheck production.

Information system security was hardly thought of as such. The focus was on batch controls for individual programs, physical access controls, and maintaining a proper environment for the reliable operation of the hardware.

1.2.4 1960s: Small-Scale Computers.

During the 1960s, before the introduction of small-scale computers, dumb[5] terminals provided users with a keyboard to send a character stream to the computer and a video screen that could display characters transmitted to it by the computer. Initially these terminals were used to help computer operators control and monitor the job stream while replacing banks of switches and indicator lights on the control console. However, it was soon recognized that these terminals could replace card readers and keypunch machines as well. Now users, identified by user IDs and authenticated with passwords, could enter input data through a cathode ray tube (CRT) terminal into an edit program, which would validate the input

and then store it on a hard drive until it was needed for processing. Later it was realized that users also could directly access data stored in online master files.

1.2.5 Transistors and Core Memory. The IBM 1401, introduced in 1960, with a core memory of 4,096 characters, was the first all-transistor computer, marking the advent of the second generation. Housed in a cabinet measuring 5 feet by 3 feet, the 1401 required a similar cabinet to add an additional 12 kilobytes of main memory. Just one year later, the first integrated circuits were used in a computer, making possible all future advances in miniaturizing small-scale computers and in reducing the size of mainframes significantly.

1.2.6 Time Sharing. In 1961, the Compatible Time Sharing System (CTSS) was developed for the IBM 7090/7094. This operating system software, and its associated hardware, was the first to provide simultaneous remote access to a group of online users through multiprogramming.[6] "Multiprogramming" means that more than one program can appear to execute at the same time. A master control program, usually called an operating system (OS), managed execution of the functional applications programs. For example, under the command of the operator, the OS would load and start application 1. After 50 milliseconds, the OS would interrupt the execution of application 1 and store its current state in memory. Then the OS would start application 2 and allow it to run for 50 milliseconds, and so on. Usually, within a second after users had entered keyboard data, the OS would give their applications a time slice to process the input. During each time slice, the computer might execute hundreds of instructions. These techniques made it appear as if the computer were entirely dedicated to each user's program. This was true only so long as the number of simultaneous users was fairly small. After that, as the number grew, the response to each user slowed down.

1.2.7 Real-Time, Online Systems. Because of multiprogramming and the ability to store records online and accessible in random order, it became feasible to provide end users with direct access to data. For example, an airline reservation system stores a record of every seat on every flight for the next 12 months. A reservation clerk, working at a terminal, can answer a telephoned inquiry, search for an available seat on a particular flight, quote the fare, sell a ticket to the caller, and reserve the seat. Similarly, a bank officer can verify an account balance and effect money transfers. In both cases, each data record can be accessed and modified immediately, rather than having to wait for a batch to be run. Today both the reservation clerk and the bank officer can be replaced by the customers themselves, who directly interface with the online computers.

While this advance led to a vast increase in available computing power, it also increased greatly the potential for breaches in computer security. With more complex operating systems, with many users online to sensitive programs, and with databases and other files available to them, protection had to be provided against inadvertent error and intentional abuse.

1.2.8 A Family of Computers. In 1964, IBM announced the S/360 family of computers, ranging from very small-scale to very large-scale models. All of the six models used integrated circuits, which marked the beginning of the third generation of computers. Where transistorized construction could permit up to 6,000 transistors per cubic foot, 30,000 integrated circuits could occupy the same volume. This lowered the costs substantially, and companies could buy into the family at a price within their

means. Because all computers in the series used the same programming language and the same peripherals, companies could upgrade easily when necessary. The 360 family quickly came to dominate the commercial and scientific markets. As these computers proliferated, so did the number of users, knowledgeable programmers, and technicians. Over the years, techniques and processes were developed to provide a high degree of security to these mainframe systems.

The year 1964 also saw the introduction of another computer with far-reaching influence: the Digital Equipment Corp. (DEC) PDP-8. The PDP-8 was the first mass-produced true minicomputer. Although its original application was in process control, the PDP-8 and its progeny quickly proved that commercial applications for minicomputers were virtually unlimited. Because these computers were not isolated in secure computer rooms but were distributed throughout many unguarded offices in widely dispersed locations, totally new risks arose, requiring innovative solutions.

1.2.9 1970s: Microprocessors, Networks, and Worms. The foundations of all current personal computers (PCs) were laid in 1971 when Intel introduced the 4004 computer on a chip. Measuring $\frac{1}{16}$ inch long by $\frac{1}{8}$ inch high, the 4004 contained 2,250 transistors with a clock speed of 108 kiloHertz. The current generation of this earliest programmable microprocessor contains millions of transistors, with speeds over 1 gigaHertz, or more than 10,000 times faster. Introduction of microprocessor chips marked the fourth generation.

1.2.10 First Personal Computers. Possibly the first personal computer was advertised in *Scientific American* in 1971. The KENBAK–1, priced at $750, had three programming registers, five addressing modes, and 256 bytes of memory. Although not many were sold, the KENBACK–1 did increase public awareness of the possibility for home computers.

It was the MITS Altair 8800 that became the first personal computer to sell in substantial quantities. Like the KENBACK–1, the Altair 8800 had only 256 bytes of memory, but it was priced at $375 without keyboard, display, or secondary memory. About one year later, the Apple II, designed by Steve Jobs and Steve Wozniak, was priced at $1,298, including a CRT display and a keyboard.

Because these first personal computers were entirely stand-alone and usually under the control of a single individual, there were few security problems. However, in 1978, the VisiCalc spreadsheet program was developed. The advantages of standardized, inexpensive, widely used application programs were unquestionable, but packaged programs, as opposed to custom designs, opened the way for abuse because so many people understood their user interfaces as well as their inner workings.

1.2.11 First Network. A national network, conceived in late 1969, was born as ARPANET[7] (Advanced Research Projects Agency Network), a Department of Defense–sponsored effort to link a few of the country's important research universities, with two purposes: to develop experience in interconnecting computers and to increase productivity through resource sharing. This earliest connection of independent large-scale computer systems had just four nodes: the University of California at Los Angeles (UCLA), the University of California at Santa Barbara, Stanford Research Institute, and the University of Utah. Because of the inherent security in each leased-line interconnected node, and the physically protected mainframe computer rooms, there was no apparent concern for security issues. From this simple network, with no thought

of security designed in, there finally evolved today's ubiquitous Internet and the World Wide Web (WWW) with their vast potential for security abuses.

1.2.12 Further Security Considerations. With the proliferation of remote terminals on commercial computers, physical control over access to the computer room was no longer sufficient. In response to the new vulnerabilities, logical access control systems were developed. An access control system maintains an online table of authorized users. A typical user record would store the user's name, telephone number, employee number, and information about the data the user was authorized to access and the programs the user was authorized to execute. A user might be allowed to view, add, modify, and delete data records in different combinations for different programs.

At the same time, system managers recognized the value of being able to recover from a disaster that destroyed hardware and data. Data centers began to make regular tape copies of online files and software for off-site storage. Data center managers also began to develop and implement off-site disaster recovery plans, often involving the use of commercial disaster-recovery facilities. Even with such a system in place, new vulnerabilities were recognized throughout the following years, and these are the subjects of much of this *Handbook*.

1.2.13 First "Worm." A prophetic science-fiction novel, *The Shockwave Rider*, by John Brunner[8] (1975), depicted a "worm" that grew continuously throughout a computer network. The worm eventually exceeded a billion bits in length and became impossible to kill without destroying the network. Although actual worms later became real and present menaces to all networked computers, prudent computer security personnel install, and regularly update, antivirus programs that effectively kill viruses and worms without having to kill the network.

1.2.14 1980s: Productivity Enhancements. The decade of the 1980s might well be termed the era of productivity enhancement. The installation of millions of personal computers in commercial, industrial, and government applications enhanced efficiency and functionality of vast numbers of users. These advances, which could have been achieved in no other way, were made at costs that virtually any business could afford.

1.2.15 Personal Computer. In 1981, IBM introduced a general-purpose small computer it called the "Personal Computer." That model and similar systems became known generically as PCs. Until then, small computers were produced by relatively unknown sources, but IBM, with its worldwide reputation, brought PCs into the mainstream. The fact that IBM had demonstrated a belief in the viability of PCs made them serious contenders for corporate use.

There were many variations on the basic Model 5100 PC, and sales expanded far beyond IBM's estimates. The basic configuration used the Intel 8088, operating at 4.77 megaHertz, with up to two floppy disk drives, each with a capacity of 160 kilobyte and with a disk-based operating system (DOS) in an open architecture. This open OS architecture, with its available "hooks," made possible the growth of independent software producers, the most important of which was the Microsoft Corporation, formed by Bill Gates and Paul Allen.

IBM had arranged for Gates and Allen to create the DOS operating system. Under the agreement, IBM would not reimburse Gates and Allen for their development costs; rather, all profits from the sale of DOS would accrue to them. IBM did not have

an exclusive right to the operating system, and Microsoft began selling it to many other customers as MS-DOS. IBM initially included with its computer the VisiCalc spreadsheet program, but soon sales of Lotus 1-2-3 surpassed those of VisiCalc. The open architecture not only made it possible for many developers to produce software that would run on the PC but also enabled anyone to put together purchased components into a computer that would compete with IBM's PC. The rapid growth of compatible application programs, coupled with the ready availability of compatible hardware, soon resulted in sales of more than 1 million units. Many subsequent generations of the original hardware and software are still producing sales measured in millions every year.

Apple took a very different approach with its Macintosh computer. Where IBM's system was wide open, Apple maintained tight control over any hardware or software designed to operate on the Macintosh so as to assure compatibility and ease of installation. The most important Apple innovations were the graphical user interface (GUI) and the mouse, both of which worked together to facilitate ease of use. Microsoft had attempted in 1985 to build these features into the Windows operating system, but early versions were generally rejected as slow, cumbersome, and unreliable. It was not until 1990 that Windows 3.0 overcame many of its problems and provided the foundation for later versions that were almost universally accepted.

1.2.16 Local Area Networks. During the 1980s, stand-alone desktop computers began to perform word processing, financial analysis, and graphic processing. Although this arrangement was much more convenient for end users than was a centralized facility, it was more difficult to share data with others.

As more powerful PCs were developed, it became practical to interconnect them so that their users could easily share data. These arrangements were commonly referred to as local area networks (LANs) because the hardware units were physically close, usually in the same building or office area. LANs have remained important to this day. Typically, a more powerful PC with a high storage capacity fixed[9] disk was designated as the file server. Other PCs, referred to as workstations, were connected to the file server using network interface cards installed in the workstations with cables between these cards and the file server. Special network software installed on the file server and workstations made it possible for workstations to access defined portions of the file server fixed disk just as if these portions were installed on the workstations. Furthermore, these shared files could be backed up at the file server without depending on individual users. By 1997, it was estimated that worldwide there were more than 150 million PCs operating as LAN workstations. The most common network operating systems (NOS) were Novell NetWare and later Microsoft IE (Internet Explorer).

Most LANs were implemented using the Ethernet (IEEE 802.3) protocol.[10] The server and workstations could be equipped with a modem (modulator/demodulator) connected to a dedicated telephone line. The modem enabled remote users, with a matching modem, to dial into the LAN and log on. This was a great convenience to LAN users who were traveling or working away from their offices, but such remote access created yet another new security issue. For the first time, computer systems were exposed in a major way to the outside world. From then on, it was possible to interact with a computer from virtually anywhere and from locations not under the same physical control as the computers themselves.

Typical NOS logical access control software provided for user-IDs and passwords and selective authority to access file server data and program files. A workstation user logged on to the LAN by executing a log-in program resident on the file server.

The program prompted the user to enter an ID and password. If the log-in program concluded that the ID and password were valid, it consulted an access-control table to determine which data and programs the user might access. Access modes were defined as read-only, execute-only, create, modify (write or append), lock, and delete, with respect to individual files and groups of files. The LAN administrator maintained the access control table using a utility program. The effectiveness of the controls depended on the care taken by the administrator, and so, in some circumstances, controls could be weak. It was essential to protect the ID and password of the LAN administrator since, if they were compromised, the entire access-control system became vulnerable. Alert information system security officers noted that control over *physical* access to LAN servers was critical in maintaining the logical access controls. Intruders who could physically access a LAN server could easily restart the server using their own version of the NOS, completely bypassing the installed logical access controls.

Superficially, a LAN appears to be the same as a 1970s mainframe with remote dumb terminals. The difference technically is that each LAN workstation user is executing programs on the workstation, not on the centralized file server, while mainframe computers use special software and hardware to run many programs concurrently, one program for each terminal. To the user at a workstation or remote terminal, the two situations appear to be the same, but from a security standpoint, there are significant differences. The mainframe program software stays on the mainframe and cannot, under normal conditions, be altered during execution. A LAN program on a workstation can be altered, for example, by a computer virus, while actually executing. As a rule, mainframe remote terminals cannot download and save files whereas workstations usually have at least a removable disk drive. Furthermore, a malicious workstation user can easily install a rewritable CD device, which makes it much easier to copy and take away large amounts of data.

Another important difference is the character of the connection between the computer and the terminals. Each dumb terminal has a dedicated connection to its mainframe and receives only data that is directed to it. A LAN operates more like a set of radio transmitters sharing a common frequency on which the file server and the workstations take turns "broadcasting" messages. Each message includes a "header" block that identifies the intended recipient, but every node (the file server and the workstations) on a LAN receives all messages. Under normal circumstances, each node ignores messages not addressed to it. However, it is technically feasible for a workstation to run a modified version of the NOS that allows it to capture all messages. In this way, a workstation could identify all log-in messages and record the user IDs and passwords of all other users on the LAN, giving it complete access to all of the LAN's data and facilities.

Mainframe and LAN security also differ greatly in the operating environment. As noted, the typical mainframe is installed in a separate room and is managed by a staff of skilled technicians. The typical LAN file server, however, is installed in ordinary office space and is managed by a part-time, remotely located LAN administrator who may not be adequately trained. Consequently, the typical LAN has a higher exposure to tampering, sabotage, and theft. However, if the typical mainframe is disabled by an accident, fire, sabotage, or any other security incident, many business functions will be interrupted, whereas the loss of a LAN file server usually disrupts only a single function.

1.2.17 1990s: Total Interconnection. With the growing popularity of LANs, the technologies for interconnecting them emerged. These networks of

physically interconnected local area networks were called wide area networks, or WANs. Any node on a LAN could access every node on any other interconnected LAN, and in some configurations, those nodes might also be given access to mainframe and minicomputer files and to processing capabilities.

1.2.18 Telecommuting. Once the WAN technology was in place, it became feasible to link LANs together by means of telecommunications circuits. It had been expensive to do this with the low-speed, online systems of the 1970s because all data had to be transmitted over the network. Now, since processing and most data used by a workstation were on its local LAN, a WAN network was much less expensive. Low-traffic LANs were linked using dial-up access for minimum costs, while major LANs were linked with high-speed dedicated circuits for better performance. Apart from dial-up access, all network traffic typically flowed over nonswitched private networks. Of the two methods, dial-up communications were considerably more vulnerable to security violations, and they remain so to this day.

1.2.19 Internet and the World Wide Web. The Internet, which began life in 1969 as the ARPANET, slowly emerged onto the general computing scene during the 1980s. Initially, access to the Internet was restricted to U.S. Government agencies and their contractors. ARPANET users introduced the concept of e-mail as a convenient way to communicate and exchange documents. Then, in 1989–1990, Tim Berners-Lee conceived of the World Wide Web and the Web browser. This one concept produced a profound change in the Internet, greatly expanding its utility and creating an irresistible demand for access. During the 1990s, the U.S. Government relinquished its control, and the Internet became the gigantic, no-one-is-in-charge network of networks it is today.

The Internet offers several important advantages: The cost is relatively low, connections are available locally in most industrialized countries, and by adopting the Internet protocol, Transmission Control Protocol/Internet Protocol (TCP/IP), any computer becomes instantly compatible with all other Internet users.

The World Wide Web technology made it easy for anyone to access remote data. Almost overnight, the Internet became the key to global networking. Internet service providers (ISPs) operate Internet-compatible computers with both dial-up and dedicated access. A computer may access an ISP directly as a stand-alone ISP client or via a gateway from a LAN or WAN. A large ISP may offer dial-up access at many locations, sometimes called points of presence (POPs), interconnected by its own network. ISPs establish links with one another through the national access points (NAPs) initially set up by the National Science Foundation. With this "backbone" in place, any node with access can communicate with another node, connected to a different ISP, located halfway around the globe, without making prior arrangements.

The unrestricted access provided by the Internet created new opportunities for organizations to communicate with clients. A company can implement a Web server with a full-time connection to an ISP and open the Web server, and the WWW pages it hosts, to the public. A potential customer can access a Web site, download product information and software updates, ask questions, and even order products. Commercial Web sites, as they evolved from static "brochure-ware" to online shopping centers, stock brokerages, and travel agencies, to name just a few of the uses, became known as e-businesses.

1.3 GOVERNMENT RECOGNITION OF INFORMATION ASSURANCE.

Certain major events in the history of information assurance (IA) center on government initiatives. In particular, IA has been strongly influenced by the development of security standards starting in the 1980s, by the publication of the landmark publication *Computers and Risk*[11] in 1991, and by the establishment of the InfraGard program in the late 1990s for protection of the U.S. critical infrastructure.

1.3.1 IA Standards.

In the late 1970s, the U.S. Department of Defense "established a Computer Security Initiative to foster the wide-spread availability of trusted computer systems."[12] The author of the initial report that later became the *Trusted Computer Systems Evaluation Criteria* (TCSEC), DoD Standard 5200.28, wrote: "Trusted computer systems are operating systems capable of preventing users from accessing more information than that to which they are authorized. Such systems are in great demand as more processing is entrusted to computers, while less information should be shared by all the system's users. With this demand comes a need to ascertain the integrity of computer systems on the market." The TCSEC was issued with a bright orange cover and became known as the *Orange Book*. Under the direction of National Computer Security Center (NCSC) director Patrick Gallagher and others, the National Security Agency (NSA) issued a series of books known as the Rainbow Series that profoundly affected the direction of IA in the United States and globally.[13]

The Rainbow Series led to similar efforts in other countries, culminating in the Common Criteria Evaluation and Validation Scheme (CCEVS), which has become the international standard for defining security levels for systems and software and for determining acceptable methods for testing and certifying system compliance with such standards.[14]

For details of the evolution of security standards, see Chapter 51 in this *Handbook*.

1.3.2 Computers at Risk. [15]

In 1988, the Defense Advanced Research Projects Agency (DARPA) asked the Computer Science and Technology Board (renamed the Computer Science and Telecommunications Board of the National Research Council [NRC] in 1990) for a study of computer and communications security issues affecting U.S. Government and industry. The NRC's System Security Study Committee published its results in a readable and informative book, *Computers at Risk: Safe Computing in the Information Age*. [16]

The committee included experts with impeccable credentials, including executives from major computer vendors such as Hewlett-Packard, DEC, and IBM; from high-technology companies such as Shearson, Lehman, Hutton Inc., and Rockwell International; universities such as Harvard and the Massachusetts Institute of Technology; and think tanks like the RAND Corporation.

A public misconception is the supposed divergence in focus of the military and of commerce: The military usually is described as concerned with external threats and the problem of disclosure, whereas businesses are said to worry more about insider threats to data integrity. On the contrary, the military and commerce need to protect data in similar ways. The differences arise primarily from (1) the sophistication and resources available to governments that try to crack foreign military systems; (2) the relatively strong military emphasis on prevention compared with commercial need for proof that can be used in legal proceedings; and (3) the fact that the military can access deep background checks on personnel, in contrast with the limits imposed on the invasion of privacy in the commercial sector.

Some of the more interesting points raised by the NRC Committee assert that:

- Because of the rapid and discontinuous pace of innovation in the computer field, "with respect to computer security, the past is not a good predictor of the future."
- Embedded systems (those where the microprocessor is not accessible to user reprogramming, as in medical imaging systems) open up greater risks from inadequate quality assurance (e.g., a software bug in a Therac 25 linear accelerator killed three patients by irradiating them with more than 100 times the intended radiation dosage).[17]
- Networking makes it possible to harm many more systems: "Interconnection gives an almost ecological flavor to security; it creates dependencies that can harm as well as benefit the community."

The committee proposed five major recommendations, summarized next:

1. Push for implementation of generally accepted system security principles including:
   - Quality assurance standards that address security considerations
   - Access control for operations as well as data (e.g., any of the menu systems which preclude access to the operating system).
   - Unambiguous user identification (ID) and authentication (e.g., personal profiles and hand−held password generators).
   - Protection of executable code (e.g., flags to show that certain object modules are "production" or "installed" and thus apply strict access control that would prevent unauthorized modification—as found in configuration control systems).
   - Security logging (e.g., logging failed file-open attempts and logon password violations).
   - Assigning a security administrator to each enterprise.
   - Data encryption.
   - Operational support tools for verifying the state and effectiveness of security measures (e.g., audit tools).
   - Independent audits of system security by people not directly involved in programming or system management of the audited system.
   - Hazard analysis evaluating threats to safety from different malfunctions and breaches of security (e.g., consequences of tampering with patient data in hospitals).

2. Take specific short-term actions now:
   - Develop security policies for your organization before there is a problem.
   - Form and train computer emergency response teams before a crisis to respond to security violations or attacks.
   - Use the Orange Book's (TCSEC, from the National Computer Security Center's Rainbow series) C2 and B1 criteria to define guidelines on security.
   - Improve software systems development by applying better quality-assurance methods.

- Contribute to voluntary industry groups developing modern security standards and implement those standards in commercial software.
- Make effective security the default in software and hardware. (Make the user explicitly disable security instead of having to enable it.)

3. Learn and teach about security:

- Build a repository of incident data.
- Foster education in engineering secure systems, both by encouraging universities to provide postgraduate training in security and by urging industry to include security training as part of software engineering projects.
- Teach beginners about security and ethics in computer usage and programming (e.g., the National Center for Supercomputing Applications is working on a research and development project to study beliefs, attitudes and behavior about ethical issues in computing in grade and high schools, colleges, and universities).

4. Clarify export control criteria and set up a forum for arbitration. (Hardware and software vendors have been complaining for years that the arbitrary imposition of severe export restrictions hampers American competitiveness in overseas markets without materially helping national security.)

5. Fund and pursue needed research in such areas as:

- **Security modularity.** The effects on security of combining modules with known security properties.
- **Security policy models.** More subtle requirements, such as integrity and availability, still are not easily represented by control structures.
- **Cost estimation.** There should be better ways of measuring the costs and benefits of security mechanisms in particular applications.
- **New technology.** Networking, in particular, leads to greater complexity (e.g., how to connect "mutually suspicious organizations").
- **Quality assurance for security:** How to measure effectiveness.
- **Modeling tools.** Standards for graphical representations of security relationships analogous to the diagrams used in functional decomposition and object-oriented methodologies for program design.
- **Automated procedures.** Audit and monitoring tools for the data center management team.
- **Nonrepudiation.** Combining the need for detailed records of user actions with the values of privacy.
- **Resource control.** How to ensure that proprietary software and data are used legitimately (e.g., preventing more than the licensed number of users from accessing a system, preventing software theft).
- **Security perimeters.** How to reconcile the desire for network interconnection with limitations due to security requirements ("If, for example, a network permits mail but not directory services... less mail may be sent because no capability exists to look up the address of a recipient").

Chapter 2 of the NRC report, "Concepts of Information Security," is a 25-page primer on information systems security that could be handed to any manager who needs to be filled in on why you propose to spend so much money protecting the

computer systems. The authors cover the fundamental aspects of IS (confidentiality, integrity, and availability); management controls (individual accountability, auditing, and separation of duties); risks (probabilities of attack or damage) and vulnerabilities (weak points); and privacy issues. In Appendix 2.2, the authors report an informal survey in April 1989 of 30 private companies in a variety of fields. The consensus among those polled included these basic standards for IS security (show these to your upper management if necessary):

- Unique IDs, block access after a maximum number of incorrect logon attempts, show last successful access at logon time, make passwords and IDs expire.
- Disallow embedded passwords during logon, make passwords invisible during entry, force minimum length (6), store passwords encrypted, scan proposed passwords to eliminate easy words.
- Permit strict control over file access.
- Detect and interdict viruses, certify software as virus-free, provide data encryption, overwrite deleted files to prevent recovery, force tight binding of production data to production programs.
- Automated time-out for inactive sessions, unique identification of terminals and workstations during logon.
- Network security monitoring, modem-locking, callback, automatic data encryption during transmission.
- Audit trails including security violations.
- Generally applicable security standards that could be used by vendors and users to evaluate different equipment and software for specific environments.

Twenty years later, focus among information assurance experts has shifted beyond the technical to emphasize organizational controls. For example, the 2003 survey of members of the Information Systems Security Association included these IS function practices by the respondents:

- Access controls, 73%
- Written information security policy, 72%
- Compliance with existing laws and regulations, 66%
- Creation of organization and process to implement policy, 59%
- Awareness and training program, 57%
- Regular monitoring, reviewing and auditing, 57%
- Business continuity planning, 57%
- Risk assessment and risk management, 56%

In 2007, Gary S. Miliefsky, noted entrepreneur, and founding member of the U.S. Department of Homeland Security proposed seven priorities for corporate information security:

1. Roll out corporate security policies.
2. Deliver corporate security awareness and training.
3. Run frequent information security self-assessments.
4. Perform regulatory compliance self-assessments.

5. Deploy corporate-wide encryption.

6. Value, protect, track and manage all corporate assets.

7. Test business continuity and disaster recovery planning.[18]

The Computer Security Division of the Information Technology Laboratory at the National Institute of Standards and Technology issued a draft reference model that included these "programmatic, integration, and system security activities that are typically a part of an information security program":

Program Security Activities

Annual and Quarterly Review and Reporting of Information Security Program
Asset Inventory
Awareness and Specialized Security Training
Continuity of Operations
Incident Response
Periodic Testing and Evaluation
Plan of Action and Milestones
Policies and Procedures
Risk Management

Integration Activities

Business Risk
Capital Planning and Investment Control (CPIC)
Configuration Management
Enterprise Architecture (EA)
Environmental Protection
Human Resources
Personnel Security
Physical Security
Privacy
Records Management
Strategic Plan
System Development Life Cycle (SDLC)

System Security Activities

Categorize the Information System
Select Security Controls
Supplement Security Controls
Document Security Controls
Implement Security Controls
Assess Security Controls
Authorize the Information System
Monitor Security Controls[19]

1.3.3 InfraGard.[20] InfraGard is a nationwide program in the United States that brings together representatives from information technology departments in industry and academia for information sharing and analysis, especially to help protect critical infrastructure against cyberattacks and also to support the Federal Bureau of Investigation (FBI) in its cybercrime investigations and education projects. [21]

The organization started in the Cleveland Field Office of the FBI in 1996 and expanded rapidly until there are now over 11,000 members in over 40 chapters. Joining InfraGard is easy and free for U.S. citizens residing in the United States. Using the Web site (www.infragard.net/chapters/index.php?mn=3), you can locate a nearby local chapter and contact your chapter officers. You can get application forms online and then send them in to the FBI liaison officer for that chapter to be vetted for admission. The FBI conducts a background check to ensure that all members are likely to be trustworthy to participate in confidential discussions of threats and vulnerabilities. Chapters usually conduct regular local meetings and organize list-servers for exchange of information among members. Many have newsletters as well.

1.4 RECENT DEVELOPMENTS. In the years since the Fourth Edition of this *Handbook* was published (in 2002), one of the key developments has been the dramatic increase in availability of inexpensive portable data storage devices. At the time of writing (2008), flash drives the size of a lipstick or smaller are available online with capacities in the dozens of megabytes for a few dollars and capacities in the gigabytes for little more. Such devices are available in a wide range of concealable formats, such as pens, music players, and even watches. Digital cameras use storage cards that can be used for data transfers; mobile phones include cameras and recording capabilities. Controlling data leakage through unauthorized connection of such devices has become a significant problem for security managers. Systems for restricting connection of devices and controlling data transfers to such storage media are spreading through government and corporate environments.

Another issue that increasingly concerns security managers is the protection of personally identifiable information (PII) from customers or data subjects. Many organizations, including government agencies, banks, and universities, have suffered serious damage from loss of control over PII and the risks of identity theft resulting from exposure of such sensitive data. Legislators are responding to public concern by increasing legal requirements for protection of PII. The use of encryption on mobile data systems such as laptop computers, personal digital assistants (PDAs), mobile phones, and integrated systems that combine many functions (e.g., BlackBerries) has become a necessity. A consequence of the growing interconnectivity of storage and communications devices is that corporate networks are no longer insulated from less secure systems. Users who connect poorly protected laptop computers or other devices to public networks, such as hotel-supplied ISPs or wireless access points in coffee shops, may return to their home offices with malware-infected systems that contaminate the entire network. Security managers are increasingly turning to integrated systems for controlling connectivity via virtual private networks and supervisory software that monitors and restricts unauthorized connections, software installations, and downloads.

A most formidable new threat lies in the international operations of mafia-like rings of computer criminals. Once such collusion stole over 41 million credit and debit card records from the giant retailer, TJX. According to information released by the U.S. Department of Justice on August 5, 208, the ring consisted of three Americans, one

Estonian, three Ukrainians, one from Belarus, two from the Peoples Republic of China, and one known only by a network "handle". The eleven were charged with conspiracy, computer intrusion, fraud, and identity theft, perpetrated by "war driving" and hacking into wireless computer networks.[22] For more about wireless network security, see Chapter 33 in this Handbook.

1.5 ONGOING MISSION FOR INFORMATION SYSTEM SECURITY.

There is no end in sight to the continuing proliferation of Internet nodes, to the variety of applications, to the number and value of online transactions, and, in fact, to the rapid integration of computers into virtually every facet of our existence. Nor will there be any restrictions as to time or place. With 24/7/365, always-on operation, and with global expansion even to relatively undeveloped lands, both the beneficial effects and the security violations can be expected to grow apace.

Convergence, which implies computers, televisions, cell phones, and other means of communication combined in one unit, together with continued growth of information technology, will lead to unexpected security risks. Distributed denial-of-service (DDoS) attacks, copyright infringement, child pornography, fraud, and theft of identity are all ongoing security threats. So far, no perfect defensive measures have been developed. This *Handbook* provides a foundation for understanding and blunting both existing vulnerabilities and new threats that inevitably will arise in the future.

Certainly, no one but the perpetrators could have foreseen the use of human-guided missiles to attack the World Trade Center. Besides its symbolic significance, the great concentration of resources within the WTC increased its attractiveness as a target. After 9/11, the importance of physical safety of personnel has become the dominant security issue, with disaster recovery of secondary, but still great, concern. This *Handbook* cannot foresee all possible future emergencies, but it does prescribe some preventive measures, and it does recommend procedures for mitigation and remediation.

1.6 NOTES

1. Many technical specialists use the term "security" to refer to logical access controls. A glance at the contents pages of this volume shows the much broader scope of information system security.
2. For further details, see, for example, www.cs.uiowa.edu/∼jones/cards.
3. See inventors.about.com/library/weekly/aa062398.htm.
4. It is notable that the IBM 1401 computer was so named because the initial model had 1,400 bytes of main memory. It was not long before memory size was raised to 8 kilobytes and then later to as much as 32 kilobytes. In 1980, the Series III minicomputer from Hewlett-Packard doubled its maximum memory from 1 megabyte to 2 megabytes at a cost of $64,000 (about $200,000 in 2008 dollars). This compares with today's personal computers, typically equipped with no less than 512 megabytes and often a gigabyte or more.
5. The term "dumb" was used because the terminal had no internal storage or processing capability. It could only receive and display characters and accept and transmit keystrokes. Both the received characters and the transmitted ones were displayed on a cathode ray tube (CRT) much like a pre–color television screen. Consequently, these were also called "glass" terminals.
6. "Multiprocessing," "multiprogramming," and "multitasking" are terms that are used almost interchangeably today. Originally, "multitasking" implied that several

modules or subroutines of a single program could execute together. "Multiprogramming" was designed to execute several different programs, and their subroutines, concurrently. "Multiprocessing" most often meant that two or more computers worked together to speed program execution by providing more resources.

7. Also known as ARPAnet and Arpanet.
8. First published 1975. Reissued by Mass Market Paperbacks in May 1990.
9. "Fixed," in contrast with the removable disk packs common in large data centers.
10. See standards.ieee.org/getieee802/802.3.html.
11. National Research Council, (1991). *Computers at Risk: Safe Computing in the Information Age* (Washington, DC: National Academy Press, 1991). Available as a searchable openbook at www.nap.edu/books/0309043883/html/index.html.
12. G. H. Nibaldi, "Proposed Technical Evaluation Criteria for Trusted Computer Systems," Publication M79-225 (Bedford, MA: MITRE Corporation, 1979).
13. For access to all the Rainbow Series documents, see www.fas.org/irp/ nsa/rainbow.htm.
14. The CCEVS Web site has extensive documentation; see www.niap-ccevs.org/cc-scheme/.
15. This section is reprinted with slight modifications by permission of the author from the original manuscript for M. E. Kabay, *The NCSA Guide to Enterprise Security: Protecting Information Assets* (New York: McGraw-Hill, 1996), Chapter 1, pp. 2–5.
16. National Research Council, *Computers at Risk*.
17. www.sunnyday.mit.edu/papers/therac.pdf.
18. G. S. Miliefsky, "The 7 Best Practices for Network Security in 2007," *Network World*, January 17, 2007; www.networkworld.com/columnists/2007/011707miliefsky.html?t51hb.
19. E. Chew, K. Stine, and M. Swanson, (2007). "Information Security Reference Data Model (DRAFT)," NIST Special Publication 800-110 (Draft); http://csrc.nist.gov/publications/drafts/sp800-110/Draft-SP800-110.pdf.
20. M. E. Kabay, "InfraGard Is Not a Deodorant," *Network World*, September 8, 2005; www.networkworld.com/newsletters/sec/2005/0905sec2.html.
21. www.infragard.net/about.php?mn=1&sm=1-0.
22. www.usdoj.gov News release of August 5, 2008.