

Critical Infrastructures You Can Trust: Where Telecommunications Fits

Fred B. Schneider, Cornell University¹

Steven M. Bellovin, AT&T Labs–Research²

Alan S. Inouye, Computer Science and Telecommunications Board³

INTRODUCTION

The security of our nation, the viability of our economy, and the health and well being of our citizens rely today on infrastructures for communication, finance, energy distribution, and transportation. All of these infrastructures depend increasingly on networked information systems. That dependence, with its new levels and kinds of vulnerabilities, is attracting growing attention from government and industry. Within the last two years, the White House Office of Science and Technology Policy, the President's National Security Telecommunications Advisory Committee, the President's Commission on Critical Infrastructure Protection, the Defense Science Board, and the General Accounting Office have each issued reports on the vulnerabilities of networked information systems.⁴ Congressional hearings,⁵ articles in the popular press, and concern about the impending year 2000 problem have further heightened public awareness. Most recently, Presidential Decision Directive 63⁶ has called for a national effort to assure the security of our increasingly vulnerable critical infrastructures.

Networked information systems (NISs) integrate computing systems, communications systems, and people (both as users and operators). The defining elements are interfaces to other systems along with algorithms to coordinate those systems. Economics dictates the use of commercial off the shelf (COTS) components wherever possible, which means that developers of an NIS have neither control nor detailed information about many system components. Moreover, there is an increasing use of system components whose functionality can be changed

remotely and while the system is running (e.g., Web servers). Users and designers of an NIS built from such system components thus cannot know with any certainty what software has entered system components or what actions those components might take.

The trustworthiness of an NIS encompasses correctness, reliability, security (conventionally including secrecy, confidentiality, integrity, and availability), privacy, safety, and survivability. These dimensions are not independent, and care must be taken so that one dimension is not obtained at the expense of another. For example, the protection of confidentiality or integrity by denying all accesses trades one facet of security—availability—for others. Integrating the diverse dimensions of trustworthiness and understanding how they interact is a central challenge in building a trustworthy NIS.

A trustworthy NIS does what people expect it to do—and not something else—despite environmental disruption, human user and operator errors, and attacks⁷ by hostile parties. System design and implementation errors must be avoided, eliminated, or somehow tolerated. It is not sufficient to address only some of these trustworthiness dimensions, nor is it sufficient simply to assemble components that are themselves trustworthy. Trustworthiness is a holistic property of an NIS.

This paper discusses two NISs: the public telephone network (PTN) and the Internet. Being themselves large and complex NISs, they not only merit study in their own right but can help us to understand some of the technical problems faced by the developers and operators of other NISs. In addition, the high cost of building a global communications infrastructure from the ground up implies that one or both of these two networks is likely to furnish communications services for most other NISs. Therefore, an understanding of the vulnerabilities of the PTN and Internet informs the assessment of the trustworthiness of other NISs.

This paper also proposes some ideas for improving the trustworthiness of the PTN and Internet, both in the short-term (by improved use of existing technologies and procedures) and in the long-term (by identifying some areas where the state-of-the-art is inadequate and research is therefore needed). Finally, some observations are offered about Internet telephony and the use of the Internet for critical infrastructures.

WHAT COMPROMISES TRUSTWORTHINESS

The extent to which an NIS comes to be regarded as trustworthy is greatly influenced by people's experiences using that system. However, generalizations from individual personal experience can be misleading, shaped in part by the popular press and information from organizations that have particular advocacy agendas. For example, a predominant cause of NIS outages might not be a good topic for newspaper stories, although anecdotes of attacks perpetrated by hackers seem to be.⁸

Only limited empirical data about systems trustworthiness is available.⁹ And the absence of scientific studies that measure dominant detractors to NIS trustworthiness makes it hard to know what vulnerabilities are the most significant or how resources might best be allocated in order to enhance a system's trustworthiness. Rigorous empirical studies of system outages and their causes need to be undertaken. Empirical studies of normal system operations are also important, because having baseline data can be helpful for detecting failures and attacks by monitoring usage.¹⁰

Environmental Disruption

Trust in an NIS is not further eroded when catastrophic natural phenomena, such as earthquakes or storms, in a region disrupt the operation of NISs only in that region. But when environmental disruption has disproportionate consequences, trust is eroded. Regional and long distance telephone outages when a backhoe accidentally severs a fiber optic cable¹¹ or rodents chewing cable insulation cause a power outage that disrupts Internet access in the Silicon Valley area¹² are just two illustrations. The good news is that the frequency and scope of accidental man-made and natural disruptions is not likely to change in the foreseeable future. Thus, tolerating today's levels of such disruptions should suffice when building a trustworthy NIS for tomorrow.

Human User and Operator Errors

Errors made in the operation of a system also can lead to system-wide disruption. NISs are complex, and human operators err: an operator installing a corrupted top-level domain name service (DNS)¹³ database at Network Solutions effectively wiped out access to roughly a million sites on the Internet in July 1997,¹⁴ while an employee uploading an incorrect set of translations into a Signaling System 7 (SS7)¹⁵ processor led to a 90 minute network outage for AT&T tollfree telephone service in September 1997.¹⁶ However, automating the human operator's job is not necessarily a solution, for it simply exchanges one vulnerability (human operator error) for another (design and implementation errors in the control automation).

Controlling a complex system is difficult, even under the best of circumstances. Whether or not human operators are involved, the geographic scope and the speed at which an NIS operates mean that assembling a current and consistent view of the system is not possible. The control theory that characterizes the operation of such systems (if known at all) is likely to be fraught with instabilities and to be highly nonlinear. Put operators into the picture and now details of the system's operating status must be distilled into a form that can be understood by humans. Moreover, there is the difficulty of designing an operator interface that facilitates human intervention and control.

System Design and Implementation Errors

The challenge of implementing software that satisfies its specification is well known, and failing to meet that challenge invariably compromises system trustworthiness. NIS software is no exception.¹⁷ An oft-cited example is the January 1990 nine-hour long outage (blocking an estimated 5 million calls) that AT&T experienced due to a programming error in software for its electronic switching systems.¹⁸ More recently, software flaws caused an April 1998 outage in the AT&T frame-relay network (a nationwide high-speed data network used by business),¹⁹ and in February 1998 the operation of the New York Mercantile Exchange and telephone service in several major East Coast cities was interrupted by a software failure in Illuminet, a private carrier.²⁰

The challenges of developing software can also be responsible for project delays and cost overruns. Software thus can undermine confidence and trust in a system long before the system

has been deployed. NIS software is especially difficult to write, because it typically integrates geographically separated system components that execute concurrently, has idiosyncratic interfaces, and is sensitive to the sequence in which tasks are executed.

Attacks by Hostile Parties

Attacks specifically directed at NISs running critical infrastructures are not frequent at present, but they do occur. According to FBI Director Louis Freeh speaking at the March 1997 Computer Crime Conference in New York City, a Swedish hacker shut down a 911 emergency call system in Florida for an hour.²¹ And in March of 1997, a series of commands sent from a hacker's personal computer disabled vital services to the Federal Aviation Administration control tower at Worcester airport.²²

Evidence abounds that the PTN and Internet are not only vulnerable to attacks but are being penetrated with some frequency. In addition, hackers seeking the challenge and insiders seeking personal gain or revenge have been successful in attacking business and critical infrastructure computing systems. Accounts of successful attacks on computer systems at military sites are perhaps the most disturbing, since one might expect tighter security there. The Defense Information Systems Agency (DISA) estimates that Department of Defense may have experienced as many as 250,000 attacks to its computer systems last year and that the number of such attacks may be doubling²³ each year. The exact number of attacks is not known because DISA's own penetration attempts on these systems indicate that only about 1 in 150 attacks is actually detected and reported.²⁴ Similarly troubling statistics about private-sector computer break-ins have been reported.²⁵

A Cross-cutting Theme: Interconnection as a Vulnerability

To a first approximation "everything" is becoming interconnected. The June 1997 Pentagon cyber-war game Eligible Receiver²⁶ demonstrated that computers controlling electric power distribution are, in fact, accessible from the Internet. The Internet will ultimately give ever larger numbers and increasingly sophisticated attackers access to the computer systems that control critical infrastructures. Thus, resisting attack is a facet of trustworthiness that, while not

a significant source of disruption today, has the potential to become a significant cause of outages in the future.

Interconnection within and between critical infrastructures further amplifies the consequences of disruptions, conditioning the trustworthiness of one system on that of another. The lesson of the Northeast Power Blackout in the late 1960s was that disruptions can propagate through a system with catastrophic consequences. Three decades later, in July 1998, a tree shorting a power line running to a power plant in Idaho brought about cascading outages that ultimately took down all three of the main California-Oregon trunk transmission and interrupted service for 2 million customers.²⁷ Was the lesson learned?

Interdependence of critical infrastructures also enables disruption to propagate. An accidental fiber cut in January 1991²⁸ blocked sixty percent of the long-distance calls into and out of New York City but also disabled air traffic control functions in New York, Washington, and Boston (because voice and data links to air traffic control centers use telephone circuits) and disrupted the operation of the New York Mercantile Exchange and several commodities exchanges (because buy and sell orders, as well as pricing information, are communicated using those circuits). The impact of such a disruption could easily extend to national defense functions.²⁹ Furthermore, a climate of deregulation is promoting cost control and product enhancements in electric power distribution, telecommunications, and other critical infrastructures—actions that increase vulnerability to disruption by diminishing the cushions of extra capacity and increasing the complexity of these systems.

TRUSTWORTHINESS OF THE PTN AND INTERNET

In some ways, the PTN and Internet are very similar. No single entity owns, manages, or can even have a complete picture of either. Both networks involve large numbers of subsystems operated by different organizations. The number and intricate nature of the interfaces that exist at the boundaries of these subsystems are one source of complexity for these networks. The increasing popularity of advanced services is a second source. Of course, the PTN and Internet are also very different in some respects, most notably with respect to each network's endpoints and their basic modes of operation: circuit-switching for the PTN versus packet-switching for the Internet. Thus, it is instructive to compare and contrast the vulnerabilities and potential fixes

of the two networks together.

The vulnerabilities of the PTN and Internet are exacerbated by the dependence of each network on the other. Much of the Internet uses leased telephone lines as its physical transport medium. Conversely, telephone companies rely on networked computers to manage their own facilities, increasingly employing Internet technology (though not necessarily the Internet itself). Thus, vulnerabilities in the PTN can affect the Internet and vulnerabilities in Internet technology can affect the telephone network.

Environmental Disruption in the PTN and Internet

Link Failures

The single biggest cause of PTN outages is damage to buried cables.³⁰ And the single biggest cause of this damage is construction crews digging without proper clearance from telecommunications companies and other utilities. The phenomenon, jocularly known in the trade as “backhoe fading,” is probably not amenable to a technological solution. Indeed, pursuant to the Network Reliability and Interoperability Council (NRIC) recommendation, the Federal Communications Commission (FCC) has requested legislation to address this problem.³¹

The impact of “backhoe fading” on network availability depends on the redundancy of the network. Calls can be routed around failed links, but only if other links form an equivalent path. Prior to the 1970s, most of the nation's phone network was run by one company, AT&T. As a regulated monopoly, AT&T was free to build a network with reserve capacity and geographically diverse, redundant routings. Multiple telephone companies compete in today's market, and cost pressures make it impractical for these telephone companies to build and maintain such capacious networks. Furthermore, technical innovations, such as fiber optics and wave division multiplexing, enable fewer physical links to carry current levels of traffic. The result is a telephone network in which failure of a single link can have serious repercussions.

One might have expected that having multiple phone companies would contribute to increased capacity and diversity in the telephone network. It doesn't. Major telephone companies lease circuits from each other to lower their own costs. This practice means that

backup capacity may not be available when needed. To limit outages, telephone companies have turned to newer technologies. Synchronous Optical Network (SONET) rings, for example, provide redundancy and switch-over at a level below the circuit layer, allowing calls to continue uninterrupted when a fiber is severed. Despite the increased robustness provided by SONET rings, the very high capacity of fiber optic cables results in a greater concentration of bandwidth over fewer paths, because of economic considerations. This means that the failure (or sabotage) of a single link will likely disrupt service for many customers.

The Internet, unlike the PTN, was specifically designed to tolerate link outages. When a link outage is detected, the Internet routes packets over alternate paths. In theory, communications should continue uninterrupted. In practice, though, there may not be sufficient capacity to accommodate the additional traffic on alternate paths. The Internet's routing protocols also do not respond immediately to notifications of link outages. Having such a delay prevents routing instabilities and oscillations that would swamp routers and might otherwise arise in response to transient link outages. But these delays also mean that, although packets are not lost when a link fails, packet delivery can be delayed. In addition to the route-damping noted here, there is a disturbing trend for Internet Service Providers (ISPs) to rely on static configuration of primary and backup routes in border gateway protocol (BGP) routers. This means that Internet routing is less dynamic than was originally envisioned. The primary motivations for this move away from less-constrained dynamic routing are a desire for increased route stability and reduced vulnerability to attacks or configuration errors by ISPs and downstream service providers (DSPs).

Congestion

Congestion occurs when load exceeds available capacity. Environmental disruptions cause increased loads in two ways. First, the load may come from outside the network—people checking by telephone with friends and relatives who live in the area of an earthquake, for example. Second, the load may come from within the network—existing load that is redistributed in order to mask outages caused by the environmental disruption. In both scenarios, network elements saturate and the consequences are an inability to deliver service, perhaps at a time when it is most needed.

The PTN is better able to control congestion than the Internet is. When a phone switch or telephone transmission facility reaches saturation, new callers receive “reorder” (i.e., “fast” busy) signals and no further calls are accepted. This forestalls increased load and congestion. PTN operations staff can even block call attempts to a given destination at sources, thereby saving network resources from being wasted on calls that are unlikely to be completed. For example, when an earthquake occurs near San Francisco, the operations staff might decide to block almost all incoming calls to the affected area codes from throughout the entire PTN.

Congestion management in the Internet is problematic, in part, because no capabilities exist for managing traffic associated with specific users, connections, sources, or destinations, and it would be difficult to implement such capabilities. All that an Internet router can do³² is to discard packets when its buffers become full. To implement fair allocation of resources and bandwidth, routers would have to store information about users and connections, something they are not built to do. Retaining such information would require large amounts of storage. Managing this storage would be difficult, because Internet routers do not normally process call tear-down messages. Furthermore, the concept of a “user”—that is, an entity that originates or receives traffic—is not part of the network or transport layers of the Internet protocols. Choking-back load offered by specific hosts (in analogy with PTN “reorder” signals) is also not an option for preventing Internet congestion, since an Internet protocol (IP)-capable host can concurrently have connections open to many destinations. Stopping all flows from the host is clearly inappropriate. Highly dynamic traffic flows between ISPs are particularly problematic. Here, very high speed (e.g., OC-12) circuits are used to carry traffic between millions of destinations

over short intervals and the traffic mix can completely change over a few seconds.

Although congestion in the Internet is nominally an IP-layer phenomena—routers have too many packets for a given link—measures for dealing successfully with congestion have been deployed in the transmission control protocol (TCP) layer.³³ Some newer algorithms work at the IP level,³⁴ but there is an inadequate knowledge base, especially for defining and enforcing flexible and varied procedures for congestion control. One suggestion involves retaining information about flows from which packets have been repeatedly dropped. Such flows are deemed uncooperative and, as such, are subjected to additional penalties;³⁵ cooperating flows respond to indications of congestion by slowing down their transmissions.

Some work has been done, but very little is known about usage patterns, flow characteristics, and other relevant parameters of current Internet traffic, much less how these patterns may evolve in the future. Having such information is likely to enable better congestion control methods. However, usage patterns are dictated by the application designs, and as new applications arise and become popular, traffic characteristics change. Today, the use of the Web has radically changed packet sizes from a time when file transfer and e-mail were the principal applications. Even within the Web environment, when a very popular Web site arises, news of its location spreads quickly and traffic flows shift noticeably.

There are two further difficulties associated with managing congestion in networks. First, there appears to be a tension between implementing congestion management and enforcing network security. A congestion control mechanism may need to inspect and even modify traffic being managed, but strong network security mechanisms will prohibit reading and modifying traffic en route. For example, congestion control in the Internet might be improved if IP and TCP headers were inspected and modified but the use of IPsec³⁶ will prevent such actions.

A second difficulty arises when a network comprises multiple, independent but interconnected providers. In the Internet, no single party is either capable of or responsible for most end-to-end connections, and local optimizations performed by individual providers may lead to poor overall utilization of network resources or sub-optimal global behavior. In the PTN, which was designed for a world with comparatively few telephone companies but in which switches can be trusted, competitive pressures are now forcing telephone companies to permit widespread interconnections between switches that may not be trustworthy. This opens telephone networks to both malicious and non-malicious failures.³⁷

Human User and Operator Errors in the PTN and Internet

“To err is human” the saying goes, and human operator errors are indeed responsible for network outages as well as for unwittingly disabling protection mechanisms that then enable hostile attacks to succeed. Located in a network operations center, operators take actions based on their perceptions of what the network is doing and what it will do, but without direct knowledge of either. In these circumstances, the consequences of even the most carefully considered operator actions can be surprising—and devastating.

With regard to the PTN, the NRIC found that operational errors caused about one in every four telephone switch failures.³⁸ Mistakes by vendors, mistakes in installation and maintenance, and mistakes by system operators all contributed. The Internet has also been a victim of operational errors, though there does not seem to be the thorough analysis of frequency or specific causes as exists for the PTN. (Anecdotal accounts are numerous.)

Exactly what constitutes an operational error may depend on system capacity. A system operating with limited spare capacity can be especially sensitive to operational missteps. For example, injecting inappropriate, but not technically incorrect, routing information led to a day-long outage of Netcom’s (a major ISP) own internal network in June 1996 as the sheer volume of resulting work overloaded the ISP’s relatively small routers. And this incident may foreshadow problems to come—many routers in the Internet are operating near or at their memory or CPU capacity. It is unclear how well the essential infrastructure of the Internet could cope with a sudden spike in growth rates.

That operator errors are prevalent should not be a surprise. The PTN and Internet are both complex systems. Large numbers of separate and controllable elements are involved in each, and the control parameters for these elements can affect network operation in subtle ways. Operator errors can be reduced when a system

- presents its operators with a conceptual model that allows those operators to predict the effects of their actions and their inaction,³⁹ and
- allows its operators to examine all of the system's abstractions, from the highest to the lowest level, whichever is relevant to the issue at hand.

The entire system must be designed—from the outset—with controllability and understandability as a goal. The reduction of operational errors is more than a matter of building flashy window-based interfaces. The graphics is the easy part. Moreover, with an NIS, there is the added problem of components with different management interfaces provided by multiple vendors. Rarely can the NIS developer change these components or their interfaces, which may make the support of a clean system-wide conceptual model especially difficult.

One approach to reducing operational errors is simply to implement automated support and remove the human from the loop. The route-configuration aids used by PTNs are an example of such automation. More generally, better policy-based⁴⁰ routing mechanisms and protocols will likely free human operators from low-level details associated with setting-up network routes. In the Internet, ISPs currently have just one policy tool: their BGP configurations.⁴¹ But even though BGP is a powerful hammer, the sorts of routing policies that are usually desired do not much resemble nails. Not surprisingly, getting BGP configurations right has proven to be quite difficult.

Finally, operational errors are not only a matter of operators producing the right responses. Maintenance practices—setting up user accounts and access privileges, for example—can neutralize existing security safeguards. And poor maintenance is an oft-cited opening for launching a successful intrusion into a system. The network operations staff at M.I.T., for example, reports that about six weeks after running vulnerability-scan software on a public UNIX workstation, the workstation will again become vulnerable to intrusion as a result of mis-configuration. Managers of corporate or university networks often cite similar problems with firewall and router configuration which, if performed improperly, can lead to access control

violations or denial of service.

System Design and Implementation Errors in the PTN and Internet

The PTN and Internet both experience outages from errors in design and implementation of the hardware and software they employ. A survey by the NRIC⁴² found that software and hardware failures each accounted for about one-quarter of telephone switch outages. This finding is inconsistent with the commonly held belief that hardware is relatively bug-free but software is notoriously buggy. A likely explanation comes from carefully considering the definition of an outage. Within telephone switches, software failures are prone to affect individual telephone calls and, therefore, might not always be counted as causing outages.

Comparable data about actual outages of Internet routers do not seem to be available. One can speculate that routers should be more reliable than telephone switches, because router hardware is generally newer and router software is much simpler. However, against that, one must ask whether routers are engineered and provisioned to the same high standards as telephone switches have been. Moreover, most failures in packet routing are comparatively transient; they are artifacts of the topology changes that routing protocols make to accommodate a failure, rather than being direct consequences of the failure itself.

One thing that is fairly clear is that the Internet's endpoints, including servers for such functions as the DNS, are its least robust components. These endpoints are generally ordinary computers running commercial operating systems and are heir to all of their attendant ills. By contrast, telephony endpoints tend to be either very simple, as in the case of the ordinary telephone, or are built to telephone industry standards.

Even without detailed outage data, it can be instructive to compare the PTN and Internet, since their designs differ in rather fundamental ways and these differences affect how software and hardware failures are handled. The PTN is designed to have remarkably few switches, and it depends on them. That constraint makes it necessary to keep all its switches running virtually all the time. Consequently, switch hardware itself is replicated and the switch software is tasked with detecting hardware and software errors. Upon detecting an error, the software recovers quickly without a serious outage of the switch itself. Individual in-progress calls may be sacrificed, though, to restore the health of the switch.

This approach does not work for all hardware and software failures. That was forcefully illustrated by the January of 1990 failure of the AT&T long distance network. That outage was due to a combination of hardware and software, and the interaction between them:⁴³

The incident began when a piece of trunk equipment failed and notified a switch of the problem. Per its design, the switch took itself offline for a few seconds while it tried to reinitialize the failing equipment; it also notified its neighbors not to route calls to it. When the switch came back on-line, it started processing calls again; neighboring switches were programmed to interpret the receipt of new call setup messages as an indication that the switch had returned to service. Unfortunately, a timing bug in a new version of that process caused those neighboring switches to crash. This crash was detected and (correctly) resulted in a rapid restart—but the failure/restart process triggered the same problem in their neighbors.

A trend that is expected to continue into the future is an increasing reliance on software, rather than on dedicated physical devices, in the management of the PTN. Modern telephony equipment, such as cross-connects and multiplexors, is programmable. A typical “leased line” is simply a programmed path through a series of cross connect boxes. Adjunct processors implement advanced services, such as call forwarding. This increasing use of computer hardware and software in the management of the PTN is a growing vulnerability; if these systems should fail or be penetrated, the reliability of the PTN will suffer. Furthermore, the reliance on familiar systems and protocols, rather than proprietary systems, decreases the learning curve for would-be attackers.

The “switches” for the Internet—its routers—are also intended to be reliable, but they are not designed with the same level of redundancy or error detection as PTN switches. Rather, the Internet as a whole recovers and compensates for router (switch) failures. If a router fails, then its neighbors notice the lack of routing update messages and update their own route tables accordingly. As neighbors notify other neighbors, the failed router is dropped from possible packet routes. In the meantime, retransmissions by endpoints preserve ongoing conversations by causing packets that might have been lost to reenter the network and traverse these new routes.

Attacks by Hostile Parties in the PTN and Internet

Attacks on the PTN and Internet fall into two broad categories, according to the nature of the vulnerability being exploited. First, there are authentication-related attacks. This category includes everything from eavesdroppers' interception of plaintext passwords to designers' misplaced trust in the network to provide authentication. In theory, these attacks can be prevented by proper use of cryptography. The second category of attacks is much harder to prevent. This category comprises attacks that exploit bugs in code. Cryptography cannot help here⁴⁴ nor do other simple fixes appear likely—the design and development of quality software is a long-standing challenge. Yet as long as software does not behave as intended, there will be opportunities for attackers to subvert systems by exploiting unintended system behavior.

Attacks on the Telephone System

Most attacks on the PTN perpetrate toll fraud. The cellular telephony industry provides the easiest target, with caller information being broadcast over unencrypted radio channels and thus easily intercepted.⁴⁵ But attacks have been launched against wireline telephone service as well. Toll fraud probably cannot be prevented altogether. Today, much of it is detected with automated traffic analysis mechanisms that flag for investigation abnormal patterns of calls, credit card authorizations, and other activities.

The NRIC⁴⁶ reports that security incidents have not been a major problem in the PTN until recently. However, the council warns that the threat is growing, for reasons that include (often indirect) interconnections of the computers that run the telephone system (called OSSs) to the Internet, an increase in the number and skill level of attackers, and the increasing number of SS7 interconnections to new phone companies. The report also notes that existing SS7 firewalls

are neither adequate nor reliable in the face of the anticipated threat. As noted earlier, this threat has increased dramatically because of the substantially lower threshold now associated with connection into the SS7 system.

ROUTING ATTACKS

To a would-be eavesdropper, the ability to control call routing can be extremely useful. Installing wiretaps at the endpoints of a connection may be straightforward, but such taps are also the easiest to detect. Interoffice trunks can yield considerably more information to an eavesdropper and with a smaller risk of detection. To succeed here, the eavesdropper first must determine which trunks the target's calls will use, something that is facilitated by viewing or altering the routing tables used by the switches. Second, the eavesdropper must extract the calls of interest from all the calls traversing the trunk; access to the signaling channels can help here.

How easy is it for an eavesdropper to alter routing tables? As it turns out, apart from the usual sorts of automated algorithms which calculate routes based on topology, failed links, or switches, the PTN does have facilities to exert manual control over routes. These facilities exist to allow improved utilization of PTN equipment. For example, there is generally a spike in business calls around 9:00 am on weekdays when workers arrive in their offices. If phone switches in, say, New York are configured to route other East Coast calls through St. Louis or points further west (where the work day has not yet started), then the 9:00 am load spike can be attenuated. However, the existence of this interface for controlling call routing offers a point of entry for the eavesdropper who can profit from exploiting that control.

DATABASE ATTACKS

OSSs and databases they manage are employed to translate telephone numbers so that the number dialed by a subscriber is not necessarily the number that will be reached. Databases are used to implement services such as toll-free numbers, call forwarding, conference calling, hunt groups,⁴⁷ and message delivery. If an attacker can compromise these databases, then various forms of abuse and deception become possible. The simplest such attack exploits network-based speed dialing, a feature that enables subscribers to enter a one or two digit abbreviation and have calls directed to a predefined destination. If the stored numbers are changed by an attacker, then speed-dialed calls could be routed to destinations of the attacker's choice. Beyond harassment, an attacker who can change speed dialing numbers can impersonate a destination or can redial to the

intended destination while staying on the line and eavesdropping. In one successful attack, the database entry for the phone number of the probation office in Del Ray Beach, Florida was reconfigured. People who called the probation office when the line was busy had their calls forwarded to a phone sex line in New York.⁴⁸

Because a subscriber's chosen long-distance carrier is stored in a phone network database, it too is vulnerable to change by attackers. Here the incentive is a financial one—namely, increased market share for a carrier. In a process that has come to be known as slamming, customers' long-distance carriers are suddenly and unexpectedly changed. And this problem has been pervasive enough that numerous procedural safeguards have been mandated by the FCC and various state regulatory bodies.

Looking to the future, more competition in the local telephone market will lead to the creation of a database that enables the routing of incoming calls to specific local telephone carriers. And, given the likely use of shared facilities in many markets, outgoing local calls will need to be checked to see what carrier is actually handling the call. In addition, growing demand for "local number portability," whereby a customer can retain a phone number even when switching carriers, implies the existence of one more databases (which would be run by a neutral party and consulted by all carriers for routing of local calls). Clearly, a successful attack on any of these databases could disrupt telephone service across a wide area.

In contrast to the Internet, the telephone system does not depend on having an automated process corresponding to the Internet's DNS translation from names to addresses.⁴⁹ One doesn't call directory assistance before making every phone call, and success in making a call is not critically dependent on the directory assistance service. Thus, in the PTN, an Internet's vulnerability is avoided but at the price of requiring subscribers to dial phone numbers rather than dialing subscriber names. Furthermore, unlike DNS, the telephone network's directory service is subject to a sanity test by its clients. If a (human) caller asks directory assistance for a neighbor's number and is given an area code for a town halfway across the country, the caller would probably doubt the accuracy of the number and conclude that the directory assistance service was malfunctioning. Still, tampering with directory assistance can cause phone calls to be misdirected.

FACILITIES

The nature of the telephone company physical plant leads to another class of

vulnerabilities. Many central offices are normally unstaffed and, consequently, they are vulnerable to physical penetration, which may go entirely undetected. Apart from the obvious problems of intruders tampering with equipment, the documentation present in such facilities (including, of course, passwords written on scraps of yellow paper and stuck to terminals) is attractive to phone phreaks.⁵⁰ A similar vulnerability is present in less populated rural areas, which are served by so-called remote modules. These remote modules perform local switching but depend on a central office for some aspects of control. Remote modules are invariably deployed in unstaffed facilities, hence subject to physical penetration.

AUTHENTICATION

Authentication is a key part of any scheme for preventing unauthorized activity. In a network containing programmable elements, authentication is an essential ingredient for protecting those elements from performing actions illicitly requested by attackers. Specifically, in the PTN, the OSSs, must be able authenticate requests in order to control changes in the configuration of the elements (cross-connects, multiplexors, etc.) comprising the network. In addition, authentication is required in order to support certain enhanced services, like CallerID.⁵¹ To prevent CallerID from subversion, all elements in the path from the caller to the recipient must be authenticated.

The need for authentication by OSSs is growing because interconnections among previously isolated networks has increased the risk of external intrusions. As the PTN's management networks convert to TCP/IP and are connected to other TCP/IP-based networks, ignoring authentication may prove disastrous. Historically, proprietary protocols and dedicated networks were used for the network's management, so knowledge of these was restricted to insiders and there was little need for authentication or authorization of requests.

Attacks on the Internet

The general accessibility of the Internet makes it a highly visible target and within easy reach of attackers. The widespread availability of documentation and actual implementations for Internet protocols means that devising attacks for this system can be viewed as an intellectual puzzle (where launching the attack validates the puzzle's solution). Internet vulnerabilities are documented extensively on CERT's Web site⁵² and at least one Ph.D. thesis⁵³ is devoted to the subject.

NAME SERVER ATTACKS

The Internet critically depends on the operation of the DNS. Outages or corruption of DNS root servers and other top-level DNS servers—whether due to failure or successful attacks—can lead to denial of service. Specifically, if a top-level server cannot furnish accurate information about delegations of zones to other servers, then clients making DNS lookup requests are prevented from making progress. The client requests might go unanswered or the server could reply in a way that causes the client to address requests to DNS server machines that cannot or do not provide the information being sought. Cache contamination is a second way to corrupt the DNS. An attacker who introduces false information into the DNS cache can intercept all traffic to a specific targeted machine.⁵⁴ One highly visible example of this occurred in July 1997, when somebody used this technique to divert requests for a major Web server to his own machines.⁵⁵

In principle, attacks on DNS servers are easily dealt with by extending the DNS protocols. One such set of extensions, Secure DNS, is based on public key cryptography⁵⁶ and can be deployed selectively in individual zones.⁵⁷ Perhaps because this solution requires the installation of new software on client machines, it has not been widely deployed. No longer merely a question of support-software complexity, the Internet, has grown sufficiently large so that even simple solutions, like Secure DNS, are precluded the sheer number of computers that would have to be modified. A scheme that involved only changing the relatively small number of DNS servers would be quite attractive. But lacking that, techniques must be developed to institute changes in a large-scale and heterogeneous network.

ROUTING SYSTEM ATTACKS

Routing in the Internet is highly decentralized. This avoids the vulnerabilities associated with dependence on a small number of servers that can fail or be compromised. But it leads to other vulnerabilities. With all sites playing some role in routing, there are many more sites whose failure or compromise must be tolerated. The damage inflicted by any single site must somehow be contained, even though each site necessarily serves as the authoritative source for some aspect of routing. Decentralization is thus not a panacea for avoiding the vulnerabilities intrinsic in centralized services. Moreover, the trustworthiness of most NISs will, like the Internet, be critically dependent both on services that are more sensibly implemented in a centralized fashion (e.g., DNS) and on services more sensibly implemented in a decentralized way (e.g., routing). Understanding how either type of services can be made trustworthy is thus instructive.

The basis for routing in the Internet is each router periodically informing neighbors about what networks it knows how to reach. This information is direct when a router advertises the addresses of the networks to which it is directly connected. More often, though, the information is indirect, with the router relaying to neighbors what it has learned from others. Unfortunately, recipients of information from a router rarely can verify its accuracy⁵⁸ since, by design, a router's knowledge about network topology is minimal. Virtually any router can represent itself as a best path to any destination as a way of intercepting, blocking, or modifying traffic to that destination.⁵⁹

Most vulnerable are the interconnection points between major ISPs, where there are no grounds at all for rejecting route advertisements. Even an ISP that serves a customer's networks cannot reject an advertisement for a route to those networks via one of its competitors—many larger sites are connected to more than one ISP.⁶⁰ Such multi-homing thus becomes a mixed blessing, with the need to check accuracy, which causes traffic addressed from a subscriber net arriving via a different path to be suspect and rejected, being pitted against the increased availability that multi-homing promises. Some ISPs are now installing BGP policy entries that define which parts of the Internet's address space neighbors can provide information about (with secondary route choices). However, this approach undermines the Internet's adaptive routing and affects overall survivability.

Somehow, the routing system must be secured against false advertisements. One

approach is to authenticate messages a hop at a time. A number of such schemes have been proposed,⁶¹ and a major router vendor (Cisco) has selected and deployed one in products. Unfortunately, the “hop at a time” approach is limited to ensuring that an authorized peer has sent a given message; nothing ensures that the message is accurate. The peer might have received an inaccurate message (from an authorized peer) or might itself be compromised. Thus, some attacks are prevented but others remain viable.

The alternative approach for securing the routing system against false advertisements is, somehow, for routers to employ global information about the Internet’s topology. Advertisements that are inconsistent with that information are thus rejected. Some schemes have been proposed, but these do not appear to be practical for the Internet. Perlman’s scheme,⁶² for example, requires source-controlled routing over the entire path. Routing protocol security is an active research area, and deserves continued support.

It is worth noting that the routing system of the Internet closely mirrors call routing in the PTN, except that, in the PTN, a separate management and control network carries control functions. Any site on the Internet can participate in the global routing process, whereas subscribers in the PTN do not have direct access to the management and control network. The added vulnerabilities of the Internet derive from this lack of isolation. As network interconnections increase within the PTN, it may become vulnerable to the same sorts of attacks as the Internet now is.

PROTOCOL DESIGN AND IMPLEMENTATION FLAWS

The design and implementation of many Internet protocols make them vulnerable to a variety of denial of service attacks.⁶³ Some attacks exploit buggy code. These are perhaps the easiest to deal with; affected sites need only install newer or patched versions of the affected software. Other attacks exploit artifacts of particular implementations, such as limited storage areas, expensive algorithms, and the like. Again, updated code often can cure such problems.

The more serious class of attacks exploit features of certain protocols. For example, one type of attack exploits both the lack of source address verification and the connectionless nature of user datagram protocol (UDP)⁶⁴ to bounce packets between query servers on two target hosts.⁶⁵ This process can continue almost indefinitely, until a packet happens to be dropped. And, while the process continues, computation and network bandwidth are consumed. The obvious remedy would be for hosts to detect this attack or any such denial of service attack,

much the same way virus screening software detects and removes viruses. But if it is cheaper for an attacker to send a packet than it is for a target to check it, then denial of service is inevitable from the sheer volume of packets. Even cryptography is not a cure: authenticating a putatively valid packet is much harder (it requires substantial CPU resources) than generating a stream of bytes with a random authentication check value to send the victim.⁶⁶

AUTHENTICATION (AND OTHER SECURITY PROTOCOLS)

Concern about strong and useable authentication in the Internet is relatively new. The original Internet application protocols used plaintext passwords for authentication, a mechanism that was adequate for casual logins, but was insufficient for more sophisticated uses of a network, especially in a local area network environment. Rather than build proper cryptographic mechanisms—which were little known in the civilian sector at that time—the developers of the early Internet software for UNIX resorted to network-based authentication for remote login and remote shell commands. The servers checked their clients' messages by converting the sender's IP address into a host name. User names in such messages are presumed to be authentic if the message comes from a host whose name is trusted by the server. Senders, however, can circumvent the check by misrepresenting their IP address⁶⁷ (something that is more difficult with TCP).

But cryptographic protocols—a sounder basis for network authentication and security—are now growing in prominence on the Internet. Link-layer encryption has been in use for many years. It is especially useful when just a few links in a network need protection. (In the latter days of the ARPANET, MILNET trunks outside of the continental United States were protected by link encryptors.) Although link-layer encryption has the advantage of being completely transparent to all higher-layer devices and protocols, the scope of its protection is limited. Accordingly, attention is now being focused on network-layer encryption. Network-layer encryption requires no modification to applications, and it can be configured to protect host-to-host, host-to-network, or network-to-network traffic. Cost thus can be traded against granularity of protection.

Network-layer encryption is instantiated in the Internet as IPsec, which is designed to run on the Internet's hosts, routers, or on hardware outboard to either. The initial deployment of IPsec has been in network-to-network mode. This mode allows virtual private networks to be created so that the otherwise-insecure Internet can be incorporated into an existing secure

network, such as a corporate net. The next phase of deployment for IPsec will most likely be the host-to-network mode, with individual hosts being laptops or home machines. That would provide a way for travelers to exploit the global reach of the Internet in order to access a secure corporate net.

It is unclear when general host-to-host IPsec will be widely deployed. Although transparent to applications, IPsec is not transparent to system administrators—the deployment of host-to-host IPsec requires outboard hardware or modifications to the host's protocol system software and that constitutes a significant impediment to deployment. Because of the impediments to deploying IPsec, the biggest use of encryption in the Internet is currently above the transport layer, as the secure socket layer (SSL) is embedded into popular Web browsers and servers. SSL, though quite visible to its applications, affects only those applications and not the kernel or the hardware. SSL can be deployed without supervision by a central authority, the approach used for almost all other successful elements of Internet technology.

Higher still in the protocol stack, encryption is found in fairly widespread use for the protection of electronic mail messages. In this manner, an email message is protected during each Simple Mail Transfer Protocol,⁶⁸ while spooled on intermediate mail relays, while residing in the user's mailbox, while being copied to the recipient's machine, and even in storage thereafter. However, no secure email format has been both standardized by the Internet Engineering Task Force (IETF) and accepted by the community. Two formats that have gained widespread support are S/MIME⁶⁹ and PGP.⁷⁰ Both have been submitted to the IETF for review. Thus, cryptography has experienced some success in improving the level of authentication in the Internet. The continued deployment of cryptography should be encouraged.

A LOOK TO THE FUTURE

Internet Telephony

What are the trustworthiness implications if, as predicted by many pundits, today's traditional telephone network is replaced by an Internet-based transport mechanism? Will telephony become even less secure, due to all the security problems with the Internet discussed

earlier in this chapter? Or will some portion of the Internet that is used only for telephony be resistant to many of the problems described in the preceding sections?

Recall that many current PTN vulnerabilities are related to either the services being provided or to the physical transport layer. Rehosting the PTN on the Internet will have no effect on these vulnerabilities. Thus, the OSSs and database lookups related to enhanced PTN services, with their associated vulnerabilities would be unaffected by the move to an Internet-based phone system. Similarly, if access to the Internet-based phone system is accomplished by means of twisted pairs (albeit twisted pairs carrying something like Integrated Services Digital Network (ISDN) or Asymmetric Digital Subscriber Line (ADSL)), then multiplexors and cross-connects of some sort will still be needed. They would likely be replaced by routers or switches, but these replacements would be at least as programmable, and at least as vulnerable.

Call routing in an Internet-based phone system would be different, but likely no more secure. At the very least, IP routing would be involved. Most probably, a new database would be introduced to map telephone numbers to domain names or IP addresses. Both, of course, raise serious trustworthiness concerns.

In at least two respects, both noted earlier in this chapter, an Internet-based phone system could be significantly more vulnerable to attack than today's PTN. The primary active elements of an Internet-based network—the routers—are, by design, accessible from the network they control, and the network's routing protocols execute in-band with the communications they control. By contrast, virtually the entire PTN is now managed by out-of-band channels. Considerable care will be needed to deliver the security of out-of-band control by using in-band communications. The other obvious weakness of the Internet is its endpoints, PCs and servers, because then attacks on them can be used to attack the phone system.

Is the Internet Ready for "Prime Time"?

Whether the Internet is "ready for business" depends on the requirements of the business. There already are numerous examples of businesses using the Internet for advertising, marketing, sales of products and services, coordination with business partners, and various other infrastructure activities. On the other hand, the Internet is also viewed—and rightly so—as being less reliable and less secure than the PTN. Specifically, the Internet is perceived as more

susceptible to interception (i.e., eavesdropping) and has proved to be more susceptible to active attacks (e.g., server flooding, Web site modification). Consequently, most Internet-savvy business users restrict what transactions they entrust to the Internet.

The Internet is also more prone to outages than the PTN. Thus, it would be unwise for utility companies and other critical infrastructure providers to abandon the PTN and rely on remote access through the Internet for controlling power distribution substations, because individual ISPs are less likely to survive local power interruptions than individual telephone companies.⁷¹

Few established business seem willing to forgo their telephone order centers for Internet-only access, although a small and growing number of newer businesses, such as Virtual Vineyards and Amazon.com, do maintain an Internet-only presence. Abandoning the PTN for the Internet seems unwise for businesses, such as brokerage houses or mail-order catalog companies, where continued availability of service is critical. For example, during the October 27-28, 1997, stock market frenzy, customers of Internet-based brokerage systems experienced unusual delays in executing trades. But the magnitude of their delays was relatively small and was commensurate with the delays suffered by telephone-based access and even some of the stock market's back-end systems. Still, it is sobering to contemplate the effect of an Internet-related failure that coincided with a spike in market activity.

Mail-order firms, brokerage houses, and others do make extensive use of the Internet as an avenue of customer access. But it is not the only avenue of access and neither the customers nor the business have become wholly dependent on this avenue. If, for example, these and similar business reduced their other avenues of access (e.g., to save money), then an Internet outage could have a significant impact. Consider a scenario in which banks acquire the capability to download customer money onto smart cards through the Internet. Over time, banks might reduce the number of automatic teller machines available (just as the numbers of physical bank branches and tellers have fallen as automated teller machines have proliferated). A prolonged failure of this Internet cash distribution mechanism could overload the few remaining available machines and tellers.

In theory, the risks associated with using the Internet can be evaluated and factored into a risk management model.⁷² Most businesses, however, are not fully cognizant of these risks nor of the return on investments in protection. As a result, the level of protection adopted by many business users of the Internet does not seem commensurate with that afforded their physical assets. For example, it seems as though the quality of burglar alarms and physical access control systems deployed by most businesses is considerably higher than the level of Internet security countermeasures they deploy.⁷³

Moreover, businesses that make extensive use of Internet technology may do so in a fashion that externalizes the risks associated with such use. If infrastructure suppliers, such as phone companies and electric and gas utilities, do not take adequate precautions to ensure the availability of their systems in the face of malicious attacks over the Internet, then the public will bear the brunt of the failure. Because many of these businesses operate in what is effectively a monopoly environment, the free market forces that should eventually correct such cost externalization may not be effective.

Of particular concern is that most of the security countermeasures adopted by business connecting to the Internet are designed only to thwart the most common attacks used by hackers. Most hackers, however, are opportunistic and display only a limited repertoire of skills. Protection against that hacker threat is insufficient for warding off more capable, determined threats, such as criminals or terrorists sponsored by national governments.

And while in one sense the Internet poses no new challenges—a system that can be

accessed from outside only through a cryptographically-protected channel on the Internet is at least as secure as the same system reached through a conventional leased line—new dangers arise precisely because of pervasive interconnectivity. The capability to interconnect networks gives the Internet much of its power; by the same token, it opens up serious new risks. An attacker who may be deflected by cryptographic protection of the front door can often attack a less protected administrative system and use its connectivity through internal networks to bypass the encryption unit protecting the real target.

¹ Fred B. Schneider is the chair of the Committee on Information Systems Trustworthiness of the Computer Science and Telecommunications Board, National Research Council that produced *Trust in Cyberspace*, National Academy Press, Washington, DC (1998). Schneider is supported in part by ARPA/RADC grant F30602-96-1-0317 and AFOSR grant F49620-94-1-0198. The views and conclusions contained herein are those of the author and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied of these organizations or the U.S. government. Author's address: Computer Science Department, Cornell University, Ithaca, New York 14853, fbs@cs.cornell.edu.

² Steven M. Bellovin is a member of the Committee that produced *Trust in Cyberspace*. Author's address: AT&T Laboratories—Research, 180 Park Avenue, Room E-215, Florham Park, NJ 07932, smb@research.att.com.

³ Alan S. Inouye is the program officer for the Committee that produced *Trust in Cyberspace*. Author's address: National Research Council, Computer Science and Telecommunications Board, 2101 Constitution Avenue, N.W., Washington, DC 20418, ainouye@nas.edu.

⁴ See *Cybernation: The American Infrastructure in the Information Age: A Technical Primer on Risks and Reliability*, Executive Office of the President, Office of Science and Technology, Washington, DC (1997); Reports from the Eight NSTAC Subcommittee Investigations, National Security Telecommunications Advisory Committee, Tysons Corner, VA (1997); *Critical Foundations: Protecting America's Infrastructures*, President's Commission on Critical Infrastructure Protection, Washington, DC (1997); *Report of the Defense Science Board Task Force on Information Warfare Defense (IW-D)*, Defense Science Board, Washington, DC (1996); and *Information Security—Computer Attacks at Department of Defense Pose Increasing Risks: A Report to Congressional Requesters*, General Accounting Office, Washington, DC (1996).

⁵ Such as testimony before the Senate Governmental Affairs Committee on "Weak Computer Security in Government: Is the Public at Risk," May 19, 1998, and testimony before the Senate Armed Services Committee on "Future Threats to the Department of Defense Information Systems: Y2K & Frequency Spectrum Reallocation," June 4, 1998.

⁶ See URL=<http://www.ciao.gov>

⁷ In the computer security literature, vulnerability, attack, and threat are technical terms. A vulnerability is an error or weakness in the design, implementation, or operation of the system. An attack is a means of exploiting some vulnerability in a system. And a threat is an adversary that is motivated and capable of exploiting a vulnerability.

⁸ The classification and restricted distribution of many government studies about vulnerability and frequency of hostile attacks, rather than informing the public about real risks, serves mostly to encourage speculation.

⁹ For example, see Neumann, Peter, *Computer Related Risks*, ACM Press, New York (1995) and the Computer Emergency Response Team (CERT)/Coordination Center, an element of the Networked Systems Survivability Program in the Software Engineering Institute at Carnegie Mellon University. See URL=<http://www.cert.org>. Also see "Attacks by Hostile Parties" later in this sub-section.

¹⁰ See Ware, Willis H., *The Cyber-posture of the National Information Infrastructure*, RAND Critical Technologies Institute, Washington, DC (1998).

¹¹ See *Computer Related Risks* (see note 9).

¹² See Neumann, Peter, "Rats Take Down Stanford Power and Silicon Valley Internet Service," *RISKS Digest*, 18(52), (1996).

¹³ The most important function of the DNS is to map host names, such as www.nas.edu, into numeric IP addresses. DNS also maps IP addresses into host names and performs other tasks.

¹⁴ See Wayner, Peter, "Human Error Cripples the Internet," *New York Times*, July 17, 1997.

¹⁵ SS7 is a protocol suite used for communication with, and control of, telephone central office switches and processors. It provides out-of-band signaling.

¹⁶ See Towson, Peter, "AT&T Database Glitch Caused '800' Phone Outage," *Telecom Digest*, 17(253), 1997.

¹⁷ See *Trust in Cyberspace* (see note 1) for an extended discussion.

¹⁸ See *Computer Related Risks* (see note 9).

¹⁹ See Mills, Mike, "AT&T High Speed Network Fails Red Cross, Banks Scramble to Adjust," *Washington Post*, April 14, p. C01, (1998).

²⁰ See Associated Press, "Phone Outages Affect East Coast," *Associated Press*, June 12, 1998; and Kalish, David E., "Mishaps

Cause Phone Outage in East,” *Associated Press*, June 12, 1998.

²¹ See Milton, Pat, “FBI Director Calls for Effort to Fight Growing Danger of Computer Crime,” *Associated Press*, March 4, 1997.

²² See Boston Globe, “Youth Faces Computer Crime Charges U.S. Attorney Says Federal Case is First Involving a Juvenile,” *Boston Globe*, March 18, 1998.

²³ Specifically, defense installations reported 53 attacks in 1992, 115 in 1993, 255 in 1994, and 559 in 1995, from *Information Security—Computer Attacks at the Department of Defense Pose Increasing Risks: A Report to Congressional Requesters*, General Accounting Office, Washington, DC, May 1996.

²⁴ See *Information Security—Computer Attacks at the Department of Defense Pose Increasing Risks: A Report to Congressional Requesters* (see note 23).

²⁵ See Hardy, Quentin, “Many Big Firms Hurt by Break-ins,” *Wall Street Journal*, November 21, 1996, p. B4; Power, Richard G., *Testimony Richard G. Power, Editor, Computer Security Institute (CSI)*, Before the Permanent Subcommittee on Investigations of the U.S. Senate, Washington, DC, June 5, 1996; and War Room Research LLC, 1996 *Information Systems Security Survey*, War Room Research LLC, Baltimore, MD (1996).

²⁶ See Gertz, Bill, “‘Infowar’ Game Shutdown U.S. Power Grid, Disabled Pacific Command,” *Washington Times*, April 16, 1998, p. A1; Myers, Laura, “Pentagon Has Computers Hacked,” *Associated Press*, April 16, 1998.

²⁷ See Sweet, William, and Linda Geppert, eds., “Main Event: Power Outages Flag Technology Overload, Rule-making Gaps,” *IEEE Spectrum* 1997 Technology Analysis and Forecast (1997).

²⁸ See *Computer Related Risks* (see note 9).

²⁹ In March 1997, DISA disclosed that a contract had been awarded to Sprint for a global telecommunications network designed primarily to carry signal intelligence data to Fort Meade (Brewin, Bob, “DISA Discloses Secret NSA Pact with Sprint,” *Federal Computer Week*, March 10, 1997). And, according to the *Report of the Defense Science Board Task Force on Information Warfare Defense* (see note 4), the U.S. government procures over 95 percent of its domestic telecommunications network services from U.S. commercial carriers.

³⁰ See Network Reliability and Interoperability Council (NRI), *Final Report of the Network Reliability and Interoperability Council*, Federal Communications Commission, Washington, DC (1997).

³¹ For both the proposed text and the letter to Congress, see URL=<http://www.fcc.gov/oet/nric>.

³² In fact, routers can transmit an ICMP (Internet Control Message Protocol) Source Quench message to advise a host of congestion, but there has never been a standard, accepted response to receipt of a Source Quench, so many hosts merely ignore such messages. In such circumstances the resources needed to construct and send the Source Quench, may be wasted, and may compound the problem!

³³ See Jacobson, V., “Congestion Avoidance Control,” *SIGCOMM* 88, Stanford, CA (1988).

³⁴ See Floyd, S., and V. Jacobson, “Random Early Detection Gateways for Congestion Avoidance,” *IEEE/ACM Transactions on Networking*, 1(4): pp. 397-413 (1993).

³⁵ See Floyd, S., and K. Fall, “Promoting the Use of End-to-End Congestion Control in the Internet,” *IEEE Transactions on Networking* (1998).

³⁶ Network-layer encryption is instantiated in the Internet as IPsec, which is designed to run on the Internet’s hosts, routers, or on hardware outboard to either.

³⁷ See the *Final Report of the Network Reliability and Interoperability Council* (see note 30).

³⁸ See Network Reliability and Interoperability Council (NRI), *Network Reliability: The Path Forward*, Federal Communications Commission, Washington, DC (1996).

³⁹ See Wickens, Christopher D., Anne S. Mavor, and James P. McGee, eds., *Flight to the Future: Human Factors in Air Traffic Control*, National Academy Press, Washington, DC (1997); Parasuraman, Raja, and Mustapha Mouloua, eds., *Automation and Human Performance: Theory and Applications*, Edited by Bary H. Kantowitz, *Human Factors in Transportation*, Lawrence Erlbaum Associates, Mahwah, NJ (1996).

⁴⁰ Policy-based routing recognizes constraints on a path other than topology and bandwidth. For example, an ISP may carry transit traffic – traffic not destined for its own customers – from some of its peers, but not from others, depending on business arrangements.

⁴¹ See Rekhter, Y., and T. Li, *A Border Gateway Protocol 4 (BGP-4)*, RFC 1771 (1995); Rekhter, Y., and P. Gross, *Application of the Border Gateway Protocol in the Internet*, RFC 1772 (1995); Traina, P., *Experience with the BGP-4 Protocol*, RFC 1773 (1993); and Traina, P., *BGP-4 Protocol Analysis*, RFC 1774 (1995).

⁴² See *Network Reliability: The Path Forward* (see note 38).

⁴³ Based on Cooper, Brinton. 1989. “Phone Hacking,” *RISKS Digest*, 8(79). Downloaded from URL=<http://catless.ncl.ac.uk/Risks/8.79.html#subj4>.

⁴⁴ See Blaze, Matt, “Afterword,” Edited by Bruce Schneier, 2nd ed., Published in *Applied Cryptography*, John Wiley, New York (1996).

⁴⁵ See Computer Science and Telecommunications Board, *The Evolution of Untethered Communications*, National Academy Press, Washington, DC (1997).

⁴⁶ See the *Final Report of the Network Reliability and Interoperability Council* (see note 30).

⁴⁷ A hunt group is like a switchboard. It enables one phone number to lead to multiple phone lines.

⁴⁸ There is even an historical precedent for such attacks. The original phone switch was invented by an undertaker; his competitor's wife was a telephone operator who connected anyone who asked for a funeral home to her husband's business (Cooper, Brinton, "Phone Hacking," *RISKS Digest*, 8(79), (1989)).

⁴⁹ This is not strictly true, since calls to certain classes of telephone numbers (e.g., 800, 888, and 900) do result in a directory lookup to translate the called number into a "real" destination phone number. In these instances, the analogy between the PTN and the Internet is quite close.

⁵⁰ A "phone phreak" is a telephone network hacker.

⁵¹ CallerID is an enhanced service that identifies the originator of a telephone call to a (suitably equipped) receiver. As this service becomes more pervasive, it will be more and more used for identification and authentication by systems employing the telephone network for communications. Here, then, is a vulnerability that can propagate from a communications fabric into an NIS that is built on top of that fabric.

⁵² The Computer Emergency Response Team (CERT)/Coordination Center is an element of the Networked Systems Survivability Program in the Software Engineering Institute at Carnegie Mellon University. See URL=<http://www.cert.org>.

⁵³ See Howard, John D., "An Analysis of Security Incidents on the Internet 1989-1995," Ph.D. thesis, Carnegie Mellon University, Pittsburgh, PA (1995).

⁵⁴ See Bellovin, Steven M., "Security Problems in the TCP/IP Protocol Suite," *Computer Communications Review*, 19(2): pp. 32-48 (1989).

⁵⁵ See "An Internet Stunt Causes Trouble for Kashpureff," *The Wall Street Journal*, Nov. 4, 1997.

⁵⁶ See Eastlake, D., and C. Kaufman, *Domain Name System Security Extensions*, RFC 2065 (1997).

⁵⁷ However, configuration management does become much harder when there is partial deployment of Secure DNS.

⁵⁸ In a few cases it actually is possible to reject inaccurate information. For example, an ISP will know what network addresses belong to its clients, and neighbors of such a router generally will believe that and start routing traffic to the ISP.

⁵⁹ See "Security Problems in the TCP/IP Protocol Suite."

⁶⁰ The percentage of such multi-homed sites in the Internet is currently low but appears to be rising, largely as a reliability measure by sites that cannot afford to be offline.

⁶¹ Such as Badger, M. R., and S. L. Murphy, "Digital Signature Protection of the OSPF Routing Protocol," published in *Proceedings of the Symposium on Network and Distributed System Security*, February, in San Diego, CA, pp. 93-102 (1996); Hauser, R., T. Przygienda, and G. Tsodik, "Reducing the Cost of Security in Link-State Routing," published in *Proceedings of the Symposium on Network and Distributed System Security*, February, in Los Alamitos, CA (1997); Sirois, K. E., and Stephen T. Kent. 1997. Securing the Nimrod Routing Architecture, published in *Proceedings of the Annual Internet Society (ISOC) Symposium on Network and Distributed System Security*, February, in Los Alamitos, CA, pp. 74-84 (1997); and Smith, B.R., S. Murthy, and J.J. Garcia-Luna-Aceves. "Securing Distance-Vector Routing Protocols," published in *Proceedings of the Annual Internet Society (ISOC) Symposium on Network and Distributed System Security*, February, in Los Alamitos, CA, pp. 85-92 (1997).

⁶² See Perlman, Radia, "Network Layer Protocols with Byzantine Robustness," Ph.D. thesis, MIT, Cambridge, MA (1988).

⁶³ See Schuba, Christoph L., Ivan Krsul, Markus G. Kuhn, Eugene H. Spafford, Aurobindo Sundaram, and Diego Zamboni, "Analysis of a Denial of Service Attack on TCP," *Proceedings of 1997 IEEE Symposium on Security and Privacy*, Oakland, CA, pp. 208-233 (1997).

⁶⁴ UDP is an Internet transport protocol.

⁶⁵ See CERT Advisory CA-96.01

⁶⁶ Encryption is even worse in this regard, as the cost of decryption is often greater than the cost of authentication and because a receiver might have to both decrypt and authenticate a packet to determine if it is valid. The Encapsulating Security Payload (ESP) protocol of IPsec counters this denial of service vulnerability by reversing the order in which these operations are applied (i.e., a receiver authenticates ciphertext prior to decrypting it).

⁶⁷ A number of different attacks on this scheme are known. This can be accomplished in a number of ways, such as sequence number guessing (Morris, Robert T., *A Weakness in the 4.2 BSD UNIX TCP/IP Software*, AT&T Bell Laboratories (1985)) or route corruption (Bellovin, Steven M., "Security Problems in the TCP/IP Protocol Suite"). Alternatively, the attacker can target the address-to-name translation mechanism (Bellovin, Steven M., "Using the Domain Name System for System Break-ins," *Proceedings of the 5th USENIX/UNIX Security Symposium*, Salt Lake City, UT, pp. 199-208 (1995)).

⁶⁸ See Postel, J., *Simple Mail Transfer Protocol*, RFC 821 (1982).

⁶⁹ See Dusse, S., P. Hoffman, and B. Ramsdell, *S/MIME Version 2 Certificate Handling*, RFC 2312 (1998); and Dusse, S., P. Hoffman, b. Ramsdell, L. Lundblade, and L. Repka, *S/MIME Version 2 Message Specification*, RFC 2311 (1998).

⁷⁰ See Zimmerman, Philip R., *The Official PGP User's Guide*, MIT Press, Cambridge, MA (1995).

⁷¹ Internet service providers have differing plans for dealing with power system failures, which may make it impossible to access computers and data following such a failure. The failure need not even be widespread. By contrast, phone networks are under central control, can easily implement backup power systems, and require very little electrical current for an ordinary phone line.

⁷² See *Trust in Cyberspace* (see note 1) for a discussion.

⁷³ See *Trust in Cyberspace* (see note 1).