

BUILDING TRUSTWORTHY SYSTEMS:

Lessons from the PTN and Internet

FRED B. SCHNEIDER

Cornell University

STEVEN M. BELLOVIN

AT&T Labs—Research

ALAN S. INOUE

Computer Science and Telecommunications Board

National and economic infrastructures are coming to depend on networked information systems, or NISs. These systems must be trustworthy—do what users and operators expect (and not something else) despite environmental disruption, human user and operator errors, attacks by hostile parties, and system design and implementation errors. Economics dictates the use of commercial off the shelf (COTS) components wherever possible, so NIS developers have neither control nor detailed information about many of their system's components. Moreover, the increasing use of components whose functionality can be extended after deployment (“plug-and-play” and other extensible operating system features) means users and designers of an NIS cannot know what software has entered system components or what actions those components might take.

Trustworthiness is a holistic property, encompassing security (conventionally including confidentiality, integrity, and availability), correctness, reliability, privacy, safety, and survivability. It is not sufficient to address only some of these diverse dimensions, nor is it sufficient simply to assemble components that are themselves trustworthy. Integrating the components and understanding how the trustworthiness dimensions interact is a central challenge in building a trustworthy NIS.

The public telephone network (PTN) and the Internet are two large, complex NISs that shed light on the technical problems faced by the developers and operators of other NISs. In some ways, the two networks are very similar. No single entity owns, manages, or can even have a complete picture of either. Both networks involve large numbers of subsystems operat-

The PTN and Internet are two large and complex networked information systems. Studying the vulnerabilities of these systems can help identify ways new research might eliminate those vulnerabilities.

ed by different organizations. The number and intricate nature of the interfaces that exist at the boundaries of these subsystems are one source of complexity; the increasing popularity of advanced services is a second source.

The high cost of building a global communications infrastructure from the ground up implies that one or both of these networks is likely to furnish communications services for most other NISs. An understanding of each network's vulnerabilities therefore informs the assessment of the trustworthiness for other NISs. For example, the Internet uses leased telephone lines as its physical transport medium while telephone companies increasingly employ Internet technology (though not necessarily the Internet itself)

Some ISPs are relying on static route configuration, making Internet routing less dynamic, hence less robust, than was originally envisioned.

to manage their own facilities. Thus, vulnerabilities in the PTN can affect the Internet and vulnerabilities in Internet technology can affect the PTN.

This article discusses vulnerabilities in the PTN and Internet, identifying ways new research might help to eliminate those vulnerabilities. The article is based on an excerpt from *Trust in Cyberspace*, the final report of the Committee on Information Systems Trustworthiness of the Computer Science and Telecommunications Board, National Research Council.¹

ENVIRONMENTAL DISRUPTIONS

Environmental disruptions range from earthquakes and storms to more localized incidents, like rodents chewing through cable insulation and accidents caused by human carelessness. The effects and, to some extent, impact of these different types of disruptions differ across the two networks.

Link Failures

The single biggest cause of PTN outages is damage to buried cables. And the single biggest cause of this damage is construction crews digging without prop-

er clearance from telecommunications companies and other utilities. The phenomenon, known in the trade as "backhoe fading," is probably not amenable to a technological solution. Its impact 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 PTN 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. Many companies compete in today's telephone market, and cost pressures make it impractical for these 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 higher levels of traffic. Failure of a single link can now have serious repercussions. Moreover, to lower costs, major telephone companies lease circuits from each other, and backup capacity thus may not be available when needed.

To limit outages, telephone companies have turned to newer technologies such as Synchronous Optical Network (Sonet) rings. Sonet rings provide redundancy and switch-over at a level below the circuit layer, allowing calls to continue uninterrupted when a fiber is severed. However, despite the increased robustness they provide, the very high capacity of fiber optic cables results in a greater concentration of bandwidth over fewer paths. 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 an Internet link outage is detected, packets are routed 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 do not respond immediately to notifications of link outages. This delay prevents routing instabilities, although it may also delay delivery of packets. But a disturbing trend has been for Internet service providers (ISPs) to rely on static configuration of primary and backup routes so that Internet routing becomes less dynamic, hence less robust, 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.

Congestion

Congestion occurs when load exceeds available capacity. Increased 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. A load increase may also come from within the network—existing load that is redistributed to mask outages caused by an environmental disruption. In both scenarios, network elements saturate and service is impaired.

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” (that is, “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. The PTN is also capable of load-sensitive routing. New calls can be routed via alternate paths based on calling patterns at a given time. Experiments with load-sensitive routing in the Internet have been unsuccessful, probably because flows are too short and traffic patterns too chaotic as compared with longer-lived and constant bit rate telephone calls.

Congestion management in the Internet is problematic, in part because no capabilities exist for managing traffic associated with specific users, connections, sources, or destinations. An Internet router can only 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 that would be expensive and difficult to do. 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. Nor is choking back the load offered by specific hosts (analogous to PTN “reorder” signals) an option, since an IP-capable host can have concurrent 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 (such as OC-12) circuits are used to carry traffic between millions of destinations over short intervals, and the traffic mix can completely change within 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. Some newer algorithms work at the IP level, but the knowledge base is inadequate here, 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;² cooperating flows respond to indications of congestion by slowing down their transmissions.

Today's Internet would have more trouble coping with the requirements for a voice channel than the PTN does.

Having more information about usage patterns, flow characteristics, and other relevant parameters of current Internet traffic, as well as how these patterns may evolve in the future, is likely to improve congestion control methods. However, usage patterns are dictated by application designs, and as new applications become popular, traffic characteristics change. For example, the growth of the Web has resulted in packets that are much larger than they were when file transfer and e-mail were the principal applications.

The Internet does have one advantage over the PTN with regard to congestion control because it supports different grades of service. The PTN offers just one service: a 56-Kbps channel with guaranteed bandwidth. Many uses of the Internet require much less bandwidth, and most protocols and applications will automatically adapt. Congestion control in that sense is automatic; end systems react to congestion by reducing the load they offer the network. Note, though, that this is as much a property of applications as of the network. Today's Internet would have more trouble coping with the requirements for a voice channel than the PTN does.

There are two further difficulties associated with managing congestion in networks. First, is 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

PREVENTING UNAUTHORIZED ACTIVITY IN THE INTERNET

Concern about strong and usable 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 insufficient for more sophisticated uses of a network, especially in a LAN 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 client 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 address (something that is more difficult with TCP).

Cryptographic protocols—a sounder basis for network authentication and security—are now gaining 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 U.S. 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, which requires no modification to applications, and 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 (VPNs) to be created so that the otherwise insecure Internet can be incorporated into an existing secure network, such as a corporate intranet. 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 allow travelers to exploit the global reach of the Internet to access a secure corporate intranet.

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 IPSec¹, the biggest use of encryption in the Internet is currently above the transport layer, as the secure socket layer (SSL²) 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 used to protect e-mail messages. An e-mail message is encrypted during each Simple Mail Transfer Protocol (SMTP), 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 e-mail format has been both standardized by the IETF and accepted by the community. Two formats that have gained widespread support are S/MIME³ and PGP⁴, both of which have been submitted to the IETF for review.

REFERENCES

1. S. Kent and R. Atkinson, "Security Architecture for the Internet Protocol," Internet Engineering Task Force RFC 2401, Nov. 1998; available online at [ftp://ftp.ietf.org/rfc/rfc2401.txt](http://ftp.ietf.org/rfc/rfc2401.txt).
2. A. O. Freier, P. Karlton, and P. Kocher, "The SSL Protocol Version 3.0," Internet Draft, 1996, work in progress; available online at <http://www.netscape.com/eng/ssl3/ssl-toc.html>.
3. B. Ramsdell, ed., "S/MIME Version 3 Message Specification," IETF RFC 2633, June 1999; available online at [ftp://ftp.ietf.org/rfc/rfc2633.txt](http://ftp.ietf.org/rfc/rfc2633.txt).
4. J. Callas et al., "OpenPGP Message Format," IETF RFC 2440, Nov. 1998; available online at [ftp://ftp.ietf.org/rfc/rfc2440.txt](http://ftp.ietf.org/rfc/rfc2440.txt).

the Internet might be improved if IP and TCP headers were inspected and modified, but the use of IPSec will prevent such actions. (for more on IPSec and other Internet security issues, see the sidebar "Preventing Unauthorized Activity in the Internet.")

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 suboptimal 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 nonmalicious failures.

USER AND OPERATOR ERROR

“To err is human” the saying goes, and human operator errors are indeed responsible for network outages as well as unwittingly disabling protection mechanisms that might then allow hostile attacks to succeed. Located in a network operations center, operators base their actions 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 devastating.

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. Many routers in the Internet are operating near or at their memory or CPU capacity, and how well the essential infrastructure of the Internet could cope with a sudden spike in growth rates is unclear. Aggressive use of topology-based address assignments in support of Classless Interdomain Routing (CIDR) has slowed the growth of routing tables.³ But this strategy is threatened by increased use of multihoming, whereby a host increases availability by connecting to more than one ISP.

Reducing operational errors requires more than building flashy window-based interfaces. Large numbers of separate and controllable elements are involved in both the PTN and the Internet, and the control parameters for these elements can affect network operation in subtle ways. Thus, to reduce operator errors, the entire system must be designed from the outset with controllability and understandability as a goal. A further difficulty is that an NIS typically will be built with components from multiple vendors and, therefore, will have many different management interfaces. Rarely can the NIS developer change these components or their interfaces, which makes the support of a clean system-wide conceptual model especially difficult.

One approach to reducing operational errors is to implement automated support and remove the human from the loop. More generally, better policy-based 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 (Border Gateway Protocol) configurations. 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. Poor maintenance practices—for example, setting up user accounts and access privileges—can neutralize existing security safeguards. Such practices can open the door to a successful intrusion into a system.

DESIGN AND IMPLEMENTATION ERRORS

A survey by the Network Reliability and Interoperability Council (NRIC) found that software and hardware failures each accounted for about one-quarter of telephone switch outages.⁴ Comparable data about actual outages of Internet routers do not seem to be available. We 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, we must ask whether routers are engineered and provisioned to the same high standards as telephone switches. 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 direct consequences of the failure itself.

One thing that is fairly clear is that the Internet's end points, including servers for such functions as the Domain Naming Service (DNS), are its least robust components. These end points are generally ordinary computers running commercial operating systems and are heir to all of their attendant ills. By contrast, telephony end points tend to be either very simple, as in the case of the ordinary telephone, or 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 must detect 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. That this

approach does not work for all hardware and software failures was forcefully illustrated by the January 1990 failure of the AT&T long-distance network, which was due to a combination of hardware and software, and the interaction between them.⁵

The PTN is expected to continue increasing its reliance on software, rather than on dedicated physical devices. Modern telephony equipment, such as cross-connects and multiplexers, 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. 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 Internet's routers are also intended to be reliable but 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 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 end points preserve ongoing conversations by causing packets that might have been lost to reenter the network and traverse these new routes.

ATTACKS BY HOSTILE PARTIES

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 here 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 tar-

get, with caller information being broadcast over unencrypted radio channels and thus easily intercepted. But attacks have been launched against wireline telephone service as well.

The NRIC reports that security incidents have not been a major problem in the PTN until recently. However, the NRIC warns that the threat is growing, for reasons that include (often indirect) interconnections of the computers that run the telephone system (called operations support systems, or OSSs) to the Internet, an increase in the number and skill level of attackers, and the increasing number of Signaling System 7 (SS7) interconnections to new phone companies. The NRIC report also notes that existing SS7 firewalls are neither adequate nor reliable in the face of the anticipated threat. 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 end points of a connection might 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 that calculate routes based on topology, failed links, or switches, the PTN has facilities to exert manual control over routes. These facilities exist to allow improved utilization of PTN equipment; however, they can also offer a point of entry for eavesdropping and other types of attacks.

Database attacks. OSSs translate telephone numbers and manage databases with which they implement services such as toll-free numbers, call forwarding, conference calling, hunt groups, and message delivery. If an attacker can compromise the 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. Attackers can

change the stored numbers, rerouting speed-dialed calls to destinations of their choice, which can then facilitate the attacker's eavesdropping.

Because a subscriber's choice of long-distance carrier is stored in a phone network database, it too is vulnerable to attack. 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. This problem has been pervasive enough in the U.S. that numerous procedural safeguards have been mandated by the FCC and various state regulatory bodies.

Increased 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 need for one or more databases (which would be run by a neutral party and consulted by all carriers for routing local calls). Clearly, a successful attack on any of these databases could disrupt telephone service across a wide area.

The telephone system does not depend on an automated process like the Internet's DNS translation from names to addresses. Most people don't call directory assistance before making every phone call, and success in making a call is not 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 subscriber names.

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 checks the puzzle's solution. Internet vulnerabilities are documented extensively on CERT's Web site (<http://www.cert.org>) and at least one PhD thesis is devoted to the subject.⁶

Name server attacks. The Internet 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 targeted machine.

Decentralization is not a panacea for avoiding the vulnerabilities intrinsic in centralized services.

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 cryptography and can be deployed selectively in individual zones.⁷ Perhaps because this solution requires the installation of new software on client machines, it has not been widely deployed. Protecting DNS servers from attack is no longer merely a question of software complexity: the Internet has grown sufficiently large so that even simple solutions like Secure DNS are precluded by the sheer number of computers that would have to be modified. A scheme that involved changing only 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 leads to other vulnerabilities. With all sites playing some role in routing, the failure or compromise of some sites must be tolerated. 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, depend both on services that are more sensibly implemented in a centralized fashion (such as DNS) and on services more sensibly implemented in a decentralized way (such as routing). Understanding how either type of service 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 accuracy 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.

Most vulnerable are the interconnection points between major ISPs, where there are no grounds 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—larger sites can be connected to more than one ISP. Such multihoming 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 multihoming 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, and Cisco has selected and deployed one in its routers. 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 and others remain viable.

An alternative for securing the routing system against false advertisements is 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, for example, requires source-controlled routing over the entire path.⁸

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.

Denial-of-service attacks. Flaws in the design and implementation of many Internet protocols make them vulnerable to a variety of denial-of-service attacks. 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) to bounce packets between query servers on two target hosts. This process can continue almost indefinitely, until a packet is dropped. Moreover, the process consumes computation and network bandwidth. 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, the sheer volume of packets can make denial of service inevitable. 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.

CONCLUDING REMARKS: INTERNET TELEPHONY?

A "trust gap" is emerging between the needs and expectations of the public and the capabilities of today's network information systems. The PTN and Internet exemplify the trend and even lead it in some ways, as the oft debated possibility of

replacing the traditional telephone network by an Internet transport mechanism illustrates.

To start, rehosting the PTN on the Internet leaves intact the many vulnerabilities related to either the services being provided or to the physical transport layer. And although call routing in an Internet-based phone system would be different, it would involve IP routing along with a new database to map telephone numbers, both of which raise new trustworthiness concerns. Furthermore, 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 using in-band communications. The other obvious weakness of the Internet is its end points, PCs and servers, because then attacks on them can be used to attack the phone system.

Looking beyond the PTN and Internet, other instances of a "trust gap" are causing headlines with some frequency. More troubling, though, is that this gap can only widen—we lack the science and technology base to build trustworthy NISs. *Trust in Cyberspace*, the basis of this article, takes a necessary first step by identifying technical problems and articulating an agenda for research to solve those problems. ■

ACKNOWLEDGMENTS

Schneider is supported in part by ARPA/RADC grant F30602-96-1-0317, AFOSR grant F49620-94-1-0198, DARPA and Air Force Research Laboratory/Air Force Material Command under agreement number F30602-99-1-0533, and National Science Foundation grant 9703470. The views and conclusions contained herein should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of these organizations or the U.S. government.

REFERENCES

1. F.B. Schneider, ed., *Trust in Cyberspace*, National Academy Press, Washington, D.C., 1999; available online at <http://www.nap.edu/readingroom/books/trust/>
2. S. Floyd and K. Fall, "Promoting the Use of End-to-End Congestion Control in the Internet," Submitted to *IEEE Trans. Networking*, Feb. 1998; available online at <ftp://ftp.ee.lbl.gov/papers/collapse.feb98.ps>.
3. V. Fuller et al., "Classless Inter-Domain Routing (CIDR): An Address Assignment and Aggregation Strategy," Internet Engineering Task Force RFC 1519, Sept. 1993; available online at <ftp://ftp.ietf.org/rfc/rfc1519.txt>.
4. Network Reliability and Interoperability Council (NRIC), *Network Reliability: The Path Forward*. Federal Communications Commission, Washington D.C., 1996; available online at <http://www.fcc.gov/oet/info/orgs/nric/>.
5. B. Cooper, "Phone Hacking," *RISKS Digest*, Vol. 8, No. 79, June 1989; available at <http://catless.ncl.ac.uk/Risks/8.79.html#subj4>.
6. J.D. Howard, "An Analysis of Security Incidents on the Internet 1989-1995," Ph.D. thesis, Carnegie Mellon Univ., Pittsburgh, Penn., 1995.
7. D. Eastlake and C. Kaufman, "Domain Name System Security Extensions," IETF RFC 2065, Jan. 1997; available online at <ftp://ftp.ietf.org/rfc/rfc2065.txt>.
8. R. Perlman, "Network Layer Protocols with Byzantine Robustness," Ph.D. thesis, MIT, Cambridge, Mass., 1988.

Fred B. Schneider is a professor of computer science at Cornell University, Ithaca, New York. He has an MS and Ph.D. from SUNY Stony Brook and a BS from Cornell. In addition to chairing the National Research Council's study committee on information systems trustworthiness and editing *Trust in Cyberspace*, Schneider is managing editor of *Distributed Computing*, co-managing editor of Springer-Verlag's Texts and Monographs in Computer Science, and serves on a number of editorial boards. He holds patents in fault-tolerant system design and is a founding member of Sun's Java Security Advisory Council. Schneider is a fellow of AAAS and ACM and is Professor at Large at University of Tromso, Norway.

Steven M. Bellovin is an AT&T Fellow at Bell Labs. He received a BA from Columbia University, and an MS and PhD in computer science from the University of North Carolina at Chapel Hill. Bellovin is the co-author of the book *Firewalls and Internet Security: Repelling the Wily Hacker* and holds several patents on cryptographic and network protocols. He served on a National Research Council study committee on information systems trustworthiness, is a member of the Internet Architecture Board, and currently focuses on designing systems that are inherently more secure.

Alan S. Inouye is a study director and program officer for the Computer Science and Telecommunications Board of the National Research Council in Washington, D.C. He received a PhD from the University of California, Berkeley, School of Information Management and Systems. Inouye has particular interests in improving the access to digital government information, adapting copyright for electronic publishing, and understanding the impact of information technology on work and workplaces.

Readers can contact Schneider at the Computer Science Department, Cornell University, Ithaca, New York 14853, fbs@cs.cornell.edu.