Policy Considerations for Data Networks

Willis H. Ware The RAND Corporation

ABSTRACT: For the purposes of the present discussion, a *network* is the telecommunications arrangement for transporting digital traffic from place to place, and the boundary is considered to be at the interface between it and subscriber systems. Both homogeneous networks and heterogeneous networks are described and discussed in detail, and security concerns and threats are presented. The focus of concern is the *heterogeneous network* (e.g., the Internet); which provides data communication services to all interested and qualifying parties. Such a network behaves in the spirit of a common carrier but is not today legally characterized as such.

Security of a network is an essential topic not only in behalf of safe, reliable, and proper functioning of the ensemble, but also because lack of it can lead to circumstances in which subscriber disappointment and rejection will threaten the goals and hopes for network acceptance and success. A review of the development of computer security and network security as national policy and standards issues is presented.

With regard to *security-related policy* that the network authority may wish (or may be obligated) to adopt, there are several categories of interest. Among them are the following:

- Policy to protect the network against specified threats and to ensure its safe, reliable, and proper functioning.
- Policy to control the nature, details, and conditions under which pro bono or for-fee service offerings to the network at large can be provided, or may be required to provide as a condition of being allowed to connect (e.g., electronic mail).
- Policy to protect subscriber traffic while it is in the custody of the network in which it is stored or while in transit through it.
- Policy to govern subscriber behavior in regard to its interaction with the network, including usage.
- Policy to protect each subscriber against improper or malicious behavior by one or more other subscribers.

With regard to networks that are funded by or related to federal interests, any or all of these may become binding through law or regulation. Other federal policies, such as limitations on the use or export of cryptography, can intrude on freedom of choice for network policy and operators.

1. Introduction

This article is both an overview of policy concerns for data networks and a discussion of the security issues that underlie them. It defines certain relevant terms, then establishes a partial taxonomy of data networks, characterizes them by various attributes, reviews security as a federal policy and standards issue, and finally suggests various policies that will be desirable, if not essential, to protect the networks that the United States will create as it moves into the age of a National Information Infrastructure (NII). This article does not recommend specific policy positions or offer draft policy statements.[1]

1.1. Terminology

Throughout this article, jargon is minimized and acronyms are expanded at least once. Intuitive similes and examples are used when feasible to facilitate understanding of constructs. A few particular words are used consistently with a specific orientation and content.

- *Recipient* (of a message or communications traffic): An individual, a computer system, a software process acting on its own, a software process acting in behalf of a user, or a network. Context will make clear which one or more of these choices is intended, but certainly, in the long run, the word must acquire a much broader meaning than a simple reference to a person.
- *Subscriber (system)*: A generic phrase for a system connected to a network, including its hardware, software, users, operational policies, capabilities, services and security controls, the nature of data that will appear in the network traffic, and so forth. It is convenient to use this term when discussing certain networks that offer various services to interested parties. As used

[1] This article is an expanded, rewritten version of one that was a focus presentation to a National Science Foundation–sponsored workshop on network policy at Amelia Island Resort, FL, January 1993. Professor Lance Hoffman, George Washington University, was the organizer and general chairman of the event.

in the network context, it is a generalization of its use in such common phrases as *magazine subscriber* or *telephone subscriber*.

- *Host*: An often-used term that might have been chosen but it tends to suggest only the hardware-software complex and to exclude other things that are of concern in discussing network connections (e.g., operational policies and services that a subscriber system might offer to the network at large).
- *Subscriber*: Used in the context of telephony, it will have the usual meaning associated with telephone usage.

The common terms *computer* and *computer system* are used where appropriate and when no special characteristics or other connotation is implied. When appropriate, *telecommunications* is used instead of *communications* to reinforce the point that actual physical arrangements to move data electronically are implied.

1.2. Definitions

There are three terms, often used imprecisely and sometimes synonymously, that must be carefully distinguished, primarily because different approaches are relevant to accommodating each but also to maintain a clear structuring of the issue.

In discussing computers or networks containing them, one needs a generic term to refer to the digitally represented material with which they deal. In a strict mathematical sense, *information* should be that term but, in common usage, *data* is the word most often used. Sometimes a distinction is made between the two in the sense that data is the raw stuff from which facts and then information are derived; but for the purpose of this discussion, such caution is not necessary. *Data* will be the preferred term to connote the digital material stored within a computer or the digital traffic in transit through a telecommunications network; in a few instances, *information* will have to be used in its strict mathematical sense.

Of direct concern to the clarity of the discussion, however, are the three terms, *confidentiality, security, and privacy*. *Integrity* is also relevant as an adjunct to security considerations.

Confidentiality is defined as a status accorded to data which indicates that, for some reason, the data is sensitive and must be protected as such. *Protection* means not only safeguarding the data against destruction or unauthorized change but also limiting access only to authorized consumers or users. *Authorized users* may be established by law, by regulation, by professional custom, by organizational policy, simply by established historical uses, or by agreement among the members of some organized community (e.g., insurance companies in the case of the Medical Insurance Bureau).

Security is the totality of safeguards present within and around a computer system or a network to ensure (1) the safeguarding of its physical and data assets and, if need be, its people and users also, (2) maintenance of confidentiality of the data and assurance of its integrity, and (3) controlled access to the system and its assets. Safeguards might include some or all of the following: technical (software or hardware or both), procedural, administrative, management, physical, and personnel. An essential feature of a system of safeguards is the access control—to data, a system, a network, resources, and software processes. In turn, this may imply control of what is done with or to data after access is achieved.

Privacy, in the data context, means the use of personal information for prescribed activities. [2] *Prescribed* can mean by law, by regulation, by organizational policy, by professional custom, by established historical uses, or by members of some organized community (e.g., the use of names and addresses for mailing lists by the Direct Marketing Association). The general intent of privacy law is to protect individuals—not legal bodies such as corporations in the United States—against harm, unwarranted intrusion, or possibly serious damage. If personal information is used for a purpose other than that for which it was initially collected, it often is said to be a *secondary use.*

Finally, for completeness, *integrity* is the knowledge that a given body of data, a system, an individual, a network, a message in transit through a network, or the like has the properties that were a priori expected of it. Note that such a definition does not require absolute accuracy, freedom from errors, or complete specification of the entity in question. It requires only that whatever something was believed to be before the fact is indeed what it proves to be after examination. In some contexts, *integrity* is taken to mean assurance against unauthorized change. Because security safeguards control access, there is a clear interplay between the interests of security and those of integrity. Some safeguards—possibly all, in some circumstances—will contribute to both end goals. Sometimes integrity is equated to data accuracy.

The immediate interest in this article is only that of security of networks and pertinent implementing policies. Confidentiality and privacy are collateral issues of concern to subscriber systems, end users, organizations, and individuals who utilize the capability of a net; these issues are discussed elsewhere.[3]

[2] There are other dimensions of privacy that are not of interest here (e.g., physical privacy or psychological privacy). For a recent discussion of information privacy, see [16].

[3] For a historical progression of the salient publications on privacy, see [1, 19, 12, 11, 4, 16].

2. Networks

An equally important detail is the definition of *network*. Use of the term is quite variable, especially by computer people who invoke it in many manifestations (e.g., local area networks (LANs), wide area networks (WANs), computer networks, networks of networks, corporate networks). To the communicator, however, a network is a transportation mechanism for the movement of electronic traffic. As such, its obligation is to move the traffic from originators to recipients as determined by the originator. Moreover, it must do so in a timely fashion and must faithfully deliver to the recipient the same information that was supplied by the originator but not necessarily in the same representation (e.g., a message might enter as a digital stream of characters but be delivered as a printed page). Ideally, a telecommunications network will also handle traffic without perturbation en route. For networks dealing with signals representing data in analog fashion, *perturbation* equates to worsening the signal-to-noise ratio; for those representing data in digital form, it will be the bit-error rate or, in general, unintended modification of the data content.

Put another way, a network—data or telecommunications—is to provide connectivity among subscribers.

For the general discussion of networks, it is convenient to allow *network* to be a somewhat loose term. It will be easier to talk about such things as *network services* and to carry on general exposition. When policy becomes the topic of discussion, the term will have to be restricted, but it will be so noted at that time.

2.1. The Network Boundary

It will prove useful to use the so-called public-switched-network (PSN) as a reference analog. The PSN has been around a long time; it can offer some lessons, points of view, and policy concerns. It is the collection of facilities that have been assembled over many decades to provide what is loosely called telephone service, "loosely" because much more than simple voice conversations are now supported by the PSN.

But what is considered to be the boundary of a network, especially in the context of a data network?

Again, usage is variable. Sometimes subscriber systems or computer-related components are considered to be within the boundary of the network but sometimes not. The PSN points the way to a possible position. Today the PSN stops at the plug-on-the-wall, although it was not always so. Prior to the Carterfone court decision, which resolved the foreign attachments issue, the telephone company held responsibility for and administered what was connected to the plug—telephone handsets, for the most part. Today a wide variety of devices can be connected to the telephone network as long as prescribed interface standards are met. Today the PSN does not conceptually include the devices that are connected to it.

The LAN industry holds the same position. It vends telecommunication arrangements that permit devices to have connectivity with one another, but its purview excludes that which is connected, except to stipulate the connection interface details and generally to require software in the connected mechanisms to facilitate the proper operation of the interface.

Unlike a LAN, the PSN does not have to reach beyond the plug to require modification of connected devices, but why should it have to? After all, devices are designed specifically to attach to the PSN and must be certified to be in compliance with federal regulations. In contrast, most computers have been designed as self-contained, stand-alone devices. Connecting them to the PSN or to a LAN or to another form of telecommunications arrangement is achieved with specially designed components (modems, communication controllers, drop-in cards) or with specialized software to straddle the difference between the detailed internal workings of the computer and the workings of the telecommunications system.

Computer people frequently use "network" in a very collective sense. For example, a company will install a "computer network"; what that really means is that a collection of computer systems and components, a collection of telecommunication systems and components, and specialized or commercial software, or both, are blended into an ensemble that is conveniently called "a network."

Such a collective use of the term is acceptable for a situation that operates under a single, cohesive management and policy. From conception, the entire thing operates on a common set of rules and policies, extending out to the end users working at terminals. Such a global implication is, however, not workable for the situation in which a telecommunications arrangement is a provider of inter-user connectivity for any organization that can meet the interface requirements, can supply the necessary investment and operational funding, and will agree to abide by the rules that govern connection to and use of (what in effect is) a common-carrier arrangement. The many subscribers need not have, and in general will not have, an overall uniform policy or even hardware/software commonality; nor will they have uniform needs and requirements for giving their computer systems electronic outreach. Similarly, telephone subscribers buy and use a wide variety of instruments to suit individual needs or desires. As long as certified devices are used, there is no other common authority controlling their behavior except law. The PSN cannot dictate what is connected to it, as long as interface standards are met. This obligation is one of the attributes implied by the legal status of *common carrier.*

In the case of the well-known Internet, *the net*, as it is colloquially called, is only the telecommunications part in a precise sense, but the term is often extended to include everything that is connected to it.[4] In truth there are computers within Internet (e.g., packet switches to provide routing and addressing information, machines in the network control center), and there are computers attached to the Internet (e.g., subscriber systems). Thus, when the term "network" is used to mean a telecommunications arrangement to which subscriber systems can attach, it must be understood to also include such computers or switches as may be necessary to make the telecommunications system function and allow it to be operationally controlled and maintained. [EN1][5]

2.2. *Other Network Attributes*

Quite aside from the question of placing the boundary, there are other ways in which networks can be characterized. Among them are the following:

- Sometimes a network will be described by the technology used (e.g., microwave vs. fiber optic vs. satellite). Some current vendors of telephone services speak of "an all fiber network," meaning that its links are fiber optic technology.
- Sometimes a network will be called *circuit switched* vs. *packet switched.* The first is the traditional telecommunications arrangement in which, at one time in history, two subscribers did, in fact, have an end-to-end direct connection—so to speak, a continuous pair of copper wires from one to the other. In current telephone networks, such things no longer exist because digital controls allow rapid re-routing of traffic to accommodate network loading or outages or physical movement of a subscriber. It is all achieved without interruption of the connection and it is transparent—that is, not visible—to the user. Such circuits are commonly referred to as *virtual circuits.*

[4] For example, see [10]. In it, Internet is not only the telecommunications network, but everything connected to it. Hence, its table of contents includes such things as Security Consciousness (of subscriber systems), Politics and the Internet, Legal Implications (of traffic on Internet), and Network Ethics. Other sections deal with various net-wide services, such as electronic mail and file transfer.

[5] References in this format are to end notes just prior to the bibliography.

Packet switching is particularly amenable to computer traffic because it breaks everything into chunks—packets—of prescribed length and individually numbers, identifies, addresses, routes, and delivers each one. Packet networks can be assembled from the circuit-oriented telephone plant, but while the end-points of the leased PSN circuits are fixed with regard to the packet-net structure, the precise routing between any two end-points might vary as the telecommunications operations require. A packet network can be, and is, assembled from the virtual circuits of the PSN.

- A network might be referred to as public or private. The distinction might be made on the basis of the community of users served, the basis of services offered, or the funding that supports the network. Many combinations can exist; for example, a blend of public and private funding but publicly available (Internet); privately funded but publicly available (PSN); privately funded with private subscribers (corporations).

3. Homogeneous vs. Heterogeneous

Even a casual acquaintance with computer-oriented networks makes clear that there is something fundamentally different between the ones that serve business and industry (e.g., airline reservations, insurance company systems, credit card transactions) and the Arpanet, which served academic and research organizations, or its successor, the Internet. The differences will result in different policy concerns and, hence, must be explored.

In the taxonomy of data networks there are two broad classes commonly referred to as *homogeneous networks and heterogeneous networks*. [EN2] The difference between them is reasonably distinct, although the edges are fuzzy and to some extent overlap. Some nets will have features of both; the progress of technology is making a careful distinction even more difficult. However, there are characteristics typically associated with each, and they will be elaborated later. In brief, the homogeneous net is technically uniform with respect to hardware, software, and architecture; it operates under a single policy authority. In contrast, the heterogeneous network is technically diverse with respect to hardware, software, and architecture and operates as a cohesive entity under the policies of many authorities.

For either kind of network, it is convenient to regard the telecommunications component as administered by a single authority, but in reality, this may never be true. Independent authorities can preside over different parts of the overall PSN from which facilities are leased; a consortium might collectively function as a single authority to establish policy, or third party resellers might exist. Whatever the governance, there will be a minimum and uniform set of policy and technical requirements across the whole PSN; otherwise, administration of the network and relations with subscribers would be very difficult.

The term “open network” might have been used instead of heterogeneous network, but “open” is already well established as “open technology” or “open systems.” Moreover, any network can, in principle, be implemented from open technology and conceivably as an open system, so it seems wise not to risk the confusion of using the word “open” in different contexts.

The homogeneous network and the heterogeneous network will each be treated separately below, from defining the characteristics through discussing security threats. While the two kinds of networks have security vulnerabilities and threats that can be similar, it is useful to discuss the whole scope of network security and help put things into perspective. The common threads are the state-of-the-art of computer security, especially the technical part, and policy issues. Keep in mind that the level of discourse is that of general policy and security matters, not the technical details of how the network is implemented.

4. Homogeneous Network Characteristics

Consider first the class of homogeneous networks. It will not only lay a foundation for appreciating what they are, but also start our consideration of security concerns. Importantly, it will also exclude a whole class of networks from the present discussion.

The term implies many things. Although there is inevitably variability among various homogeneous networks, historically a homogeneous network would:

- Operate under a policy that derives from a single authority, a very important aspect. The authority extends to all data-communication arrangements per se, to all users, for all computer systems throughout the net—in brief, to the entire thing.

- Operate under a common policy and management structure that flows from the single authority, again an important observation.

- Be typically designed as an entity and implemented as an entity, although possibly incrementally in phases.

- Not permit arbitrary changes on the part of system users. All change is coordinated, approved, and controlled by the management authority.

- Have a community of users who are generally constrained in regard to what each can do. The capabilities provided to user segments in the overall community will depend on such things as job assignment, position in the management hierarchy, and perhaps physical location. Access to system resources or data is generally imposed by system security controls—possibly technical ones, possibly procedural ones (e.g., physical limitation to accessing terminals), perhaps terminal characteristics (e.g., no graphics activity from a terminal that the system knows as a text-only terminal).

- Have commonly a uniform technical base, hardware, software, and architecture; or at least a base that is technically compatible throughout (e.g., plug-compatible mainframes or disk drives).

- Frequently have a few large computers or computer systems concentrated in one or a very few centers, which are interconnected and collectively share and support the system workload.

- Generally, but not necessarily, have a set of terminals locally or remotely associated with each computer system; but also have a network connection to other (if any) computer systems and perhaps to remote terminals.

- Usually have dedicated point-point telecommunications, although these might operate in both a high-speed continuous data stream and packet mode.

- Often have dial-in or dial-out arrangements for some users—which might be other systems or simply individuals with mobile job assignments or off-site obligations.

- Usually have a stable connectivity among the participating components and systems; e.g., impromptu arbitrary connections usually would not occur. The analog of a telephone subscriber calling any other subscriber at an arbitrary time for an arbitrary purpose does not exist. The connectivity is generally known at design time although some flexibility in connectivity may occur and certainly changes can take place.

- Usually have a traffic pattern among systems that was reasonably fixed or known or prescribed by overall system requirements. The patterns tend to be slowly changing.

While such characteristics typify a homogeneous network, the progress of technology dilutes the notion of highly centralized installations, which, however, continue to exist in many major network applications. Such advances as the spreading use of LANs of all kinds, the aggregation of LANs into WANs, distributed databases resident on multiple servers geographically dispersed, clusters of workstations and terminals serviced from nearby servers as well as telecommunications links to large centers, and distributed regional computer systems instead of one or a few large centers all act to yield an implementation architecture and structure that is more diverse than the conceptual overview.

A salient characteristic of a homogeneous network is that it usually supports a community of users who are all contributing to common organizational goals; it is specified, designed, and implemented as such. If connectivity to other extra-organizational networks becomes necessary, it will be a special arrangement and not all users are likely to have such privileges. "Homogeneous" applies to the entire ensemble—the telecommunications arrangement, everything connected to it, management and policy, control of users.

The class of homogeneous networks is not a topic of interest in the present discussion. If one of them supports a private-sector entity, its operational policy will be determined by the policies of the entity and relevant law. If one of them supports a government entity, its operational policy will be determined by the agency policies as influenced or established by relevant law or executive order. For homogeneous networks, there is no doubt as to the nature and source of policy constraining and governing its operation.

Also note that the telecommunications infrastructure of a homogeneous data network can, and commonly does, support other organizational needs than data transmission (e.g., voice, facsimile, video conferencing). Such extensions are growing as the capability of commercial telephone networks continues to increase dramatically and the economics of shared use becomes attractive.

4.1. Homogeneous Data Networks vs. the PSN

Before divestiture AT&T would have been the classic example of a homogeneous network, although limited largely to voice. Although AT&T per se did not own, and therefore did not directly manage, all telephone assets in the country, it consolidated many small systems into one entity that did own a major portion of the national system. Its influence, market strength, and research advances encouraged, if not forced, all other purveyors of telephone services to adopt AT&T standards and switching practices. In effect, for all practical purposes, the United States had a homogeneous network for telephone services dominated by and technically controlled by AT&T as a central management authority. As the principal force in United States telephony for decades, AT&T also worked out and standardized the interface connectivity to other countries.

In the telephone network prior to divestiture but after resolution of the foreign-attachments dispute, subscriber-terminals were initially telephone sets and later consumer electronics devices; the boundary of the telecommunications component of the overall network stopped at the wall plug (or jack). The network provided connectivity among nationwide subscribers who input their connection, billing, and other requirements to the network initially by voice, later by dial, and currently by touch-tone keypad.

Prior to the emergence of such consumer devices as facsimile machines and personal computers, the single service was voice supplemented by necessary support (e.g., directory service). Today, the range of services that can be called upon by a contemporary telephone subscriber goes far beyond simply establishing connectivity. As a result of the digitization of the communications per se and of putting operational control of the network in the hands of compatible computers from many vendors, all sorts of ancillary features have become available (e.g., call waiting/forwarding, speed dialing, caller ID, call tracing, last-call return, specialized billing arrangements, special ringing; and in the long run, more elaborate services such as records maintenance for the subscriber, or database services).

Post-divestiture, of course, the situation has became somewhat more complex conceptually and in reality. The extensive standards already in place as a result of AT&T's long domination and central authority made the technical transition reasonably smooth. The PSN is well on its way to becoming a different kind of network.

4.2. Examples of Homogeneous Data Networks

There are many well-known homogeneous data networks not commonly thought of by that term. They are widely deployed and generally known, serving business, government, and research. While functionally meeting the general outline of the characteristics above, some of these have changed markedly from the historical implementation as a consequence of new technology and system upgrades. Most notably, the centralized mainframe single-database arrangement has given way to regionalized distributed database arrangements often hosted on clusters of workstations sharing a server via a LAN.

Among them are the following:

SABRE and APOLLO—airline reservation systems.

SSA network—a single computer-center, supporting terminal service and database access for 30,000 terminals in 1300 field offices. It is scheduled to be upgraded by installing workstations and LANs in field offices with some of the master database on local servers.

IRS network—multiple regional sites, partially networked, now scheduled for upgrade and will serve all field offices. Also scheduled for a major upgrade, probably similar to that of the SSA but with regional major processing centers as well.

Healthcare networks—several public ones (run by the Health Care Financing Administration) and private ones (e.g., Electronic Data Systems, Blue Cross/Blue Shield) supporting Medicare, corresponding state programs, and private health care delivery. The individual networks are usually not interconnected although individual ones function as networks.

Bank card and credit card networks—nationwide and international, generally with a few regional sites and permanent or dial-up telephone links to merchants worldwide and to other network sites.

Insurance company networks—usually one or two computer centers supporting terminals in agent offices.

Computer vendor networks—usually from one to many computer centers supporting sales offices, corporate and other worldwide needs; sometimes regional centers as well.

Automobile company networks—for order, inventory, and distribution of repair parts, centralized records for vehicles of the vendor, engineering design, production control, manufacturing support.

ATM networks—to provide very limited financial transaction capability to users, but can span many states.

Financial industry networks—can be worldwide directly or via other cooperating networks; many corporations, banks, and other financial institutions participate.

LANs—although the technical base of connected terminals or workstations, might be compatible rather than uniform.

4.3. Bulletin Board and Database Services

Database and other public services—such as Medline, Compuserve, LEXIS/NEXIS—are really a variation of the homogeneous net structure. They differ from the prior characterization in some details; for instance, they usually do not have a uniform technical base at the end-user level, especially when the service is offered to the public at large. To accommodate this variability, they include flexibility to support many types of terminals or computer types or systems at the subscriber level.

They consist usually of one or a few computer centers that provide user access over dial-in arrangements, although major users might have permanent connections; they can have regional centers as well. User capabilities and privileges are rigidly defined and limited, although they can be very extensive. Access is usually a fee-for-service arrangement; sometimes, however, a database will be publicly available (e.g., library services in a large urban area). The offered services may go far beyond database service (e.g., electronic mail and discussion groups on CompuServe). This class of services is likely to be as close to homogeneous as a public-subscriber data network will become. Even then, the analogy is not precise in some ways because the telecommunications part may not be under the authority and control of the organization offering the service, but rather will be one chosen by the subscriber. The individual normally will use the PSN or one of the value-added-services (e.g., Telnet) that provide data transmission services with telecommunications arrangements based on facilities leased from the PSN.

4.4. Major Security Considerations

Typical security concerns for a homogeneous network include the following:

- Physical vulnerability of the system resources. The computers and the communications components and the people must be protected against deliberate attack or sabotage intended to physically damage or destroy them, and against accidental damage (e.g., bombing, fire, burst pipes, severed wiring).
- Personnel control and trustworthiness. Security risks attach to casual unescorted visitors, to maintenance personnel of unknown trustworthiness, to concentration of excessive system privilege and capability in one person (e.g., an unsupervised operator who has complete access to the system during the graveyard shift).
- Security safeguards for the computer systems embedded within the system. This, in the large, is a topic separable from that of the communications vulnerabilities. It is a reasonably well-understood subject; security safeguards and precautions focus on the software components, particularly the operating system and the database system.
- Communications security (COMSEC) to protect the traffic in transit through the telecommunication parts of the network against interception. An additional concern in some circumstances can be the concealment of traffic-flow patterns. These are well-understood subjects, and appropriate security measures are available (e.g., encryption devices, hardened cable runs, sophisticated modulation schemes such as spread spectrum).
- Dial-in connections (or ports) to the network. These are a common entry point for network penetrators who methodically search telephone numbers

to identify modem connections. Safeguards exist (e.g., dial-back arrangements).

- Computers embedded in the network, not to support individual end users but rather to control, monitor, and operate the network. Sometimes they are wholly within the telecommunications portion of the network (e.g., circuit or packet switches), sometimes in a network control center. Such switches generally have dedicated tasking but there can be extensive collateral capability to facilitate maintenance, especially remote maintenance. Switch software is generally not accessible to programmers, but, in well-controlled circumstances, only for trouble shooting, repair, software upgrades; may be operated unattended or maintained remotely; and can often receive new software downloaded via the network from the control center.

- Protection of remote maintenance arrangements. Maintenance dial-in ports can be particularly dangerous since maintenance will usually require extensive access to the system software, even to or with specialized maintenance software. Hence, an unauthorized entrant through a maintenance port will often have greater run of a system than through a user port.

- Protection against crackers. This has already been indirectly noted in the discussion of dial-in connections for users or for maintenance. More sophisticated attacks are possible by a penetrator (e.g., actual wire tapping to acquire passwords or other operational details about the system), especially if the miscreant is already an authorized user of the network (e.g., injecting worms or viruses), or if he or she has acquired user status by some subterfuge (e.g., stealing a password, being given a password by an accomplice).

- Protection against insider actions. These are unauthorized actions by authorized users of the system. Depending upon the level of access by the authorized user, the damage could range from changing or stealing entries from the database to planting subtle delayed action malicious software (e.g., to delete a specified database at some time in the future or upon some event in the future) to embedding viruses and worms that circulate widely.

- Arrangements for restart and recovery. Not normally thought of as a security issue, nonetheless partial or complete system collapse "from natural causes" will disrupt and deny service. Thus, appropriate procedures, specialized actions, and personnel training are essential safeguards to restore the system to operational status and to do so without risking the secure status of the system.

4.5. Typical Threats

The typical threats are implied by the discussion of security above, but for completeness, they include the following:

- Physical attack intended to physically damage or destroy.
- Penetration, either externally by crackers or internally by dissident or disgruntled employees.
- Exploitation of remote maintenance features.
- Disruption or denial of service from any cause (e.g., severance of a communications cable, loss of primary power, partial system collapse in the PSN)
- Unauthorized actions by authorized users of the system. The risks become more severe for misbehavior by system programmers, system analysts, or maintenance specialists.
- Attacks against a computer system intended to impair its functioning as opposed to attacks against the telecommunications intended to sever the center from the rest of the system.
- Deliberately induced failures designed to disrupt service, to mask some other subversive act, to create the impression of system unreliability, to destroy end-user or management confidence in the system.
- Bypassing security safeguards (for example, introducing spurious software through the floppy drives of desk-top terminals).

5. Heterogeneous Network Characteristics

Excluding telecommunications, the principal distinction of the heterogeneous network is the absence of everything that is uniform in a homogeneous network. There is no single or compatible technical hardware or software foundation among the subscribers who connect. There may be some vendor's system that is popular and used by many subscribers; to this extent, there will be technical compatibility. No single authority sets policy over all the subscribers. No single operational policy dictates what every terminal user can do, although some commonality of user behavior will tend to be established by the nature of services provided by subscriber systems.

The telecommunication arrangements and the interface to them are specifically designed and intended to accept subscriber systems or computer components of many kinds and vendors and to be able to provide connectivity among them.

Importantly, except possibly for the telecommunications portion, a heterogeneous network is not designed and implemented as a cohesive unit; rather there will be intentional flexibility and standardization to accept subscriber systems as they come along, whatever they may be. The only stipulation for connection is to meet interface standards, to abide by the policies imposed by whatever authority controls the right to connect, and to pay any relevant fees and charges.

In effect, a heterogeneous network (1) is usually public in its outreach and serves broad individual and organizational interests, (2) is in effect a common carrier—although not legally—telecommunications arrangement for interconnecting subscriber data systems of various kinds, sizes, and capabilities, and (3) sometimes will provide services other than just connectivity, possibly for a fee (e.g., database access, electronic mail, database search).

The salient difference between the homogeneous and the heterogeneous network is captured in the phrase "what and who can be connected." Other differences flow from this single distinction. The homogeneous network accepts connections as prescribed by the organizational authority, usually corporate, that designed and implemented it; it accepts only devices that the authority stipulates. The heterogeneous network accepts all comers who can and will meet its operational policies and connection obligations. The distinction can blur, however, because a corporate entity might well design and implement its network with attributes of the heterogeneous network if such an arrangement happens to suit its needs.

It is also clear that many things are the same for the two kinds of networks. Both can have dial-in arrangements for users (e.g., people, other systems, other networks, remote maintenance). Both may have some of the telecommunications implemented as point-to-point circuit links; both can have packet links as well. A computer or system in either is likely to support its own set of terminals and end users, as well as connectivity to remote terminals and users.

However, the differences are important to note, especially as they will influence the policy issues, but they are not always sharply drawn.

- While both classes of networks have a telecommunications component, that of the heterogeneous network is likely to have an architecture readily amenable to growth and change as subscribers materialize. It is commonly implemented incrementally and a priori is expected to grow and change.

- The telecommunications component of the heterogeneous network operates effectively under a unified policy albeit not originating in a single authority but from multiple cooperating authorities or from historical events (e.g., the PSN as an example); each individual subscriber system operates under its respective authority, policy, and management structure.

- There is no single authority and source of policy over individual subscriber systems. Each offers the services pertinent to its local community of users but meets network interface standards in order to be able to connect. In effect, network service is an adjunct service to other local services. Moreover, subscriber systems are free to make changes to themselves as they wish so long as communications standards and conventions are not violated.

- The communications authority may level restrictions on a subscriber system; e.g., data rates, data quantity, data content (e.g., no sensitive traffic).

- Individual end users are autonomous and generally free to do anything that the local or remote subscriber systems support, but they are subject to individual user privileges and to network restrictions. Unless a user violates network agreements (e.g., saturates the telecommunications by moving huge files during busy peak hours and hence denies service to others), he or she is largely unconstrained except by criminal law, personal ethics, and morality.

- There will not be a uniform or compatible technical base in the subscriber systems—hardware, software, architecture. Rather there will be multiple hardware/software bases among the subscribers. A subscriber may be another network, a collection of terminals and servers connected to (say) a LAN, or even a single workstation.

  Communications protocols, software, and standards make intercommunication feasible and possible. For example, subscribers to Internet connect IBM systems with various operating systems, or VAX systems with VMS or AT&T/UNIX or BSD/UNIX, or Hewlett-Packard systems with their operating systems, or SUN systems with SUN/UNIX. All, however, support a common communications interface, including protocols, and have collectively learned how to interoperate successfully.

- Relative to a homogeneous network, there will often be a huge number of subscriber systems attached to a heterogeneous network. A big homogeneous net might have a few tens or thousands of computers or systems supporting many tens of thousands of terminals; the Internet has many millions of subscribers.[6] Obviously, small heterogeneous networks can be built and may well be in the future.

- The connectivity among participating subscribers resembles that of the telephone network; each subscriber can potentially connect to any other one at will. What connectivity does exist at any moment depends on the interests and activities of the users. The telecommunications will support any connection; any two subscribers that agree to communicate can do so. Thus, the traffic through the network is defined by user habits and actions, not by a prescribed workload supporting an organization.

- Traffic patterns among subscriber systems are generally not known initially but develop as the users work from day to day. They are not established by an a priori requirements analysis of the workload to be supported or by an operational schedule.

- General characteristics of traffic flow are known historically but are not predictable except on a statistical basis or history. They can be quite variable and fluctuating; for example, a national event might focus a lot of attention and message traffic on a particular subscriber or individual.

- Subscribers are typically spread not only over large geographic expanse (e.g., throughout the United States, worldwide) but also are physically located in a very large number of organizations.

- Because of the diversity of subscriber systems and the absence of a single overall policy authority, inter-system security features become necessary. For example, systems must identify themselves to one another, may have to authenticate the identification, and possibly exchange explicit information about the kind or amount of traffic to exchange. Even now, not every system will allow remote file transfers or remote identification of users. This aspect is likely to become much more important as the NII evolves because of the diversity of users expected. For example, the procurement traffic of the DOD supported by electronic commerce must be delivered only to specific subscribers.

- System-to-system privileges and processes will be defined by mutual agreement and advertised as services-available-to-all. Such decisions are made

[6] As of early 1993, the subscriber systems connected to Internet were estimated to be 1.3 million distributed over 21,000 domains. A "domain" is analogous to a zip code; a subscriber system, to an individual mailing address within a zip code. The "domain server" is responsible for knowing of all subscribers within its domain and for providing routing information to control traffic movement among subscribers.

by the individual subscribers and can include such things as archives of software, archives of all the electronic mail of a discussion group, remote identification of end-user status, or file transfer capabilities.

The distinction between homogeneous networks and heterogeneous networks is not rigidly defined. Each can have features of the other, but there are attributes that tend normally to be characteristic of one, but not the other.

5.1. Examples of Heterogeneous Networks

There are many fewer heterogeneous networks than corporate homogeneous networks.

- Internet—the contemporary consortium that evolved from the original Arpanet plus later the Milnet and NSFnet and a number of worldwide nets. It is far and away the best known and largest heterogeneous network.
- The descendants of Arpanet around the world. Most major countries have at least one national network patterned after the publicly available technical details of the original work. Often such national networks exist through the sponsorship of a consortium of educational institutions or the common interests of a research community.
- Experimental regional networks now in the pilot phase that are anticipated to evolve into (what the United States is calling) a National Information Infrastructure.

5.2. Security Concerns of Heterogeneous Networks

The security concerns of a heterogeneous network are generally all of those of the homogeneous network plus additional ones that arise from inter-system requirements to identify and authenticate one subscriber to another and to define the data categories or programming processes permitted to be exchanged. To illustrate this point, observe that a telephone subscriber assumes that the switch will make the correct connection; there is no echo back from the dialed number that it indeed is the correct one. The subscriber depends on voice recognition or other responses from the dialed recipient to establish the correctness of the connection.

Facsimile transmissions through the PSN resemble telephone calls, except that the far end provides no identification of itself to the sender but sometimes will acknowledge successful receipt of traffic. In the parlance of the control engineer, the connection process is largely open-loop.

It is customary in telecommunications to use a "handshaking protocol" to make sure that the proper link has been established. Handshaking is, so to speak, a back and forth digital "conversation" that exchanges prescribed data elements such as identity, preferred data rates, and nature of error control. In the case of encrypted links, the handshaking, among other things, gets the encryptors at the two ends into time synchrony and may involve exchange of an encryption key. In the case of modem connections, digital signals are exchanged that result (among other things) in selection of the data rate that will be used and any error-control features.

In a general heterogeneous network, especially one serving all subscribers meeting connection requirements, there is a security risk to open-loop arrangements. A message or a file can be transmitted to the wrong recipient, and if the missent material is confidential or corporate-proprietary, it can potentially be a serious matter. Thus, system-to-system handshaking of some sort is generally necessary to identify each to the other as protection against misrouted traffic. The mutual identification might be at the "front door" of the system—its network attachment point—or it might be at an "inside door"—at the software process receiving the traffic (e.g., electronic mail software). [EN3] In the longer run, system-to-system encryption may well become a preferred approach to intersystem security. Systems might also include a separate step to authenticate the asserted identity.

As homogeneous networks proliferate and as subscriber-provided services to other subscribers expand, there will probably also have to be system-to-system—or perhaps process-to-process or even user-to-user—exchange of data relative to what is permitted to flow between them. This amounts to an additional level of security concern about network services, some of which might be restricted but others not, some of which might involve fee-for-service, some of which might involve confidentiality protection or integrity assurance for data.

For example, electronic mail is a very widely available and used Internet service. At some time, the present mail systems will be enhanced or replaced with ones that can provide confidentiality of transmission from originator to recipient. [EN4]

As an example of control of access—a security issue—to a service, consider electronic mail. The policy issue will become: may any user of a subscriber system have access to protected mail; do we depend on the judgment of the individual to decide when it is appropriate? And the collateral policy issue: for what categories of messages should protection be automatic? Today the answer to the first would be "yes," and the answer to the second would be "the question hasn't been addressed." There are arguments against automatic protection of all mail. Among the considerations are the following:

- Not all mail needs to be protected; perhaps even most mail as the service is now used, need not be. But as other communities of users come on to heterogeneous networks, there may well be categories of messages (e.g., exchange of incomplete research results, discussion of job offers, submission of personal references, exchange of corporate proprietary information such as the details of a proposal bid) that should be protected. The attitudes of mail users have much influence on the need for confidentiality, and privacy law that may be passed will also influence the situation. Legal decisions that involve electronic mail might also have some effect (e.g., someone sues for libel on the basis of an electronic mail message).

- Real costs will be involved in providing confidentiality. Since encryption and possibly special formatting will be involved, processing a protected message will take more computer-system resources, and this equates to additional load on the system and, hence, some additional cost. The chatter in on-line discussion groups would not seem to warrant the incremental cost of protection.

- Export controls on cryptographic devices and software will also limit the security that can be provided to world-wide mail and to other services.

Today network-wide services are already enormously useful, even though limited in scope. The full impact of their security implications is far from obvious. Some form of access control to services is probably required, but it is likely to have somewhat different details and scope than the phrase "access control" has historically meant in the context of a computer operating system that controls access of users to a database (i.e., the meaning of access control as it appears in discussions of trusted technology). [6]

Also to the point, "user" in the network context has a much broader meaning. It is no longer just an individual; it might be a remote system asking for services; it might be a process on a remote system asking for services; it might be any one of these coming through several intermediate systems. The user behind such requests might well not be visible or even logged on at the time of the request. Yet, the system that is being asked to respond must have assurances that the request is legitimate and authorized and perhaps that it will be paid for. Perhaps the right phrase for the generalized inter-system, sometimes sequential, steps in handshaking is "credentialing." We do not yet understand fully just what will be required and under what circumstances, nor do we understand what credentials one system may have to pass to another, nor do we understand what levels of trust may have to be assumed between subscriber systems or satisfactorily established by one to the other.

5.3. Typical Threats

Security threats for heterogeneous networks are generally the same as for homogeneous nets except for a magnification aspect. There is no central authority and enforcement mechanism to maintain "network law and order." To be safest, each subscriber will have to automatically be suspicious of every other subscriber wishing to connect to it. Handshaking will have to establish mutual trust for the connection of the moment, and it may have to be established through one or more intermediate systems. This situation implies that more sophisticated penetrations and attacks are certainly possible; at minimum there will be more sources from which they can arise.

The heterogeneous network is more open in the sense that more subscribers from more organizations have access to it, and collectively they are not under a common discipline as would be the users of a corporate homogeneous network. Thus, its systems and its data potentially have more exposure, with the result that some security threats will be of more concern than they might be in a homogeneous situation; e.g., widespread propagation of malicious software.

6. Present Focus

The target of interest for the balance of the discussion is the heterogeneous data network that provides connectivity to a broad range of prospective users and systems and is funded by government funds, private-sector funds, or fees. It is "public"—in the access context—in contradistinction to "private"; this implies that the network is generally available to qualified user organizations, not necessarily to the general public at large, but possibly so; it is not intended to support the needs of any single organization.

While the term "network" has been used in a broad connotation, it must now be restricted to just the telecommunications portion of the overall ensemble; this will follow its usage in such terms as Arpanet or Internet or LAN or PSN. The network of concern will be a telecommunications arrangement for providing connectivity among a set of subscriber systems, and it will be called heterogeneous because it is designed to accept a connection from any subscriber system that meets its interface standard and otherwise qualifies.

For our present purposes, the network boundary will be taken as the interface to subscriber systems. Why at this location? Why exclude the connected subscribers?

- A heterogeneous telecommunications arrangement can support any subscriber that can meet its interface requirements, agrees to traffic flow and other possible restrictions, qualifies in other details and will pay costs and/or fees. Aside from these few stipulations, the network authority has no cognizance of subscriber organizations. Thus, whatever the policy appropriate for the operation of the network, it cannot be imposed upon subscribers. Moreover, since the expectation is that hundreds of thousands, even millions, of systems will subscribe, there is no practical way to establish an overall uniform policy for the totality of telecommunications plus all subscriber systems. It would be inappropriate to even try, much less work out all the legalities that could arise.

- The Arpanet at first, and now the Internet, has existed and prospered under just such an arrangement. The self- administering, self-policing approach has proved workable; but it should be noted that the category of individuals involved through subscriber systems are generally computer literate, come from academic, research, or corporate organizations, and usually behave responsibly. Such fortunate circumstances are not guaranteed for the future or in the NII.

- The forthcoming NII, as we perceive its early stages, is in the image of the Internet.

- Such a posture makes the heterogeneous network essentially the data analog of the PSN for voice communications and other services provided by the telephone industry.

- Looking to the future in which there may well be many commercially offered common-carrier data interconnection networks, the homogeneous network is the only feasible general-user, subscriber-oriented arrangement. The country can never design something which is big enough to accommodate all the projected future demands yet will be implemented with the uniformity and centralized control of the homogeneous network, as characterized previously. Starting from scratch, a country with even the technical assets of the United States would be very hard put to design a telephony system and get it all correct. Moreover, we should learn from evolution, not ignore its lessons.

- Furthermore, the diversity of services that is expected will come from entrepreneurs offering them over the network, as well as from the experimental work of scientists and engineers. There has to be a convenient means to interconnect with a network over which the offerings can be hawked and made available.

What governance is appropriate in such a situation? The present guidance for the Internet is a cooperative collection of committees. To the extent that a "network authority" exists, it is distributed in any sense of the word. Yet, it has functioned well and been the guiding force that matured and evolved Arpanet into Internet. The presence of extensive government funding in the early development days, together with the oversight and technical monitoring of ARPA (later, DARPA and now again ARPA) with its many contracts, undoubtedly brought cohesiveness and responsibility to the participation by "the many." To the extent that an overall "network policy" now exists for Internet, it is the result of the composite judgment, dedicated service, and contributions of many individuals and organizations, and in some measure, of professional societies. Perhaps this is a way for the future; perhaps, some central forum must be formally established to determine overall network policies. We shall have to see what governance will be appropriate as the National Information Infrastructure evolves.

It is useful to view the PSN and AT&T as illustrative analogs. Under AT&T the PSN was essentially a homogeneous network. Post divestiture the PSN has acquired many aspects of a heterogeneous network (e.g., telephone sets come from many vendors; there are varied attachments such as modems, answering machines, and other specialized devices; switches internal to the network are technically compatible but come from many vendors). Operational authority is vested in the network operators; subscribers need only agree to interface obligations to obtain a connection but, beyond that, are not constrained in their actions and behavior except as it might be criminal or morally offensive. Government, in behalf of the general citizenry, watches over the whole thing and exerts influence on usually economic matters and service-related questions.

The PSN no longer functions under the de facto central authority and policies of AT&T as influenced by federal policy and regulations. Although the technical arrangements and standards originating with AT&T prior to divestiture still generally govern, the present PSN functions under policy set partly by the Federal Communications Commisson (FCC), partly by the Public Utility Commissons (PUCs) of the many states, and partly by federal or state law, or both.

The PSN is also getting much more flexible; like the heterogeneous network, a variety of subscriber services are becoming available (e.g., choice of carrier, how to pay for a call, multiple simultaneous conference connections, reporting of the calling party phone number, storage of frequently called numbers within the system, speed dialing, connection of arbitrary equipment to the network so long as it meets prescribed standards, and in the long run, more general data-related services). Overall, however, the decision to offer enhanced services is influenced partly by the technical circumstances of the regional operating telephone companies, partly by decisions of PUCs, partly by economic concerns, partly by business decisions.

7. Framework for Security Discussion

Because security in its technical connotation is ultimately the driver for some of the policies that will be required for successful functioning of a heterogeneous data network, it is useful to review the situation with regard to federal government involvement and standards.

By way of review, the concern is with networks that deal in data, not ones whose primary purpose is voice connections although some might support both. Various communications technologies (e.g., fiber, satellite, microwave, cable) might be used in implementing the network. In fact, many vendors are likely to be involved and may change from time to time. Generally, the focus of discussion is directed to packet-switched networks, which sometimes might contain circuit-switched portions or have circuit-switched extensions. Usually there will be traffic switches within the telecommunications arrangement; sometimes the subscriber system will select the appropriate link or routing and dispatch the traffic, but most often the network will handle the routing automatically. The network need not be a single entity funded from a common source; in fact, the Internet is a collection of individual smaller networks somewhat in the image of regional operating telephone systems and funded by network-access fees from subscribers. There are, however, backbone telecommunications that are government funded.

Geographic extent of the network can be from regional to national to international.

Finally, attention will be particularly concerned with government-funded but widely available networks whose general goal is to provide high data-rate connectivity among universities, laboratories, and other organizations that have a collective interest in and importance to national interests and goals.

We have set the boundary of "the network" at the interface to a subscriber system. It is understood that special hardware and/or software might be required to make the connection and that some of the software may have to reside in the connected subscriber system.

Within the boundary, our concern will be security as a technical issue plus policies that are required to support the technical actions.

8. Computer Security Review

For general background and context, it is useful to present a brief review of computer security, a term that was introduced prior to networking and was intended generally for the stand-alone computer system supporting its local community of terminals. It has gradually been extended in meaning to include arrangements that include more than one computer or more than one computer center, and it is often used in a very generic sense when speaking of any configuration of computer and communications components.[7]

With regard to systems having one center that supports terminals hard-wired to the center, computer security as a technical issue began in the Department of Defense in the mid 1960s. In 1967 at Atlantic City, it was presented to the general computer community during a session at the Spring Joint Computer Conference. [17] Shortly thereafter, a DOD-sponsored committee wrote a definitive document that structured the overall issue, made various recommendations for progress and policy, and is surprisingly today still a relevant and useful discussion of the topic. [18]

However, at that time, there was little understanding of security safeguards for operating systems, and hence the DOD so-called "DSB report" said little about them.[8] During the 1970s the USAF and (then) ARPA supported a variety of research projects generally addressing operating systems, developing relevant concepts for incorporating safeguards in them, and for actually building them.

Near the end of the decade there was enough understanding to support a series of workshops that collectively yielded a set of criteria published by the DOD's (now) National Computer Security Center as "Technical Criteria for Security Evaluation of Computers"—the so-called TCSEC or Orange Book. [6] It defined various classes of security scope, specified the safeguards relevant to each—the so-called security features—and identified a specified level of confidence for the validity and operational functionality of the safeguards—the so-called assurance levels.

While the TCSEC was never explicitly restricted to software, it has, in fact, had the most effect on operating systems, database managers, and a few other major software components. It has had little effect on hardware; in fact, there is no comprehensive statement or standard specifying the hardware safeguards that computer equipment of any kind should contain. There are hardware features that originated from those necessary to make time-sharing systems work successfully

[7] For a more thorough review of computer security, see [15, 5].

[8] This report is popularly known as the "Ware report."

(e.g., a privileged mode of operation, memory segmentation and protection); these have been generally sufficient to support security software.

The quid pro quo between government and the hardware vendors was that the government would test products at no cost—formally, the process is called evaluation—if the vendors would develop the products with their own funds. Over a period of a decade or so, most vendors in the computer industry have submitted products to the NCSC for evaluation, and many are now on the Evaluated Products List. [7]

It is essential to note the mindset and ancestry of the effort that eventuated in TCSEC; namely, defense interests, their structuring of the security problem, their unique needs, and importantly their perception of the threat. Even though there was a substantial effort to persuade civil government agencies and private industry that evaluated products were suitable for their needs, in fact the TCSEC has had little direct effect outside of defense. It has had significant indirect effect in that the civil and private sectors can acquire commercial products that are better designed from a security point of view, do contain useful safeguards, and are, overall, better products as a result of the formal evaluation process.

The defense effort was directed primarily at a stand-alone system, primarily its software; that is, it addressed systems containing one or a few interconnected computers, collectively serving a community of terminals that might be nearby or distant—the arrangement that has been called above a heterogeneous network.

The TCSEC was "interpreted" to create derivative criteria for other software components (e.g., one for Data Base Systems). [13] There was also a network interpretation nicknamed the Red Book, which will be touched on later. [14] In effect, an interpretation says: "This is how requirements of the TCSEC are to be understood in the context of . . . "

8.1. Civil Sector Security Needs

Through the 1980s it became progressively clear that the dominant threat components for civil agencies and private organizations were quite different from the defense perception and that the TCSEC classes—therefore products evaluated against it—lacked features that would be useful and desirable outside defense.[9]

While there still is no formal statement of the computer security threat to guide designers and operators of systems in civil government agencies or in the private sector, there are some generally accepted sources of concern. One such

[9] For example, see [3].

concern—and perhaps the most prominent—is the insider threat: the unauthorized actions of authorized users. It is hard to guard against, and the DOD in effect sidesteps it by considering the process of background investigations for its people and subsequent issuance of security clearances, plus the military hierarchical authority structure to jointly be a satisfactory surrogate for establishing trustworthiness.

Another point of difference is that of data integrity. This is a major concern for private industry, but it has largely been ignored by the DOD effort. Another point of difference is the structure of the access controls; the label-oriented ones for defense needs are not obviously the right ones for commercial users.

Subsequent to the TCSEC and undoubtedly because of its influence, other countries also created sets of evaluation criteria; all reflected the TCSEC to some extent although there were major divergences in how some things were handled. Canada has its Canadian Trusted Computer Product Evaluation Criteria; the United Kingdom, Germany, France, and the Netherlands started similar efforts but subsequently all have been combined into a joint European effort called ITSEC—Informational Technical Security Evaluation Criteria. Most recently, Japan has produced draft criteria. [9]

8.2. The New Federal Criteria

Eventually, it was realized that something beyond the TCSEC was essential for extra-defense organizations; namely, the non-classified community of government that deals with sensitive but unclassified information.[10] A joint effort between the National Computer Security Center (NCSC) of the National Security Agency and the National Institute of Standards and Technology (NIST) of the Department of Commerce was begun in 1991 to develop a new Federal Criteria.[11]

The expected outcome was to be a document, eventually to become a series of Federal Information Processing Standards, that would be applicable to all of government, more flexible and comprehensive than the TCSEC, attractive to private sector users, and harmonized with other country's efforts, notably with the ITSEC. [8]

The proposed schedule called for a final "version 1.0" draft in approximately September 1993. In following years, international harmonization was to have

[10] "Non-classified" is a convenient term to include the composite interests of civil government, unclassified defense government, and the private commercial and business sector.

[11] The joint agreement stipulating this undertaking was signed and announced in December 1990. The strategic plan for implementing it, including the formation of a Federal Criteria Working Group, was completed in November 1991.

taken place with eventually a "World Technical Criteria" accepted and published by the International Standards Organization (ISO).

Most recently, NSA and NIST have announced that the new Federal Criteria will not be published as a finished document, although the present draft might be consolidated into a subsequent draft. [12] Instead, based on the willingness of the four ITSEC nations (Netherlands, France, United Kingdom, Germany) and Canada to consider a consolidated criteria, there is now underway a so-called Common Criteria. It is to be produced by a six-person editorial board on a schedule that calls for its appearance in the early part of 1994. This is a very ambitious schedule; going into the activity, the most complete and polished document is that of Canada. [2]

Depending upon differences between the TCSEC and the anticipated Common Criteria, vendors will have a lesser or larger task in responding. Under the announced schedule of an initial draft by early 1994, a reasonable expectation would be that products conforming to some of the classes in the Common Criteria might appear in 1995-1996 at the earliest. Products conforming to classes that would require extensive new development by the vendors could not be expected to appear until very late in the decade. All projections depend heavily on the mechanisms and processes that are incorporated in the Common Criteria.

It is important to appreciate what efforts such as the Common Criteria, or any criteria, can be expected to achieve. It will set standards for security safeguards in computers and communications products that are intended to counter a defined threat. In fact, the general thrust of all new efforts is to provide flexibility for the intended end user to be specific about the threat to be countered and to define safeguards accordingly.

The significant point is that any criteria now visualized will produce evaluated components, not complete systems. It is still the responsibility of the organization acquiring a complete system to make its own risk assessment, to make its own threat assessment, to select relevant components and fit them together into a total system, to provide appropriate physical and personnel safeguards, to create the administrative/procedural/management adjuncts in which to embed the system, and finally to conduct such analyses and tests as may be necessary to assure that the completed implementation does, in fact, satisfactorily counter the threats projected.

[12] This announcement was made at the July (1993) meeting of the Computer System Security and Privacy Advisory Board, which functions under the statutory authority of the 1987 Computer Security Act and has broad permissive reporting authority. Its primary mission is to be alert for latent societal or government problems that might arise from the widespread use of computer and communications technology.

This last step—management acceptance that the completed system does indeed do what was expected of it—is essential and cannot be delegated to an external party.

9. Network Security

The story for network security is much shorter and, in fact, sparse. The only official document from government is the Network Interpretation of the TCSEC Criteria—the Red Book. [14] It has little relevance to some kinds of networks, some relevance to others; it is not a complete treatment of nets by any means. It addresses homogeneous nets extensively, but it says very little about heterogeneous networks and then only in an appendix.

There is no work currently underway on United States federal network standards. Given that all available federal computer-security government assets are seemingly involved in pursuing a Common Criteria (for components), it is unlikely that network security standards will be addressed at the federal level before the end of the century. If the NII moves rapidly enough, this may change, but for the moment no guidance exists or is projected for the foreseeable future.

Hence, the design of each network—homogeneous or heterogeneous—is an individual effort by the organization desiring it. Some components are available with security controls and some components meet federal standards (i.e., have been evaluated), but there is nothing for overall design nor is there any assurance that all necessary components for some secured system or network are available with appropriate safeguards in them.

10. Policy Considerations

With the term "network" defined and with the boundary drawn at the interface to subscriber systems, what policy considerations are appropriate to a homogeneous network?

First of all, there are clearly policy issues not related to its security that will be of intense political interest for the NII. Most notable will be the question of access to the network. Who may be allowed to connect to it? How can discrimination be avoided? What sort of organizations may subscribe? Will government provide subsidies to support equality of access? Will subscribers, especially those that provide services, be under any compulsion of affirmative action steps? This collection of issues, commonly called the "equity issue," is not of concern in this discussion of homogeneous data-network security.

Another policy issue, also not of concern here, relates to classes of service. For example, will a group of packets be delivered to the recipient in the same sequence as dispatched or as they happen to arrive? Will delivery be assured within some prescribed time? If there is a fee structure, does it depend on the class of service used? For example, a fee structure might be a part of services policy (e.g., charge more for "return-receipt delivery"; charge less for off-peak delivery). At some point, traditional economic views and issues are bound to enter data network considerations and policy.

10.1. Security Policy

In the context of network security, there are various broad categories of policy concerns for heterogeneous networks.

- Those relating to the operation and self-protection of the network per se. Among them might be restrictions on users, conditions for accepting a user system, interface standards for users, services that the net provides to all users or to users for additional fees, and intra-net security. In effect, a policy that declares "Here is how the network will be operated; here are the terms and conditions for joining it."
- Those relating to the protection and handling of subscriber traffic while in the custody of the network.
- Those relating to continuity of service or, to use a currently popular term, network availability.
- Network safeguards to protect subscriber systems.
- Those that govern subscribers per se whose systems may be not only users of the network but also may provide services to it. Among them might be descriptions or limitations of services that each subscriber may provide, services each must provide, legal obligations that each might have, equal-access user rights each must assure, costs or cost-sharing each must bear, and security obligations each must assume as a condition of joining the net. In effect, a policy declaration that "If you join the network, here are the operational, service, fiscal, and legal obligations that you must meet."

With regard to security explicitly, it is possible that it would be only of limited interest to the policy authority responsible for the network, its funding, and its management. Security in the broad sense might not now be a public policy issue; only continuity of service might be a concern although this is closely related to security and, in some respects, merges with it. For example, continuity of service and security might be of concern if the services provided by the network are deemed to be important aspects of public safety or public interest. If poor security were deemed inimical to fulfillment of the purposes and goals that the net was intended to fulfill, security might place very high in importance.

10.2. Network Self-Protection

With regard to safeguarding the network itself, security policy at minimum must concern itself with reliable continuous operational service. Otherwise users will be deterred from depending on it, and the full impact of the network's presence, the societal consequences, and the payoff hoped from it will not materialize.

As with considering security safeguards for any computer or telecommunications system, the first order of business is to postulate the threat that will oppose the network. Having done so, then security policy might contain such typical provisions as the following:

- There must be physical protection for all components of the network.
- If the network switches (packet or circuit) are to be serviced and maintained remotely, hacker and intrusion protection must be provided and carefully attended.
- If the network control center has access to the software of the switches via the network (e.g., for downloading new software, repairing bugs), controls must be in place to avoid this being a security weakness, especially a single point-of-failure risk.
- Safeguards must be in place to prevent the behavior of one or a collusion of users from disrupting the network (e.g., by flooding it with traffic, by overloading switches with intentionally incorrect messages). In this regard, note that collusion might be de facto, in the sense that a group of subscribers may not a priori have been solicited to conduct improper behavior but rather, in responding to the actions of a rogue subscriber or a rogue program, effectively have been co-opted into becoming members of a collusion.
- Safeguards may be needed to protect the network against disruptive attempts of other kinds (e.g., deliberate actions by a subscriber system to penetrate a switch).
- Safeguards must be in place to avoid disruptive effects that might be mounted through the network control center.

Some of these might be regarded as aspects of network reliability or continuity of service; whatever they are called, the protective mechanisms must be present.

10.3. Traffic Protection

The opposite side of network protection is the policy that speaks to protection and handling of the subscriber traffic while it is in the custody of the network (that is, stored within it or in transit through it).

If electronic intercept is perceived to be a threat, especially on international telecommunication links, communication security will be stipulated as a part of network security policy. Depending on the threat that is perceived as facing the network, protection might be limited to encryption of only the backbone communication links. If the interception threat were seen to be more pervasive, end-to-end encryption might be stipulated as network policy. In this event, "end-to-end" would mean from the interface of the system that originates the traffic to the interface with the system to which it is addressed. Network obligation stops at its boundary which, as in the context of this discussion, would by agreement be the interface between a subscriber system and network. [EN5] Specifically, encryption would not extend to individuals who originate or receive traffic, unless the subscriber systems themselves provide for it.

Technical means exist to provide end-to-end encryption through packet switches as well as in link switches; technology is not an impediment. Interestingly, insistence by the network authority on encryption can make possible other important assurances to the subscribers (e.g., even during maintenance operations on network switches during which someone might see message traffic, it will be protected by encryption against inadvertent disclosure). There are collateral questions of managing the encryption keys so that no one within the network who might see message traffic can have access to keys; this is largely an administrative and procedural matter with technological support.[13]

10.4. Subscriber Protection

Another aspect of network security is the protection of each subscriber system against all others by safeguards within the network. Off hand, inter-system matters

[13] Some standards already exist for key management in support of the Data Encryption Standard (Federal Information Processing Standard FIPS 171, Key Management using ANSI X9.17, April 27, 1992) and others are being developed by the IEEE 802.10 Working Group.

might not appear to be an aspect for network concern, but consider the following possibility.

A subscriber system must not be able to co-opt one or more others into a coordinated attack on the network. If attacks on one subscriber system by another can imply consequences for or attacks on the telecommunications network, then it becomes a concern for the network per se. For example, suppose one subscriber system were able to "persuade" several others to ping-pong hundreds or thousands or millions of transmissions back and forth. It could result in saturation of network resources and, hence, denial of service to other users. Alternatively, the same sort of ring-around could result from an accidental glitch in software.

It is not clear what security measures must be in place within the network to guard against such problems, but possibilities exist. For example, a condition of connection might be a stipulation on maximum data rate into the network, perhaps staged over various periods of time, or a maximum rate of message injection. Perhaps the network ought to monitor such things and either throttle the flow from an offending source, cut it off, deliver a warning message, or at minimum bring it to the attention of the control center.

Whatever technical or procedural measures might be developed, if such protections are to be provided by the network in behalf of subscriber systems, a policy statement to that effect will be required; we may also need an enforcement mechanism.

10.5. Subscriber-Control Policy

Such a policy is broadly concerned with defining, characterizing, and perhaps restricting subscriber behavior and service offerings so that the network resources are equitably available to all subscribers, and so that none can abuse the network. In a no-fee network, this is likely to be of high importance because some systems will always try for more than a fair share of network capability. In a fee-for-service network, such behavior can be controlled through economic pressures (e.g., with a sliding scale of charges).

There are many dimensions of subscriber-system behavior that the network might wish to stipulate or control by policy. For example, there might be federal law that subscribers would have to agree to honor (e.g., non-discrimination, liability, equal-access rights to the network). There may have to be policy statements about services that one subscriber system proposes to offer all others, or policy that requires all subscribers to provide a particular service (e.g., electronic mail). Violation of copyright protection might become a network policy concern, given that there is so much exchange and downloading of software, documents, and published materials.

10.6. Other Examples

Here are two more detailed illustrations of possible policy considerations.

In today's heterogeneous networks, there is minimal protection against traffic arriving at the wrong place or the wrong connections being made. In the case of connectivity, it is roughly like the telephone system; incorrect or incomplete addresses are trapped and rejected by en-route control. Legitimate but erroneous addresses simply result in "wrong connections." Electronic mail approximates the postal system; an unknown or incorrectly identified recipient causes the message to be "returned to the sender."

This may prove to be too loose for networks of the future. To gain efficiency in the network, it may be necessary to require an originating subscriber system and the recipient system to be mutually assured that the intended party exists at the specified address before beginning transmission. If this were to be required from all subscribers who support mail services, it would amount to a policy declaration by the network about subscriber behavior and service details.

The PSN is moving in such a direction. In the time of mechanical switches, a long distance call moved sequentially through each switch. As each found an available link (i.e., "trunk" in telephone parlance), it would be seized and held until all downstream connections were made. Thus, the links found early in a call set-up could be idle for a long time and, hence, not generating revenue. In the most modern inter-switch signaling systems, the availability of a path is established before system resources are committed to call set-up.[14]

In networks of the future, not every user will be entitled to all traffic on it, nor will every subscriber system be so entitled either. Each originator/recipient pair may have to mutually assure one another that whatever level of sensitivity (e.g., sensitive/unclassified, proprietary, organizational-confidential) is attached to the traffic from the originator is appropriate to the recipient subscriber system and its ability to assure delivery only to the indicated addressee. Even if traffic sensitivity were not made explicit by formal labels, this principle would have to be honored. In this connection, note the role that user-to-user encryption can play in isolating the traffic and assuring confidentiality.

This point will become much more important as different classes of users begin to populate heterogeneous networks (e.g., academic users will want to protect possible patent opportunities by keeping their research out of sight to corporations; corporations will not want competitors peeking into their traffic). The growth in

[14] Signaling System 7, developed by AT&T.

the number of commercial subscribers to Internet may force this issue into prominence before the NII comes into operational status.

10.7. User Security Controls

This last policy concern is probably the most awkward and trying; namely, what security safeguards will the network authority insist that a subscriber provide in his system before being allowed to connect? These would be safeguards acting in behalf of the network itself as well as in behalf of the family of network subscribers. They are not to protect a subscriber system per se in its internal workings but to protect the network from such things as incorrect routing of traffic or improper behavior of systems that would prejudice the operation or user image of the network (e.g., a dissident system floods another with traffic and seriously intrudes on its normal service ability). The issue goes beyond that of interface standards.

There is a weak analogy in the PSN; a subscriber may not impair the PSN's proper functioning for all other subscribers by connecting an offending device to it or injecting excessive signals or by taking any other abusive action. There is also a historical progression in the PSN that may be a useful model for the future of homogeneous networks. Prior to the Carterfone decision, no foreign attachments were allowed on the PSN. Subsequent to it, the consumer was required to obtain certification that his foreign device did meet telephone interface standards. Eventually, the certification step passed to the vendors, who attest that a device does meet stipulated interface requirements, in this case ones promulgated by the Federal Communications Commission. Connecting to a data network of the Internet kind is much more complex than attaching to the PSN, but it is possible that we will see devices and equipment come with vendor certification that they are ready for network connection.

Today the matter is not really addressed; tomorrow it almost certainly will have to be as heterogeneous networks come into wider existence and use.

11. Network Services

Even though it is convenient for purposes of understanding and exposition to draw the network boundary conceptually at the subscriber interface and to separate policy issues correspondingly, such a clean dichotomy is not entirely feasible. We have already noted examples of policy matters that are properly the purview of the network authority (using "network" in its delimited sense) but at the same time impose rules of behavior or agreements on subscribers. So to speak, the network authority will levy policy behavior on the subscribers in "behalf of the greater good of the Network"—now using "network" in the comprehensive sense.

The question of general Network Services muddles the distinction even more. Such services are ones that might be provided within the telecommunications part of the network and under the aegis of the network authority (e.g., directory services, white or yellow pages service, subscriber identifications, routing information), or they might be provided as a courtesy, or a fee-for-service feature, to all subscribers. In today's Internet, such services already exist but not as yet on a fee basis.

Examples include the subscriber system that offers anonymous electronic mail dispatch. Send mail to it, and it resends the mail but without identity of the source; so to speak, it is a mail-front. Other subscribers offer network search capabilities such as the increasingly popular `archie` service that regularly canvases all hosts supporting anonymous-ftp, searches their holdings, and assembles a comprehensive database of them. Then `archie` servers can search the composite database to locate items of interest and report the sources to a particular user query. Increasingly, subscriber systems are putting archival and library holdings under ftp or list-server accessibility, or putting local databases under the purview of general network-wide search services such as `gopher`, Wide Area Information Service (WAIS), or World Wide Web (WWW).[15]

Other subscriber systems contribute some essential network-support service (e.g., a domain server for a cooperating group of subscribers, some specialized directory or information service). Today, such things are provided sometimes by agreement with the network consortium that operates the Internet, or sometimes by agreement with other subscriber systems and as a pro bono contribution to the network. At other times, a subscriber will simply announce the availability of a service-for-all and operate it for the convenience of the entire Internet community. Formal policy issues are not explicitly of concern; everything runs on the good will and honest behavior of subscriber systems. On the other hand, many of these more sophisticated services are in the nature of academic research efforts.

In the longer view of heterogeneous networks serving a very wide and disparate variety of subscriber clientele, not all of which might be as disciplined as the present Internet family, such informal arrangements may prove inadequate. A plethora of commercial subscribers offering competing services might take things out of control. If so, then some formal authority structure will be required to oversee formalized arrangements, and policy declarations will have to be established and enforced.

[15] See [10] for details of these services.

Perhaps we will need contract arrangements between the network authority and any subscriber that provides or offers general services of any kind to the network, especially if on a for-fee basis. Probably, we will have to create some payments mechanism and payments-settlement process to exchange fees collected for services, or for use of copyrighted materials. Perhaps we will need an oversight mechanism to make sure that contract terms are honored. We hope that we will not need enforcement and penalty mechanisms to compel performance, but at this point, the possibility must be acknowledged.

For many reasons, therefore, there will be an unavoidable policy interface between the network authority responsible for the governance and operation of a heterogeneous network and its subscribers. The precise details and the implementing mechanism are far from clear, but the necessity seems evident.

12. External Policy Influence

Other federal policies might impact the policy position that a network can or must take. The most notable and current one is the national dialogue beginning to take place on the use of cryptography outside the defense and military sphere. There is growing pressure for generally available exportable cryptography to protect communications and to provide message integrity and authentication. For some communities, such as the financial one, it is important to know that messages have not been modified in transit; this would be called providing an integrity check. For other communities, such as electronic commerce, it is essential that there be means for digitally signing messages so that the signature cannot be repudiated (that is, the recipient of a message knows with certainty who originated it and the originator cannot deny that he or she did so). In the submission of bids and proposals, there must be a means for protecting them against viewing by unauthorized eyes.

Such advanced features, obviously important to the spread of digitally based business and the conduct of government over heterogeneous networks, all depend on the science of cryptography. Some efforts are underway to provide federal standards, but other parts of the dialogue have yet to begin.[16] So far the government has taken no general position on the basic question of "cryptography for the masses" but it may be forced to do so because market pressures from United States vendors engaged in international commerce may cause it, or technological advances may threaten to finesse government action, or both.

[16] One effort is the NIST-proposed Digital Signature Standard which, as of this manuscript date of October 1, 1993, is unresolved and hanging on patent awkwardnesses. [20]

Law enforcement is also concerned about technical advances in communications and about widespread use of cryptography. Traditional wiretapping approaches could be impeded; telephone conversations might be encrypted. To this end, the Department of Justice and the Federal Bureau of Investigation jointly proposed during 1992 a government policy that would require providers of telephone service to install technical features to facilitate continued wiretap capability. Congressional action did not occur, but the matter may come before the 103rd Congress.

More recently, the government announced by presidential order on April 16, 1993, "escrowed-key technology" which at present is a voluntary effort for telephony only. [17] The microchips that perform the encryption broadcast a unique serial number on each call which, together with hidden permanent secret keys in the chips plus other secret keys held by trusted escrow organizations, allow law enforcement agencies, after obtaining a court-approved warrant, to deduce the session key in use and decrypt the transmission. This development is almost certain to come before the 103rd Congress during 1993.

Whatever position the federal government eventually takes in regard to cryptography for personal or private sector use, or for incorporation into mass-market software, the resultant policy may well restrict the use of cryptography or the kind that can be provided within networks and by subscribers. Federal action on cryptography can intrude on policy positions that the NII or other network authorities may wish to take.

Endnotes

EN1. This is a satisfactory assumption conceptually, but it can raise practical problems. For example, some of the (domain) servers in the Internet are provided by a subscriber facility and are colocated with other components of the subscriber system. Even though a separate and dedicated computer might be used for the server, its security status is that of the subscriber facility. Hence, if subscriber security is poor, so also is likely to be that of the servers; the security and safety of the Internet is risked. If the domain function is not on a dedicated machine but is rather a software function within a subscriber host, the security of the Internet will be even more dependent on that of the subscriber system.

[17] This is the so-called escrowed-key cryptography (associated with the nicknames Clipper, Skipjack, and Capstone), which is an Administration-proposal to provide encryption but with assured access by law enforcement authorities under court-control authorizations. In this connection, there is an NIST-proposed Escrowed Encryption Standard that is also unresolved as of the manuscript date. [20]

EN2. The Trusted Network Interpretation [14] noted (pp. xiii-xiv) two network views: the interconnected accredited AIS view and the single trusted system view. The bulk of the TNI deals with the second, which came to be referred to as a homogeneous network; only Appendix C with the first, which came to be referred to as a heterogeneous network. At the time of publication (July 1987) the technical status generally supported such a distinction; in the subsequent years, the onrush of LANs, workstations, servers, etc., has blurred but not eradicated the distinction. The two categories are still useful, especially as they relate to and drive policy considerations.

EN3. In the Internet, the inter-system security is variable and incomplete. For some things, there is little protection (e.g., bogus electronic mail can readily be sent and received). On the other hand, there are some operational features that yield a weak form of security (e.g., mail addressed to an unknown recipient is returned to the sender with a statement to that effect). For other things, checking occurs at the process level. For example, the sending system inserts its network address into every packet, and it also sends its network name as part of the initial handshaking; the receiving system can validate that the name and the address correlate correctly. Some error-checking is built into the routing tables and the routing software. However, the watchwords are “checking” and “error control,” not “security control.” There is clearly overlap between the two kinds of controls, and features that support the proper and reliable operation of the Internet contribute to security. However, genuine security controls would be carefully protected against subversion by accidental or malicious events.

EN4. There are several experimental or development software packages and a few commercial packages available. At the Internet level, Privacy Enhanced Mail (a misnomer because it is really confidentiality enhanced mail) is in the early stage of deployment. It is defined by Internet documents RFCs 1421 (PEM Specifications), 1422 (Certificate Management Infostructure), 1423 (Algorithm Identifiers), and 1424 (How to Use PEM to Do Remote Issuance of Certificates). “RFC” equates to Request for Comment, a mechanism by which the Internet community floats a technical proposal, receives comments, revises the proposal, and eventually converges on (what amounts to) an Internet Standard.

EN5. This point has subtle overtones. Undoubtedly, as with LAN connections, some of the interface software will reside within the subscriber systems. Moreover, subscriber systems might provide services that are essential to network operation, notably routing servers. The issue of drawing an interface sharply and technically soundly needs careful examination before we drift into an arrangement that will prove impossible to secure.

References

1. Paul Baran, *Communications, Computers and People*, The RAND Corporation, P-3235, November 1965. Also: *AFIPS Conference Proceedings*, vol. 27, Part II, 1965 Fall Joint Computer Conference; Spartan Books, Washington, D.C.
2. *Canadian Trusted Computer Product Evaluation Criteria*, Communications Security Establishment, Government of Canada, April 1992.
3. D. D. Clark and D. R. Wilson, "A Comparison of Commercial and Military Computer Security Policies," *Proceedings of 1987 IEEE Symposium on Security and Privacy*, Oakland, California, August 27–29, 1987. Published by the IEEE Computer Society, Order No. 771, IEEE Catalog Number 87CH2416-6.
4. "Compilation of State and Federal Privacy Laws," *The Privacy Journal*, Washington, D.C., 1989.
5. *Computers at Risk—Safe Computing in the Information Age*, National Research Council, National Academy Press, 1991.
6. *Department of Defense Trusted Computer System Evaluation Criteria*, National Computer Security Center, National Security Agency, Ft. George G. Meade, Maryland; CSC-STD-001-83, August 15, 1983. Also adopted as a Department of Defense Standard 5200.28STD in 1985.
7. *Information Systems Security Products and Services Catalogue*, U.S. Government Printing Office, stock number 908-027-00000-1. Published quarterly.
8. *Information Technical Security Evaluation Criteria*, printed and published by the Department of Trade and Industry, Controller, HMSO; London, June 1991.
9. *Japanese Computer Security Evaluation Criteria—Functionality Requirements*, draft version 1.0, Ministry of International Trade and Industry, August 31, 1992.
10. Ed Krol, *The Whole Internet—User's Guide and Catalog*, Sebastopol, CA, O'Reilly and Associates, Inc., 1992.
11. Personal Privacy in an Information Society, Report of the *Privacy Protection Study Commission*, July 1977. There are also five appendices on specialized topics, including a discussion of how the 1974 Privacy Act had been working.
12. Records, Computers, and the Rights of Citizens, Report of the *Secretary's Advisory Committee on Automated Personal Data Systems*, U.S. Department of Health, Education and Welfare, July 1973. DHEW Publication (OS) 73-94.
13. *Trusted Database Management System Interpretation*, NCSG-TG-021, U.S. Government Printing Office, stock number 008-000-00582-6.
14. *Trusted Network Interpretation*, National Computer Security Center, National Security Agency, Ft. George G. Meade, Maryland; NCSC-TG-005, July 31, 1987. U.S. Government Printing Office stock number 008-000-00486-2.
15. Willis H. Ware, Computer Security Policy Issues: From Past Toward The Future, Conference Keynote Presentation, Seventh Annual Computer Security Applications Conference, San Antonio, TX, December 4, 1991.

16. Willis H. Ware, “The New Faces of Privacy,” *The Information Society*, Vol. 9. No. 3, July–September 1993, pp. 195–211. Also, P-7831, the RAND Corporation, October 1993.

17. Willis H. Ware, “Security and Privacy in Computer Systems,” *AFIPS Conference Proceedings*, Vol. 30, 1967, pp. 279-300; also “Practical Solutions to the Privacy Problem,” pp. 301–304.

18. Willis H. Ware (ed.), Security Controls for Computer Systems, Report of Defense Science Board Task Force on Computer Security, published for the Office of the Secretary of Defense by The RAND Corporation, Santa Monica, California, as a classified document February 1970; reissued by RAND as an unclassified publication R-609-1, October 1979.

19. Alan Westin and Michael Baker, *Databanks in a Free Society*, Quadrangle Press, 1972.

20. An extensive on-line collection of materials on cryptographic issues can be found in the EFF/CPSR Internet system called ftp.eff.org. It supports anonymous ftp for downloading documents.