United States General Accounting Office

GAO

Report to the Chairman, Subcommittee on Telecommunications and Finance, Committee on Energy and Commerce, House of Representatives

June 1989

COMPUTER SECURITY

Virus Highlights Need for Improved Internet Management

Accession Number: 1921

Publication Date: Jun 12, 1989

Title: Computer Security: Virus Highlights Need for Improvement Internet Management

Personal Author: Hoy, J.B.; Brewer, M.T.; Peterson, B.A.; Dittmer, G.; et al.

Corporate Author Or Publisher: U.S. General Accounting Office, GAO, Washington, DC 20548 Report Number: GAO/IMTEC-89-57

Descriptors, Keywords: Policy Network Virus Security Research Telecommunication Violation Detection Threat Computer Vulnerability

Pages: 048

Cataloged Date: Sep 11, 1989

Document Type: HC

Number of Copies In Library: 000001

Record ID: 20792

GAO

United States
General Accounting Office
Washington, D.C. 20548

Information Management and
Technology Division

B-233721

June 12, 1989

The Honorable Edward J. Markey
Chairman, Subcommittee on Telecommunications
and Finance
Committee on Energy and Commerce
House of Representatives

Dear Mr. Chairman:

This report responds to your October 14, 1988, request and subsequent agreements with your office to (1) describe the Internet virus incident, (2) examine issues relating to Internet security and vulnerabilities, and (3) discuss factors affecting the prosecution of computer virus crimes. The report contains recommendations to the President's Science Advisor, Office of Science and Technology Policy, aimed at improving security through the creation of an Internet security focal point.

As arranged with your office, unless you publicly announce the contents of this report earlier, we plan no further distribution until 30 days after the date of this letter. At that time, we will send copies to the appropriate House and Senate Committees, the five key federal agencies involved in Internet research networks, the National Institute of Standards and Technology, the National Security Agency, and other interested parties. This report was prepared under the direction of Jack L. Brock, Director. Major contributors are listed in appendix III.

Sincerely yours,

Ralph V. Carlone

Ralph V. Carlone
Assistant Comptroller General

Executive Summary

Purpose

In November 1988, a computer program caused thousands of computers on the Internet—a multinetwork system connecting over 60,000 computers nationwide and overseas—to shut down. This program, commonly referred to as a computer virus or worm, entered computers and continuously recopied itself, consuming resources and hampering network operations.

Concerned about Internet security and the virus incident, the Chairman, Subcommittee on Telecommunications and Finance, House Committee on Energy and Commerce, asked GAO to

- provide an overview of the virus incident,
- examine issues relating to Internet security and vulnerabilities, and
- describe the factors affecting the prosecution of computer virus incidents.

Background

The Internet, the main computer network used by the U.S. research community, comprises over 500 autonomous unclassified national, regional, and local networks. Two of the largest networks are sponsored by the National Science Foundation and the Department of Defense. In addition, three other agencies operate research networks on the Internet. Over the past 20 years, the Internet has come to play an integral role in the research community, providing a means to send electronic mail, transfer files, and access data bases and supercomputers.

There is no lead agency or organization responsible for Internet-wide management. Responsibility for computer security rests largely with the host sites that own and operate the computers, while each network is managed by the network's sponsor, such as a federal agency, university, or regional consortium.

Plans are for the Internet to evolve into a faster, more accessible, larger capacity network system called the National Research Network. The initiative to upgrade the Internet—described as a "super highway" for the research community—stems from a report by the Office of Science and Technology Policy. This Office, headed by the President's Science Advisor, has a broad legislative mandate to coordinate and develop federal science policy.

In recent years, the public has become increasingly aware of computer virus-type programs that can multiply and spread among computers. The Internet virus differed from earlier viruses (which primarily

attacked personal computers) in that it was the first to use networks to spread, on its own, to vulnerable computer systems.

Federal laws exist that address computer crimes, but none are specifically directed at virus-type incidents. In addition, 48 states have enacted laws dealing with computer crime.

Results in Brief

Within hours after it appeared, the Internet virus had reportedly infected up to 6,000 computers, clogging systems and disrupting most of the nation's major research centers. After 2 days, the virus was eradicated at most sites, largely through the efforts of university computer experts. After the virus incident, multiple intrusions (not involving viruses) at several Internet sites added to concerns about security.

These incidents highlighted such vulnerabilities as (1) the lack of an Internet focal point for addressing security issues, (2) security weaknesses at some sites, and (3) problems in developing, distributing, and installing software fixes (i.e., repairs to software flaws).

While agencies and groups have taken actions to enhance security, GAO believes that many of the vulnerabilities highlighted by the virus and subsequent intrusions require actions transcending those of individual agencies or groups. For this reason, GAO believes a security focal point should be established to fill a void in Internet's management structure.

Several factors may hinder successful prosecution of virus-type incidents. For example, since there is no federal statute that specifically makes such conduct a crime, other laws must be applied. In addition, the technical nature of such cases may hinder prosecution.

Principal Findings

Internet Virus Incident

The onset of the virus was extremely swift. Within an hour after it appeared, the virus was reported at many sites, and by early morning, November 3, thousands of computers were infected at such sites as the Department of Energy's Lawrence Livermore National Laboratory, the National Aeronautics and Space Administration's Ames Research Center, the Massachusetts Institute of Technology, Purdue University, and Cornell University.

The virus spread over networks largely by exploiting (1) two holes (flaws) in systems software used by many computers on the networks and (2) weaknesses in host site security policies, such as lax password management.

The primary effects of the virus were lost computer processing and staff time. However, while apparently no permanent damage was done, a few changes to the virus program could have resulted in widespread damage and compromise of sensitive or private information.

Vulnerabilities Highlighted

The lack of an Internet security focal point created difficulties in responding to the virus. For example, problems were reported in communicating information about the virus to sites, coordinating emergency response activities, and distributing fixes to eradicate the virus.

The virus also exploited security weaknesses at some sites. For example, the incident showed that some sites paid insufficient attention to security issues, such as proper password usage, and lacked system management expertise for dealing with technical issues.

In addition, problems were highlighted in developing, distributing, and installing software fixes for known flaws. For example, vendors are not always timely in repairing software holes that may create security vulnerabilities. Further, even when fixes are available, sites may not install them, through either neglect or lack of expertise. In the subsequent intrusions, intruders entered several computer systems by exploiting a known software hole. In one case, the vendor had not supplied the fix for the hole, and in the other, the fix was supplied but not installed.

Since the virus incident, agencies and groups have taken actions, such as creating computer emergency response centers and issuing ethics statements to heighten users' moral awareness. These actions are an important part of the overall effort needed to upgrade Internet security. However, GAO believes that a focal point is needed to provide the oversight, coordination, and policy-making capabilities necessary to adequately address the Internet's security vulnerabilities. Because no one organization is responsible for Internet-wide management and the Office of Science and Technology Policy has taken a leadership role in initiating plans for a National Research Network, GAO believes that the Office would be the most appropriate body to coordinate the establishment of a security focal point.

Prosecution Problems

To prosecute computer virus-type incidents on the federal level, such laws as the Computer Fraud and Abuse Act of 1986 (18 U.S.C. 1030) or the Wire Fraud Act (18 U.S.C. 1343) might be used. However, the 1986 act, the law most closely related to computer virus-type cases, is untried with respect to virus-type incidents, and contains terms that are not defined. Also, the evidence in such cases tends to be highly technical, which may hinder prosecution.

Recommendations

To help ensure the necessary improvements to Internet-wide security are achieved, GAO recommends that the President's Science Advisor, Office of Science and Technology Policy, coordinate the establishment of an interagency group, including representatives from the agencies that fund research networks on the Internet, to serve as the Internet security focal point. This group should

- provide Internet-wide security policy, direction, and coordination;
- support ongoing efforts to enhance Internet security;
- obtain the involvement of Internet users, software vendors, technical advisory groups, and federal agencies regarding security issues; and
- become an integral part of the structure that emerges to manage the National Research Network.

Agency Comments

As requested, GAO did not obtain official agency comments on this report. However, the views of officials from the Defense Department, the National Science Foundation, and the Office of Science and Technology Policy were obtained and incorporated in the report where appropriate.

Contents

Abbreviations

CERT	Computer Emergency Response Team
DARPA	Defense Advanced Research Projects Agency
FCCSET	Federal Coordinating Council on Science, Engineering and Technology
FRICC	Federal Research Internet Coordinating Committee
GAO	General Accounting Office
HHS	Department of Health and Human Services
IMTEC	Information Management and Technology Division
MIT	Massachusetts Institute of Technology
NASA	National Aeronautics and Space Administration
NCSC	National Computer Security Center
NIST	National Institute of Standards and Technology
NSF	National Science Foundation
OSTP	Office of Science and Technology Policy
PC	personal computer

Introduction

On Wednesday, November 2, 1988, a virus[1] appeared on the Internet, the main computer network system used by U.S. researchers. The virus reportedly infected up to 6,000 computers, consuming resources and hampering network operations. The Internet, an unclassified, multi-network system connecting over 500 networks and over 60,000 computers nationwide and overseas, has come to play an integral role within the research community. A user on any one of the thousands of computers attached to any Internet network can reach any other user and has potential access to such resources as supercomputers and data bases. This chapter presents an overview of the Internet—how it evolved, how it is used and managed, and what plans there are for its further development—as well as a description of the events surrounding the Internet virus.

Internet Evolves From an Experimental Network

The Internet began as an experimental, prototype network called Arpanet, established in 1969 by the Department of Defense's Defense Advanced Research Projects Agency (DARPA). Through Arpanet, DARPA sought to demonstrate the possibilities of computer networking based on packet-switching technology.[2] Subsequently, DARPA sponsored several other packet-switching networks. In the 1970s, recognizing the need to link these networks, DARPA supported the development of a set of procedures and rules for addressing and routing messages across separate networks. These procedures and rules, called the "Internet protocols," provided a universal language allowing information to be routed across multiple interconnected networks.

From its inception, Arpanet served as a dual-purpose network, providing a testbed for state-of-the-art computer network research as well as network services for the research community. In the 1980s, the number of networks attached to Arpanet grew as technological advances facilitated network connections. By 1983 Arpanet had become so heavily used that Defense split off operational military traffic onto a separate

[1]Although there is no standard definition, technical accounts sometimes use the term "worm" rather than "virus" to refer to the self-propagating program introduced on November 2. The differences between the two are subtle, the essential one being that worms propagate on their own while viruses, narrowly interpreted, require human involvement (usually unwitting) to propagate. However, their effects can be identical. We have chosen to use the term virus in deference to popular use.

[2]Packet switching is a technique for achieving economical and effective communication among computers on a network. It provides a way to break a message into small units, or packets, for independent transmission among host computers on a network, so that a single communication channel can be shared by many users. Once the packets reach their final destination, they are reassembled into the complete message.

system called Milnet, funded and managed by the Defense Communications Agency. Both Arpanet and Milnet are unclassified networks. Classified military and government systems are isolated and physically separated from these networks.

Building on existing Internet technology, the National Science Foundation (NSF), responsible for nurturing U.S. science infrastructure, fostered the proliferation of additional networks. In 1985, NSF made the Internet protocols the standard for its six supercomputing centers and, in 1986, funded a backbone network—NSFnet—linking the six centers.[3] NSF also supported a number of regional and local area campus networks whose network connections were facilitated through NSF funding.[4] As of September 1988, there were about 290 campus networks connected to NSFnet through about 13 regional networks. Many of these networks also connect to Arpanet.

Other federal agencies fund research networks. The Department of Energy, the National Aeronautics and Space Administration (NASA), and the Department of Health and Human Services (HHS) operate networks on the Internet that support their missions.

This loosely organized web of interconnected networks—including Arpanet, Milnet, NSFnet, and the scores of local and regional networks that use the Internet protocols—make up the Internet. The Internet supports a vast, multi-disciplinary community of researchers, including not only computer scientists but physicists, electrical engineers, mathematicians, medical researchers, chemists, and astronomers.

Researchers use the Internet for a variety of functions; electronic mail, which provides a way of sending person-to-person messages almost instantaneously, is the most frequent use. Using electronic mail, researchers separated by thousands of miles can collaborate on projects, sharing results and comments daily. Other uses of the Internet include file transfer and remote access to computer data banks and supercomputers. Access to supercomputers has had a dramatic impact on scientific endeavors; experiments that took years to complete on an ordinary computer can take weeks on a supercomputer. Currently, use of the

[3]A backbone network is a network to which smaller networks are attached. Arpanet and Milnet are also backbone networks.

[4]Regional networks include partial-statewide networks (e.g., Bay Area Regional Research Network in northern California), statewide networks (e.g., New York State Educational Research Network), and multistate networks (e.g., Southern Universities Research Association Network).

Internet is generally free-of-charge to individuals engaged in government-sponsored research.

Rapid Growth of the Internet

The Internet's transition from a prototype network to a large-scale multinetwork has been rapid, far exceeding expectations. In the past 5 years, its growth has been particularly dramatic. For example:

- In late 1983, the Internet comprised just over 50 networks; by the end of 1988, the number had grown to over 500.
- In 1982, about 200 host computers were listed in a network data base; by early 1987, there were about 20,000, and by early 1989 the number exceeded 60,000.[5]
- An October 1988 NSF network publication estimated that there were over half a million Internet users.[6]

Funding for Internet operations comes from the five agencies (DARPA, NSF, Energy, NASA, and HHS) involved in operating research networks and from universities, states, and private companies involved in operating and participating in local and regional networks. A 1987 Office of Science and Technology Policy (OSTP) report estimated federal funding to be approximately $50 million. A national information technology consortium official estimated that university investments in local and regional networks are in the hundreds of millions of dollars; state investments are estimated in the millions and rapidly growing.[7]

Management in a Decentralized Environment

Management of the Internet is decentralized, residing primarily at the host site and individual network levels. Early in the Internet's development, responsibility for managing and securing host computers was given to the end-users—the host sites, such as college campuses and federal agencies, that owned and operated them. It was believed that the host sites were in the best position to manage and determine a level of security appropriate for their systems. Further, DARPA's (Arpanet's developer and the major federal agency involved in the Internet in its early years) primary function was in fostering research in state-of-the-art technology rather than operating and managing proven technology.

[5]Host computers, which include supercomputers, mainframes, and minicomputers, are the machines, attached to the networks, that run application programs.

[6]NSF Network News, No. 5, NSF Network Service Center, Oct. 1988.

[7]Industry also invests in local and regional networks; however, the amount of that investment could not be determined.

At each host site, there may be many host computers.[8] These computers are controlled by systems managers who may perform a variety of security-related functions, including

- establishing access controls to computers through passwords or other means;
- configuration management, enabling them to control the versions of the software being used and how changes to that software are made;
- software maintenance to ensure that software holes (flaws) are repaired; and
- security checks to detect and protect against unauthorized use of computers.

Operational Management at the Network Level

Each of the Internet's more than 500 networks maintains operational control over its own network, be it a backbone network (such as NSFnet), a regional network, or a local area network. Distributed responsibility allows for use of different technologies as well as different types of administration. Each network is autonomous and has its own operations center that monitors and maintains its portion of the Internet. In addition, some of the larger networks maintain information centers that provide information on network use and resources.

No Internet-Wide Management

No one agency or organization is responsible for overall management of the Internet. According to a DARPA official, decentralization provided the needed flexibility for the Internet's continuing growth and evolution. Within the Internet, networks operated by government agencies serve as backbones to connect autonomous regional and local (campus) networks. Agency backbone networks were established with agency missions in mind, and their structures and modes of operation generally reflect individual agency philosophies.

In the fall of 1987, representatives of the five federal agencies—DARPA, NSF, Energy, NASA, HHS—that operate Internet research networks joined forces to form the Federal Research Internet Coordinating Committee (FRICC). The objectives of this informal group include coordinating network research and development, facilitating resource sharing, reducing operating costs, and consolidating requirements for international connections of the participating agencies. Currently, FRICC is involved in developing plans to upgrade the Internet and improve services.

[8]For example, at the University of California, Berkeley, there are over 2,000 host computers.

Future of the Internet

The Internet, long characterized by growth and change, is evolving into an enhanced, upgraded system to be called the National Research Network. Plans are for the enhanced network system to serve as a superhighway that would run faster, reach farther, and be more accessible than any other computer network system in the world.

The National Research Network will include a number of high-speed networks, including NSFnet, Defense Research Internet, and other research networks funded by NASA, Energy, and HHS.[9] The networks will use a shared, cross-country, high-capacity link called the Research Interagency Backbone.

The initiative for an upgraded network stemmed from two high-level studies prepared by the Office of Science and Technology Policy and an ad hoc committee of the National Research Council.[10] OSTP has a broad mandate to coordinate and develop federal science policy. Within OSTP, the Congress established the Federal Coordinating Council on Science, Engineering and Technology (FCCSET) to initiate interagency consideration of broad national issues and coordinate government programs.

Both studies noted the critical importance of a modern, high-speed research network in providing for research and technology development. They concluded that current network technology did not adequately support scientific collaboration and that U.S. networks, commercial and government-sponsored, were not coordinated, had insufficient capacity, and did not assure privacy. The studies recommended that a national research network be established to improve network capabilities. The Chairman of the FCCSET Subcommittee on Networking has asked FRICC to develop a coordinated, multi-agency implementation plan for the National Research Network.

FRICC has taken some initial steps toward upgrading the Internet. FRICC's NSF representative has agreed to take the lead in organizing the National Research Network, coordinating multiagency efforts and the development of long-term management plans. In early 1989, NSF sent out a request for proposals to provide and manage the Research Interagency Backbone.

[9]Within the next few years, Arpanet will be replaced as an all-purpose network by NSFnet. A Defense Research Internet will be created for experimental work in computer networking.

[10]A Research and Development Strategy for High Performance Computing, Office of Science and Technology Policy (Washington, D.C., Nov. 1987), and Toward a National Research Network, National Research Network Review Committee, National Academy Press (Washington, D.C., 1988).

Internet Virus Spread Over Networks to Vulnerable Computers

The Internet virus, which entered computers and continuously recopied itself, was not the first virus-type program to infect computers. However, it differed from earlier viruses in several key respects. First, previous viruses were almost always limited to personal computers (PCs), whereas the Internet virus infected larger systems, such as minicomputers, workstations, and mainframes. In addition, the Internet virus was the first to spread over a network automatically (i.e., without requiring other programs or user intervention to transmit it).

The networks themselves (i.e., the communications hardware and software that connect the computer systems) were not infected by the virus; rather, they served as a roadway enabling the virus to spread rapidly to vulnerable computers. In transit, the virus was indistinguishable from legitimate traffic and, thus, could not be detected until it infected a computer. The principal symptoms of the virus were degradation of system response and loss of data storage space on file systems.

How the Virus Spread

The Internet virus spread largely by exploiting security holes in systems software based on the Berkeley Software Distribution UNIX system and by taking advantage of vulnerabilities in host site security policies.[11] UNIX is the most commonly used operating system on the Internet—a University of California, Berkeley, researcher estimated that about three-quarters of the computers attached to the Internet use some version of UNIX. Machines infected were VAX and Sun-3 computer systems.[12]

The virus propagated by using four methods of attack:[13]

Sendmail: A utility program that handles the complex tasks of routing and delivering computer mail. The virus exploited a "debug" feature of sendmail that allowed a remote operator to send executable programs. After issuing the debug command, the virus gave orders to copy itself.

[11]UNIX is a registered trademark of AT&T Laboratories. Berkeley distributes its own version of UNIX, and a number of other systems manufacturers have selected the Berkeley UNIX version as the basis for their own operating systems. The virus did not attack the operating system's "kernel" that manages the system; rather, it exploited flaws in peripheral service or utility programs.

[12]VAX and Sun-3 computers are built by Digital Equipment Corporation and Sun Microsystems, Inc., respectively.

[13]See appendix I for a more detailed account of the security flaws the virus exploited.

Fingerd: A utility program that allows users to obtain public information about other users, such as a user's full name or telephone extension. A hole in the program allowed the virus to propagate to distant machines.

Passwords: The virus tried different methods to guess user passwords. Once the virus gained access through a correct password, it could masquerade as a legitimate user and exercise that user's privileges to gain access to other machines.

Trusted hosts: Trusted host features provide users convenient access to each other's resources. This is not a software hole; it is a convenience sometimes used on local networks where users frequently use services provided by many different computers. By using these features, the virus spread quickly within local networks once one computer had been penetrated.

Chronology of the Virus

The onset of the virus was extremely swift. The first reports of the virus came from several sites at 9 p.m., Eastern Standard Time, on Wednesday, November 2. An hour later, the virus was reported at multiple Internet sites, and by early morning, November 3, the virus had infected thousands of computer systems.

Most of the nation's major research centers were affected, including Energy's Lawrence Livermore National Laboratory; NASA's Ames Research Center; the University of California, Berkeley; the Massachusetts Institute of Technology (MIT); Carnegie Mellon University; Cornell University; Purdue University; and many others. The virus also affected sites on Milnet and several overseas sites. As noted earlier, the Internet is an open, unclassified network; the virus did not affect classified government or operational military systems.

Once the virus was detected, many sites disconnected their computers from the Internet, leaving only one or two computers running to communicate with other sites and to permit study of virus activity. By Thursday, November 3, the sendmail and fingerd holes had been identified, and by late that night, the Computer Systems Research Group at the University of California, Berkeley, had posted patches on network bulletin boards to mend the holes.[14]

[14]A patch is a modification made to an object program. Patches to the sendmail hole had been posted on Thursday morning.

By Friday evening, the virus had been eliminated at most sites. At a November 8 virus post-mortem conference, hosted by the National Security Agency's National Computer Security Center (NCSC), attendees concluded that the virus had been analyzed and eradicated by computer science experts located primarily at university research institutions, with U.S. government personnel playing a small role.

Objectives, Scope, and Methodology

In response to an October 14, 1988, request of the Chairman, Subcommittee on Telecommunications and Finance, House Committee on Energy and Commerce, and subsequent agreements with his office, the objectives of our review were to

- describe the virus incident,
- examine issues relating to Internet security and vulnerabilities, and
- discuss factors affecting the prosecution of computer virus incidents.

In addition, we sought to identify federal research directed specifically at viruses and to provide an overview of research that may improve security on open networks, such as the Internet.

To understand the nature, structure, and management of the Internet and to determine events surrounding the Internet virus and related security issues, we reviewed:

- Reports, analyses, and briefings prepared by NCSC, DARPA, the Defense Communications Agency, NSF, NASA, and the Department of Energy.
- Academic analyses prepared by individuals associated with MIT, Purdue University, and the University of Utah.
- Accounts of the virus and its aftermath in scientific publications, industry journals, and newspapers.

We discussed the virus incident, implications of an open network environment, security issues, the need for increased centralized management, and the National Research Network with:

- Officials from the agencies listed above as well as from the National Institute of Standards and Technology (NIST), OSTP, FCCSET, FRICC, the Office of Management and Budget, and the General Services Administration.
- Officials representing systems software vendors, including the Computer Systems Research Group of the University of California, Berkeley; Sun Microsystems, Inc.; and Digital Equipment Corporation.

- Network users representing federal and academic sites, including Harvard University, MIT, NASA's Ames Research Center, Energy's Lawrence Livermore National Laboratory, and the University of California, Berkeley.
- Officials from private sector security companies in the Washington, D.C., area and California and from SRI, International, which operates the Defense-funded Network Information Center.

To obtain a perspective on factors affecting the prosecution of computer virus offenses, we discussed the relevant laws with officials of the Federal Bureau of Investigation, Department of Justice, and Secret Service. We also discussed these issues with representatives of the Colorado Association of Computer Crime Investigators and the University of Colorado's Computer Law Center.

We discussed research aimed at improving computer and open network security with officials from government agencies and systems software vendors cited above; with members of the Internet Activities Board, a technical group concerned with Internet standards; and with officials from Bolt, Beranek, and Newman, Inc., which maintains Arpanet's Network Operations Center. We did not develop a complete inventory of current research, nor did we evaluate its potential effectiveness.

Our work was performed in accordance with generally accepted government auditing standards. We performed our work primarily between November 1988 and March 1989 in Washington, D.C., and at research institutions and vendor locations in Massachusetts and California. We discussed the contents of a draft of this report with DARPA, NSF, and OSTP officials, and their comments have been incorporated where appropriate. However, as requested, we did not obtain official agency comments.

Virus Focuses Attention on Internet Vulnerabilities

Although the virus spread swiftly over the networks to vulnerable computers, it apparently caused no permanent damage. However, the virus highlighted vulnerabilities relating to (1) the lack of a focal point for responding to Internet-wide security problems, (2) host site security weaknesses, and (3) problems in developing, distributing, and installing software fixes. A number of agencies and organizations have taken actions since the virus to address identified problems. However, we believe that these actions alone will not provide the focus needed to adequately address the Internet's security vulnerabilities.

Impact of Virus

The virus caused no lasting damage; its primary impact was lost processing time on infected computers and lost staff time in putting the computers back on line. The virus did not destroy or alter files, intercept private mail, reveal data or passwords, or corrupt data bases.

No official estimates have been made of how many computers the virus infected, in part because no one organization is responsible for obtaining such information. According to press accounts, about 6,000 computers were infected. This estimate was reportedly based on an MIT estimate that 10 percent of its machines had been infected, a figure then extrapolated to estimate the total number of infected machines. However, not all sites have the same proportion of vulnerable machines as MIT. A Harvard University researcher who queried users over the Internet contends that a more accurate estimate would be between 1,000 and 3,000 computers infected.

Similar problems exist in trying to estimate virus-related dollar loss. The total number of infected machines is unknown, and the amount of staff time expended on virus-related problems probably differed at each site. The Harvard University researcher mentioned earlier estimated dollar losses to be between $100,000 and $10 million.

Estimated losses from individual sites are generally not available. However, NASA's Ames Research Center and Energy's Lawrence Livermore National Laboratory, two major government sites, estimated their dollar losses at $72,500 and $100,000, respectively. These losses were attributed primarily to lost staff time.

Although the virus is described as benign because apparently no permanent damage was done, a few changes to the virus program could have resulted in widespread damage and compromise, according to computer experts. For example, these experts said that with a slightly enhanced

program, the virus could have erased files on infected computers or remained undetected for weeks, surreptitiously changing information on computer files.

Vulnerabilities Highlighted by Virus

In the aftermath of the virus, questions have been raised about how the virus spread, how it was contained, and what steps, if any, are needed to increase Internet security. These questions have been the subject of a number of post-virus meetings and reports prepared by government agencies and university researchers.[1]

On the basis of these assessments, we believe that the virus incident revealed several vulnerabilities that made it easier for the virus to spread and more difficult for the virus to be eradicated. These vulnerabilities also came into play in later intrusions (not involving a virus) onto several Internet sites in November and December. The vulnerabilities—lack of a focal point for addressing Internet-wide security problems; security weaknesses at some host sites; and problems in developing, distributing, and installing systems software fixes—are discussed below.

Lack of a Focal Point to Address Internet-Wide Security Problems

During the virus attack, the lack of an Internet security focal point made it difficult to coordinate emergency response activities, communicate information about the virus to vulnerable sites, and distribute fixes to eradicate it.

A Defense Communications Agency account of the virus cited a series of problems stemming from the lack of a central, coordinating mechanism. For example:

- Although the virus was detected at various sites, users did not know to whom or how to report the virus, thus hindering virus containment and repair.
- There were no plans or procedures for such an emergency. People used ad hoc methods to communicate, including telephone or facsimile. In many instances, sites disconnected from the Internet. While effective in the short run, this action also impeded communications about fixes.

[1]Major meetings included (1) a November 8 NCSC-hosted meeting to review the virus attack and its aftermath, attended by over 75 researchers and administrators from government and academia, and (2) a December 2 meeting of UNIX vendors and users, hosted by NCSC, NIST, and a users group.

- It was unclear who was responsible for protecting networks from viruses, resulting in confusion among user, network, and vendor groups.

The confusion surrounding the virus incident was echoed by many Internet users. For example:

- A Purdue University researcher concluded that user response to the virus was ad hoc and resulted in duplicated effort and failure to promptly disseminate information to sites that needed it.[2]
- At Energy's Los Alamos National Laboratory, researchers reported that they received conflicting information on fixes. Because they did not have a UNIX expert on site, they had difficulty determining which fix was reliable.
- At Harvard University, researchers expressed frustration at the lack of coordination with other sites experiencing the same problems.

In a report resulting from NCSC's post-mortem meeting, network sponsors, managers, and users from major sites—including Defense's Army Ballistic Research Laboratory, Energy's Lawrence Livermore National Laboratory, DARPA, Harvard, MIT, and the University of California, Berkeley—called for improved communications capabilities and a centralized coordination center to report problems to and provide solutions for Internet users.

Host Security Weaknesses Facilitated Spread of Virus

Key to the Internet's decentralized structure is that each host site is responsible for establishing security measures adequate to meet its needs. Host computers are frequently administered by systems managers, typically site personnel engaged in their own research, who often serve as systems managers on a part-time basis.

According to virus incident reports as well as network users, weaknesses at host sites included (1) inadequate attention to security, such as poor password management, and (2) systems managers who are technically weak.

Inadequate Attention to Security

Discussions of computer security frequently cite the trade-offs between increased security and the sacrifices, in terms of convenience, system function, flexibility, and performance, often associated with security

[2]Eugene H. Spafford, The Internet Worm Program: An Analysis, Department of Computer Sciences, Purdue University, Nov. 1988.

measures. In deciding whether to establish additional security measures, systems managers must often be willing to make sacrifices in these areas. According to Internet users from academia, government, and the private sector, systems managers at research sites often are not very concerned with security.

One example of a trade-off between security and convenience involves trusted host features on UNIX that allow users to maintain a file of trusted computers that are granted access to the user's computer without a password. The trusted host features make access to other computers easier; however, they also create potential security vulnerabilities because they expand the number of ways to access computers.

The virus took advantage of the trusted host features to propagate among accounts on trusted machines. Some sites discourage use of the trusted host features; however, other sites use them because of their convenience. One Internet user observed that users do not like to be inconvenienced by typing in their password when accessing a trusted computer, nor do they want to remember different passwords for each computer with which they communicate.

Another example involving inadequate attention to security is in password management. According to an NSF official, a major vulnerability exploited by the virus was lax password security. The official stated that too few sites observe basic procedures that reduce the risk of successful password guessing, such as prohibiting passwords that appear in dictionaries or other simple word lists and periodically changing passwords.

The relative ease with which passwords can be guessed was discussed in an analysis of the Internet virus done by a University of Utah researcher.[3] He cited a previous study demonstrating that out of over 100 password files, up to 30 percent were guessed using just the account name and a couple of variations.

Careful control over passwords often inconveniences users to some degree. For example, an article in Computers and Security, an international journal for computer security professionals, notes that computer-

[3]Donn Seeley, A Tour of the Worm, Department of Computer Science, University of Utah, Nov. 1988. Unpublished report.

generated passwords tend to be more secure than user-selected passwords because computer-generated passwords are not chosen by an obvious method easily guessed by an intruder. However, computer-generated passwords are generally more difficult to remember.[4]

Systems Managers Who Are Technically Weak

A number of Internet users, as well as NCSC and Defense Communications Agency virus reports, stated that the technical abilities of systems managers vary widely, with many managers poorly equipped to deal with security issues, such as the Internet virus. For example, according to the NCSC report, many systems managers lacked the technical expertise to understand that a virus attacked their systems and had difficulty administering fixes. The report recommended that standards be established and a training program begun to upgrade systems manager expertise.

Problems in Developing, Distributing, and Installing Software Fixes

Systems software is generally very complex. A major problem programmers face in software design is the difficulty in anticipating all conditions that occur during program execution and understanding precisely the implications of even small changes. Thus, systems software often contains flaws that may create security problems, and software changes often introduce new problems.

Internet users and software vendors frequently cited problems relating to inadequacies in developing, distributing, and installing corrections to identified software holes. Holes that are not expeditiously repaired may create security vulnerabilities. The Internet virus incident and two later Internet intrusions highlighted problems in getting vendors to develop and distribute fixes and in having host sites install the fixes.

Problems With Vendors

A number of network users representing major Internet sites said that vendors should be more responsive in supplying patches to identified software holes. For example, more than 1 month after the virus, several vendors reportedly had not supplied patches to fix the sendmail and fingerd holes.

Most vendors, when notified of a hole, send users a patch to repair the hole or wait until their next software revision, at which time the hole (as

[4]Belden Menkus, "Understanding the Use of Passwords," Computers and Security, Vol. 7, No. 2, April 1988.

well as any other identified flaws) will be corrected. However, since a revision may take up to 6 to 9 months to release, the latter approach may leave systems vulnerable to security compromise for long periods. According to Internet users, critical security patches should be provided as quickly as possible and should not be delayed until the next release of the software.[5]

Officials of one major vendor pointed out the problems they faced in distributing patches expeditiously. According to these officials:

- Their company sells computers with three or four different architectures, each with several versions of the UNIX operating system. When a fix is needed, they have to distribute about 12 different patches, making it difficult to develop and release patches quickly.
- Patches have to be carefully screened so that new holes will not be inadvertently incorporated. The officials noted that the quality assurance this screening provides is an important part of their business because their reputation depends on the quality of their software.
- Vendors have a hard time keeping track of customers who do not have service maintenance contracts. In addition, some systems are sold through contractors and the vendors may not know the contractors' customer bases.
- Disseminating a patch to thousands of users can cost a company millions of dollars.

The vendor officials said they considered these factors in determining how to implement a patch.

Berkeley's Computer Systems Research Group, which distributes its version of UNIX, has a software policy that differs from that of many other vendors. Berkeley generally provides source code along with the UNIX object code it sells to users.[6] However, Berkeley's policy is unusual—most vendors treat source code as proprietary and it is typically not provided to users. With source code, an experienced systems manager may be able to fix holes without waiting for the vendor to supply a patch or a system revision.

[5]According to a Defense official, this problem is compounded by the fact that sites not subscribing to software maintenance/support may not receive any new releases.

[6]Source code is the program written by the programmer. It is translated (by a compiler, interpreter, or assembler program) into object code for execution by the computer.

Berkeley routinely transmits fixes to UNIX users and vendors through networks and bulletin boards. While this may result in timely fixes, it can also create security vulnerabilities. In particular, when a fix is widely disseminated, information about a vulnerability is also made apparent. Thus, there is a race between intruders seeking to exploit a hole and systems managers working to apply the fix.

This dilemma was highlighted in multiple intrusions, which occurred in November and December 1988, at several Internet sites, including Lawrence Livermore National Laboratory and Mitre Corporation. In these instances, intruders exploited vulnerabilities in a UNIX utility program, called FTPD, that transfers files between Internet sites.[7]

Berkeley had sent out patches for the FTPD hole in October 1988. However, other UNIX vendors had not released patches for the hole. Mitre officials reported that their systems managers applied the Berkeley patch on many of their computers, but not on the computer penetrated by the intruders. Lawrence Livermore officials reported that they applied patches to computers that use Berkeley UNIX. However, the vendor for its other computers had not supplied a patch before the intrusion. Lawrence Livermore did not have source code for the other vendor's machines, so they had to wait for the vendor's patch.

According to a Defense official, the intruders most likely tried to gain access to many machines until they found those machines to which patches had not been applied. Once the intruders penetrated the FTPD hole, they installed "trap doors" by adding new accounts and modifying systems routines, which allowed them continued access after the FTPD holes were closed. Officials from the Federal Bureau of Investigation and from sites involved in the intrusions said that the intruders have been identified and the case is under investigation. Reportedly, aside from the trap doors, no files were altered, and no classified systems were affected.

Problems in Installing Software Fixes

Even when a vendor distributes fixes, there is no assurance that sites will install them. Internet users and managers at several major university research and government sites cited the following reasons as to why fixes were not expeditiously installed:

[7]As discussed, the Internet virus exploited vulnerabilities in two other UNIX utility programs, sendmail and fingerd.

- Systems managers vary in their ability and motivation to manage their systems well.
- System managers often serve on a part-time basis, and time spent on systems management takes away time from research.
- System revisions may contain errors, so some systems managers are reluctant to install the revisions.
- System revisions may be expensive if the system is not on a maintenance contract.
- Some sites do not know who their system managers are and, thus, have problems ensuring that fixes get distributed and installed.

As discussed earlier, problems and confusion resulted when sites had to respond to the Internet virus. Although Berkeley posted a fix to both the sendmail and fingerd holes within 2 days after the onset of the virus and Sun Microsystems reportedly published a fix within 5 days, almost a month after the virus a number of sites reportedly still had not reconnected their host computers to the Internet.

Actions Taken in Response to Virus

In response to the Internet virus, DARPA, NIST, NCSC,[8] and a number of other agencies and organizations have taken actions to enhance Internet security. These actions include developing computer security response centers, coordinating meetings, preparing publications to provide additional guidance, and publishing statements of ethics.[9]

Computer Security Response Centers Established

In the wake of the virus, many Internet users, site managers, and agency officials have voiced concerns about problems in responding to and preventing emergencies, such as the Internet virus. To address these concerns, some agencies are developing computer security response centers to establish emergency and preventative measures.

The first center, the Computer Emergency Response Team (CERT), was established by DARPA in mid-November 1988. CERT's mandate is broad—it is intended to support all of the Internet's research users. DARPA views CERT as a prototype effort for similar organizations in other computer

[8]NIST is responsible for developing standards and guidelines for the security of unclassified federal computer systems. It performs these responsibilities with the National Security Agency's technical advice and assistance. The National Security Agency (of which NCSC is a part) is responsible for the security of classified information in the defense and national security areas, including that stored and processed on computers.

[9]In addition, agencies are engaged in ongoing research aimed at improving network and computer security. An overview of these activities is presented in appendix II.

communities. Also, CERT is seen as an evolving organization whose role, activities, and procedures will be defined as it gains experience responding to Internet security problems.

According to DARPA, CERT's three main functions are to provide

- mechanisms for coordinating community response in emergencies, such as virus attacks or rumors of attacks;
- a coordination point for dealing with information about vulnerabilities and fixes; and
- a focal point for discussion of proactive security measures, coordination, and security awareness among Internet users.

CERT has no authority, although it can make recommendations. CERT officials recognize the need to establish credibility and support within the Internet community so that its recommendations will be acted upon.

CERT's nucleus is a five-person coordination center located at the Software Engineering Institute at Carnegie Mellon University in Pennsylvania.[10] CERT has enlisted the help of over 100 computer specialists who are on call when problems arise in their areas of expertise. In addition, CERT is developing working relationships with government organizations, including NCSC, NIST, Energy, and the Federal Bureau of Investigation, and with vendor and user groups. CERT expects to rely on DARPA funding until its value is recognized by the Internet community and alternate funding mechanisms are established—probably within 3 to 5 years.

The Department of Energy began setting up a center at Lawrence Livermore National Laboratory in February 1989. This center is to focus on proactive preventive security and on providing rapid response to computer emergencies within the agency. The center plans to develop a data base of computer security problems and fixes, provide training, and coordinate the development of fixes. In addition, the center is considering developing software to assist in network mapping and to assure proper system configuration.

[10]The objective of the institute, which is a Federally Funded Research and Development Center, is to accelerate the movement of software technology into defense systems.

Meetings Held and Guidance Issued

NIST is coordinating interagency meetings to (1) draw on agency experience and develop a model for agencies to use in setting up response/coordination centers and (2) educate others on the model that is developed. NIST has also set up a computer system that may be used as a data base for computer problems and fixes and as an alternate means of communication in case the Internet's electronic mail system becomes incapacitated. In addition, NIST is planning to issue guidance this summer that will discuss threats inherent to computers and how such threats can be reduced.

NCSC plans to distribute three security-related reports discussing (1) viruses and software techniques for detecting them, (2) the role of trusted technology in combating virus-related programs, and (3) security measures for systems managers. NCSC is also providing an unclassified system to serve as an alternate means of communications in case the Internet's electronic mail system is not working.

Ethics Statements Released

The Internet Activities Board, a technical group comprising government, industry, and university communications and network experts, issued a statement of ethics for Internet users in February 1989. Many Internet users believe there is a need to strengthen the ethical awareness of computer users. They believe that a sense of heightened moral responsibility is an important adjunct to any technical and management actions taken to improve Internet security.

The Board endorsed the view of an NSF panel that characterized any activity as unethical and unacceptable that purposely

- seeks to gain unauthorized access to Internet resources;
- disrupts the intended use of the Internet; or
- wastes resources, destroys the integrity of computer-based information, or compromises users' privacy.

The Computer Professionals for Social Responsibility and various network groups have also issued ethics statements encouraging (1) enforcement of strong ethical practices, (2) the teaching of ethics to computer science students, and (3) individual accountability.

Conclusions

In the 20 years in which it evolved from a prototype DARPA network, the Internet has come to play an integral role in the research and development community. Through the Internet, researchers have been able to

collaborate with colleagues, have access to advanced computing capabilities, and communicate in new ways. In providing these services, the Internet has gone beyond DARPA's original goal of proving the feasibility of computer networking and has served as a model for subsequent public data networks.

Since there is no lead agency or organization responsible for Internet-wide policy-making, direction, and oversight, management on the Internet has been decentralized. We believe this is because, at least in part, Internet developments were driven more by technological considerations than by management concerns and because decentralized authority provided the flexibility needed to accommodate growth and change on an evolving network. However, we believe that the Internet has developed to the point where a central focus is necessary to help address Internet security concerns. These concerns will take on an even greater importance as the Internet evolves into the National Research Network, which will be faster, more accessible, and have more international connections than the Internet.

The Internet virus and other intrusions highlighted certain vulnerabilities, including

- lack of a focal point in addressing Internet-wide security issues, contributing to problems in coordination and communications during security emergencies;
- security weaknesses at some host sites; and
- problems in developing, distributing, and installing systems software fixes.

Since the virus, various steps have been taken to address concerns stemming from the incident, from creating computer security response centers to issuing ethics statements to raise the moral awareness of Internet users.

We support these actions and believe they are an important part of the overall effort required to upgrade Internet security. Host sites may need to take additional actions to heighten security awareness among users and to improve identified host level weaknesses, such as lax password management.

However, many of the vulnerabilities highlighted by the virus require actions beyond those of individual agencies or host sites. For this reason, we believe that a security focal point should be established to fill a

void in the Internet's management structure and provide the focused oversight, policy-making, and coordination necessary at this point in the Internet's development.

For example, we believe that concerns regarding the need for a policy on fixes for software holes would be better addressed by a security focal point representing the interests of half a million Internet users than by the ad hoc actions of host sites or networks. Similarly, a security focal point would better ensure that the emergency response teams being developed by different Internet entities are coordinated and that duplication is lessened.

There are no currently available technical security fixes that will resolve all of the Internet's security vulnerabilities while maintaining the functionality and accessibility that researchers believe are essential to scientific progress. Similarly, there is no one management action that will address all of the Internet's security problems. However, we believe concerted action on many fronts can enhance Internet security and provide a basis for security planning on the National Research Network.

FRICC, an informal group made up of representatives of the five agencies that operate Internet research networks, is attempting to coordinate network research and development, facilitate resource sharing, and reduce operating costs. However, no one agency or organization has responsibility for Internet-wide management and security. The Office of Science and Technology Policy, through its Federal Coordinating Council on Science, Engineering and Technology, has, under its mandate to develop and coordinate federal science policy, taken a leadership role in coordinating development of an interagency implementation plan for the National Research Network. Therefore, we believe that the Office, through FCCSET, would be the appropriate body to coordinate the establishment of a security focal point.

Recommendation

We recommend that the President's Science Advisor, Office of Science and Technology Policy, through FCCSET, coordinate the establishment of an interagency group to serve as an Internet security focal point. This group should include representatives from the federal agencies that fund Internet research networks.

As part of its agenda, we recommend that this group:

- Provide Internet-wide policy, direction, and coordination in security-related areas to help ensure that the vulnerabilities highlighted by the recent incidents are effectively addressed.
- Support efforts already underway to enhance Internet security and, where necessary, assist these efforts to ensure their success.
- Develop mechanisms for obtaining the involvement of Internet users; systems software vendors; industry and technical groups, such as the Internet Advisory Board; and NIST and the National Security Agency, the government agencies with responsibilities for federal computer security.
- Become an integral part of the structure that emerges to manage the National Research Network.

Factors Hindering Prosecution of Computer Virus Cases

The Internet incident is a recent example of the growing number of instances in which computers, or their information or programs, have been the target of sabotage or attack. As of March 23, 1989, there have been no indictments in the Internet virus case. Because it is an open matter, Justice officials would not provide any specific information about the case.

There are some factors that may hinder prosecution of computer virus-type incidents. For example:

- There is no federal statute that specifically makes such conduct a crime, so other federal laws must be applied to computer virus-type cases.
- The technical nature of computer virus-type cases may hinder prosecution.

As yet, there have been no federal prosecutions of computer virus-type incidents.

No Statute Specifically Directed at Viruses

No federal law is specifically directed at computer virus-type incidents. Thus, the ability to prosecute such cases depends on whether conduct associated with a particular incident, such as unauthorized access or destruction of records, falls within an existing statute.

The Computer Fraud and Abuse Act of 1986 (18 U.S.C. 1030) is the act most closely directed at computer crimes. The most relevant provisions in the act relating to virus-type incidents make it a crime for individuals to

- intentionally,[1] without authorization, access a federal computer or a federally used computer if such access affects the government's operation of the computer;
- knowingly,[2] and with intent to defraud, access a federal interest computer[3] or exceed authorized access, where such access furthers the

[1]The term "intentionally" means that the outcome was an objective of the conduct.

[2]The term "knowingly" means that the actor was aware that the result was practically certain to follow from the conduct.

[3]The act defines federal interest computers as ones exclusively used by the government or a financial institution, or if not exclusively so used, used by the government or a financial institution and the conduct constituting the offense affects the financial institution's or the government's operation of the computer, or a computer that is one of two or more used in committing the offense, not all of which are in the same state (18 U.S.C. 1030(e)(2)).

intent to defraud and obtains anything of value, unless the object of the fraud and the thing of value consists only of the use of the computer; or

- intentionally, without authorization, access and by such conduct alter, damage, or destroy information in any federal interest computer or prevent the authorized use of such computer or information and thereby (A) cause losses aggregating $1,000 or more to one or more others during any one year or (B) modify or impair, or potentially modify or impair, the medical examination, diagnosis, treatment, or care of one or more individuals.

The act defines some relevant terms, but not others. For instance, the act defines "exceeds authorized access" as access to a computer with authorization and use of such access to obtain or alter information in the computer that the accessor is not entitled to obtain or alter (18 U.S.C. 1030(e)(6)). However, the act does not define "access," "information," or "prevents the authorized use."

Because some of the terminology has not been defined, it is not clear whether all virus-type cases would fit within the act's scope. For instance, it is unclear whether the introduction of a virus into a system by electronic mail, a nominally authorized means of entry, would constitute unauthorized access as contemplated by the statute. Nor is it clear that a virus that merely slowed a system's response time would prevent its authorized use.

There are also obstacles in applying other federal laws to virus-type incidents. For example, it is possible to view the creation and use of counterfeit passwords (used, for example, in the Internet incident) as a violation of the Credit Card Fraud Act of 1984 (18 U.S.C. 1029). This statute prohibits the production or use of counterfeit or unauthorized access devices with the intent to defraud. However, the act's legislative history[4] suggests that it is intended to address financial and credit abuses, and it is not certain that its prohibitions could be extended to nonfinancial incidents.

Another law that has been suggested for use in prosecuting virus-type incidents is the Wire Fraud Act (18 U.S.C. 1343). This act prohibits the introduction into interstate or foreign commerce of radio, wire, or television communications intended to further a fraudulent scheme. However, applying this statute to virus-type incidents may be complicated by the

[4]See House Report 894, 98th Congress, 2d Session; Senate Report 368, 98th Congress, 2d Session.

absence of traditional fraud elements, such as the effort to obtain something of value.

In addition to federal laws, computer crimes may be prosecuted under state laws. Forty-eight states have adopted legislation dealing with computer crimes, and the other two are currently considering such legislation.[5] State laws vary widely in terms of coverage and penalties. For instance, some state laws:

- Include provisions that specifically define information stored in computers as property. This definition facilitates prosecution under traditional statutes governing property crimes.
- Authorize victims to sue for violations of the statutes.
- Provide for forfeiting (that is, permanently taking away) the violator's computer property used in the crime as part of the penalty. Federal statutes do not provide for such a remedy or penalty.

Technical Nature of Virus-Type Incidents May Hinder Prosecution

The technical nature of computer virus-type incidents may hinder prosecution. Even when a violation can be clearly established, the evidence is likely to be arcane and technical, and prosecutors may not have the background and training needed to deal with it proficiently. Moreover, even if prosecutors are prepared to deal with the evidence, it is not likely that the court and jury would be similarly capable of assessing complex computer-related evidence. Consequently, prosecutors would need to devote additional resources and effort in preparing to communicate the substance of the case. This difficulty was described by the court in a 1985 software copyright case involving similar types of evidence:

"This fact-rich case has presented difficult issues for resolution, particularly since the intellectual property at issue is computer programming, a form not readily comprehended by the uninitiated. The challenge to counsel to make comprehensible for the court the esoterica of bytes and modules is daunting."[6]

Another potential problem in prosecuting virus-type incidents is that pretrial discovery may be burdensome and raise problems regarding access to sensitive computer records or security systems.[7] For example,

[5]Statistics were not readily available regarding the extent to which state laws have been used for prosecuting computer virus-type cases.

[6]Q-CO Industries, Inc. v. Hoffman, 625 F.Supp. 608, 610 (1985).

[7]The term "discovery" refers to pretrial legal procedures that can be used by one party to obtain facts and information from the other party in order to assist in preparation for trial.

in a recent Texas case involving a virus-type incident,[8] the defense moved for access to the victim company's backup tapes containing confidential records. The issue was ultimately resolved by giving the defendant access to the data over one weekend, with physical control of the tapes remaining in the company's hands. However, it is possible that similar requests for access to computer files or even security systems could deter prosecution in future incidents.

Proposed Legislation on Computer Viruses and Related Offenses

Two bills have been introduced in the Congress dealing with computer viruses and related conduct. These bills contain language addressing computer-virus type incidents. In addition, they provide for a private right of action authorizing the injured party to sue for a violation. Neither of the bills includes a forfeiture penalty.

The proposed Computer Virus Eradication Act of 1989 (H.R. 55) adds a new provision to the Computer Fraud and Abuse Act of 1986 prohibiting the introduction of commands or information into a computer program knowing that they may cause loss, expense, or risk to the health or welfare of the computer's users or to persons who rely on information contained in the computer program. The bill also prohibits individuals from knowingly transferring a program containing such instructions in circumstances where the recipient is unaware of the program or its effects. The bill provides for criminal penalties and fines and authorizes victims to sue for a violation of the statute.

The second bill, the Computer Protection Act of 1989 (H.R. 287), prohibits the knowing and willful sabotage of the proper operation of a computer hardware system or associated software that results in loss of data, impaired computer operation, or tangible loss or harm to the computer's owner. This bill also provides for criminal penalties and fines and authorizes the victim to sue for a violation of the statute.

In addition to these bills, which have been referred to the Judiciary Committee, Department of Justice officials said they are considering draft legislation to better address virus-type incidents.

Conclusions

Federal laws are not specifically directed at virus-type incidents. The law most relevant to such incidents is untested with respect to virus-

[8]Texas v. Burleson, unreported. Our discussion is derived from an unpublished case summary prepared by the Office of the Criminal District Attorney, Tarrant County, Texas.

type offenses and contains terms that are not defined. To date, no federal computer virus-type cases have been tried. In addition, the technical nature of computer virus-type incidents may hinder the prosecution of such cases. Legislation directed at computer virus-type incidents could eliminate the uncertainty regarding the applicability of current laws.

History of Computer Viruses

Computer viruses and worms are generally described as programs that can infect, replicate, and spread among computer systems.[1] The effects of viruses and worms have ranged from an unexpected message flashed on a computer's screen to destruction of valuable data and program files. Although computer viruses are a relatively recent threat, there are many varieties or strains that may infect computer systems.

Vulnerabilities in PC Design and Use Have Been Exploited by Viruses

Historically, most viruses have attacked personal computers rather than other systems, such as minicomputers, workstations, and mainframes. A Defense official said that the principal reason for this is that the first generation of PCs, due to their hardware and systems software design, are intrinsically vulnerable. For example:

- Early generation PCs do not have the same hardware and software capabilities for managing system resources that workstations and larger scale systems do. PCs were originally intended to serve only one user, and limitations on user privileges were not incorporated into PCs' accessing schemes.
- Most PCs do not differentiate among users and, therefore, every person who operates a PC has access to all resources.
- With PCs, the programs that enable the computer to operate are unprotected; they are stored on the same hard disk as the operator's files and there are few limitations on accessing program files.

In addition, PCs are often used in offices, where access is not monitored or recorded. Diskettes are shared among computer users, and networking is becoming common practice in organizations that use PCs. These operating conditions enable virus-type programs to spread among computers with relative ease.

According to Defense agency officials, creating a PC virus requires only moderate programming skills and access to a PC. These and other basic security weaknesses often make PC virus prevention, detection, and eradication difficult.

How Viruses Spread

Viruses are often spread among PCs by sharing infected computer diskettes, down-loading infected programs from electronic bulletin boards,

[1] Viruses are closely related to computer worms—they both spread and reproduce and their effects can be identical. The primary distinction between the two is that a worm is self-replicating and self-propagating, while a virus requires human assistance (usually unwitting) to propagate. Virus propagation can occur by sharing diskettes, forwarding mail messages, or other means.

or using infected software packages. For example, viruses may spread when an infected diskette is loaded into a computer. The virus may copy itself from the infected diskette onto the PC's hard disk. When other diskettes are inserted into the infected machine, they also become infected. These newly infected diskettes can then infect other computers that they come in contact with. This cycle continues until the virus is detected and eliminated. In the PC community, computers can be reinfected many times by the same virus and, even after viral attacks, may be left just as vulnerable as before. Therefore, virus attacks in the PC community may last for months or years. Recently, networks have also been used to transmit viruses among personal computers.

Viruses and other similar programs can be designed to trigger a wide variety of actions. For example, they can destroy files and hinder or stop computer operations. Viruses may also be designed to remain dormant until certain conditions occur. When the designated condition is met, the virus activates to achieve its intended purpose. For example, some viruses have been reported to trigger an action on a specified day, such as Friday the 13th, or after being recopied a certain number of times. Such threats can be difficult to address because they can create a false sense of security and hinder detection and recovery by infecting backup files. Viruses can also have less severe consequences. For example, they may create a message on the computer monitor, creating a nuisance and interrupting activities but not causing any damage.

Examples of Viruses

Viruses are tailored to attack specific systems and spread in different ways. Following are examples of well-known PC viruses:

- The 1986 "Pakistani Brain" virus was reportedly implanted in software packages as a warning or threat to those who recopy software. It infected IBM PCs and compatibles and copied itself onto diskettes that were inserted into infected systems. The virus contained the message "Welcome to the dungeon. Beware of this VIRUS. Contact us for vaccination." The message also included an address and phone number of the two brothers in Pakistan who originally distributed the software.
- The "Scores" virus of 1987 attacked Macintosh PCs. This virus infected utility programs and then transferred copies of itself onto program files located on diskettes inserted into the infected machines. The Scores virus caused system slowdown and printing problems.
- The "Lehigh" virus, discovered in 1987 at Lehigh University, attacked IBM PCs and compatibles. It infected PC operating systems and copied itself onto diskettes inserted into the machines. It was programmed to

infect four disks and then to destroy the computer's file system. It reportedly infected several hundred computers, many of which lost all the data on their disks.

The "Christmas Tree" virus of 1987 attacked IBM mainframes through an international network. It used electronic mail services to send copies of itself to network users. It displayed a holiday message on the receiver's screen and then mailed itself to others. The virus spread like an electronic chain letter through many kinds of communication links, including satellites and ocean cables, reportedly infecting computers in over 130 countries. This virus caused both denial of services and system shutdowns.

While there are many different kinds of computer viruses, there are also a number of commercial programs that can discover specific viruses through such methods as comparing storage requirements of an uninfected file with the actual storage space being occupied at any time by the file. Software packages used to discover specific viruses already present in computers include "Disk Watcher," "Protec," and "Condom."[2] However, according to Defense officials, because computer viruses are not recognizable based solely on their behavior or appearance, their detection cannot be completely assured. Currently, NCSC is evaluating such packages. In addition, officials said that because of the intrinsic vulnerabilities of most PCs, viruses can be written to circumvent most PC software security features.

The Internet Virus

The Internet incident, in which a virus-type program attacked computers through computer networks, demonstrates the potential extent and swiftness of propagation of self-replicating programs over networks. The Internet virus was the first to use several security weaknesses to propagate autonomously over a network. It was designed to attack Sun-3 and VAX computer systems that used system software based on Berkeley Software Distribution UNIX. It incorporated four primary attack methods to access thousands of computers connected by network communication lines. Two attack methods relied on implementation errors in network utility programs, a third method gained system access by guessing passwords, and the last method exploited local network security assumptions to propagate within the local networks. Because of the independent and flexible nature of its attack strategy,

[2]There are other software packages aimed at preventing initial viral infections.

the Internet virus was able to affect many systems within a short period.[3]

Infection Through Software Holes

The Internet depends on network utility programs, including remote login, file transfer, message handling, and user status reporting, to support communication between users. However, software security holes in two utility programs, sendmail and fingerd, enabled the Internet virus to propagate over the networks.[4]

Sendmail is a utility program that implements the Internet's electronic mail services by interacting with remote sites according to a standard mail protocol. The Internet virus used a weakness in sendmail involving a feature called "debug." This optional debug feature was designed into the original software as a convenience to programmers who tested network operations. According to Defense officials, the debug feature is not necessary for standard operations and should have been turned off in normal program distribution. However, through an apparent oversight, it was left activated on some releases. In those cases, the virus could exploit the debug command to send components of itself to remote hosts. It reproduced itself repeatedly as the computer received the virus components and constructed and executed the code.

Fingerd is a utility program that is intended to help remote users by providing public information about other network users. For example, fingerd can be used to determine which users are logged on to a specific computer. The program collects information from and delivers information to network users.

The virus exploited a security flaw in fingerd's procedure to collect information from remote network locations. In this instance, the virus sent more characters than fingerd had space to hold, thus overflowing the memory space allocated for storage of input parameters. Once outside this storage space, the virus overwrote the original program with portions of the virus code and was able to assume control of fingerd. Masquerading as fingerd and using fingerd's privileges, the

[3]PCs were not infected because they are not host computers on the Internet.

[4]The Internet virus exploited implementation errors in two utility programs that enable users to use network services. It did not attack or affect the computers' operating systems—the programs that control the computer's operation and access to resources.

virus could access, alter, or destroy any file that fingerd could. However, the virus was not destructive. It simply reproduced itself without damaging programs or data.

Passwords

The Internet virus also accessed systems by guessing user passwords. Many of the Internet's host computers store passwords (in encrypted form) and users' names in public files, a situation the virus exploited. The Internet virus encrypted potential passwords and compared them to the encrypted password stored in the computer's files. If they matched, the virus was able to gain access, posing as a legitimate user. It tried various passwords, including

- the user's first or last name,
- the last name spelled backwards, and
- the user's name appended to itself.

In addition, the virus contained a list of 432 potential passwords that it also encrypted and compared to the password file. Examples of such passwords include algebra, beethoven, tiger, unicorn, and wizard. The program also used words from the on-line dictionaries of the infected computers on the networks. Finally, access was attempted without using a password.

Trusted Host Features

Local area network managers can offer trusted host privileges to specific users on designated computers. These features are useful if a user wants to access his or her account frequently from another location. However, once the Internet virus infected computers on local area networks it was able to spread to other computers by exploiting these privileges. It used the feature to identify computers that had additional accounts accessible through known names and passwords. By using trusted host privileges, the virus was able to infect more Internet computers.

The virus also used trusted host privileges to identify which machines on the local networks could be accessed from other machines. The program was thus able to access many computers connected by the local networks. A Defense official compared the access policy on many of the Internet's local networks to security in an office building. For instance, in some buildings, visitors must pass through a security check at the entrance. Once inside, not every door in the building is locked because it is presumed that occupants have already passed the initial security test

when they entered the building. The Internet virus took advantage of the local area network's assumption that it was a legitimate process and spread to other machines within the local network.

Internet Virus Recovery

The Internet virus was eradicated from most host computers within 48 hours after it appeared, primarily through the efforts of computer experts at university research institutions. Patches were disseminated to sites to close the sendmail hole and fingerd holes. Once these holes were closed, the Internet virus could not reinfect the same computers providing the virus was not still present in trusted host computers.[5]

[5]According to a Defense official, many sites temporarily discontinued use of trusted host features until they were assured that the virus had been eradicated.

Research Aimed at Improving Computer and Open Network Security

Although DARPA, NIST, and NCSC sponsor or conduct considerable computer security-related research, none of these agencies are doing research specifically aimed at computer viruses.[1] According to NCSC officials, NCSC analysis of virus-type programs has been comparatively limited, with knowledge about such programs largely confined to simple examples drawn primarily from experiences with PC attacks and only recently extended toward large host and network examples. These agencies are, however, engaged in research that is aimed at enhancing computer and network security and that is, to varying degrees, applicable to open network environments, such as the Internet.

Computer Security Concerns Include Restricting Data Access and Ensuring Data Integrity

Computer and computer network security includes

- restricting data access to prevent disclosure of classified or sensitive information to unauthorized users and
- ensuring data integrity to protect data from unauthorized or accidental change or destruction.

A number of Internet users said that the government—particularly the Defense Department—has traditionally been more concerned about restricting data access than ensuring data integrity. For example, NCSC developed the "orange" and "red" books to describe computer systems that provide different degrees of access control.[2]

Current systems that meet stringent security requirements do so through physical isolation and providing access only to authorized individuals. To meet such requirements, sacrifices must be made in system function, performance, and cost, which are often unacceptable in an open network environment.

[1]NCSC is, however, evaluating commercial antiviral PC software packages. According to an NCSC official, the evaluation results will be distributed internally in spring 1989.

[2]NCSC's Trusted Computer System Evaluation Criteria, commonly referred to as the "orange book," describes criteria for evaluating computer security. These criteria describe the technical characteristics of a secure stand-alone computer system. The Trusted Network Evaluation Criteria, referred to as the "red book," describes criteria for evaluating network security.

Overview of Some Research and Projects That May Improve Security

The challenge in security research is to develop ways to increase security while minimizing the dollar, convenience, and performance costs associated with such security measures. Internet users, network sponsors, and vendors cited the following examples of research and methods that may improve computer and network security. These include (1) cryptographic methods and technology to permit users to send messages that can be understood (decrypted) only by the intended recipient, (2) improving controls on routing messages over the Internet, and (3) improving operating system quality to decrease program flaws and other security vulnerabilities.

Cryptographic Methods

Cryptography—the science of coding information to restrict its use to authorized users—can help ensure data integrity and confidentiality. NIST has designated one cryptographic approach, the Data Encryption Standard, as a Federal Information Processing Standard. This method involves a symmetric algorithm, which means the same "key" is used to both code and decipher data.[3] Research and development have produced advances in using cryptographic methods in such areas as public-key encryption, Kerberos authentication system, and portable access devices.

Public-Key Encryption

Unlike symmetric key systems, public-key encryption systems use two different keys for encrypting and decrypting data. Each user has a secret key and a public one. A sender uses the recipient's public key to send a message, and the recipient uses a private key to decode it. Since only the recipient holds the secret key, the message can be communicated confidentially. If the message is intercepted, or routed incorrectly, it cannot be decrypted and read. In addition, the message can carry additional information that assures the recipient of the sender's identity.

One method of implementing a public-key encryption system is based on a mathematical algorithm, developed by R. Rivest, A. Shamir, and L. Adleman at MIT, called the RSA algorithm. This algorithm is based on the mathematical difficulty of deriving prime factors.[4] Given an integer of more than 100 digits in length, it is very difficult to calculate its prime factors.

[3]An algorithm is the set of rules that describes the encryption process.

[4]A prime number can be divided only by itself and the number 1, without leaving a remainder.

Recently, the Internet Activities Board proposed standards based on a combination of the RSA algorithm and NIST's Data Encryption Standard. The proposed standards describe a hybrid cryptographic system intended to enhance the privacy of electronic messages exchanged on the Internet and to authenticate the sender's identity. The hybrid system uses symmetric cryptography to encrypt the message and public-key cryptography to transmit the key.

Each Internet user who uses the RSA algorithm will also receive an electronic certificate, electronically signed by a trusted authority. A computer security expert compared the certificate to a driver's license issued by the Department of Motor Vehicles. In the latter case, the Motor Vehicles Department is the trusted authority providing assurance to whomever checks the license. An Internet Activities Board official stated that this service should be available in late 1989.

Kerberos Authentication System

"Kerberos"[5] is a cryptographic-based challenged response system used at MIT to authenticate users and host computers. According to an MIT researcher, the system is intended to allow any two machines on a network to conduct secure and trusted communications, even when the network is known to be penetrated by intruders and neither machine has any intrinsic reason to trust the other. This system maintains passwords in a single secure host called a key-server. Because passwords are only present inside this key-server, the system is less vulnerable than if passwords were passed over the network. Individual machines make use of the key-server to authenticate users and host computers. Other groups, such as Berkeley's Computer Systems Research Group and Sun Microsystems, are also considering implementing this system to strengthen security.

Portable Access Control Devices

One small credit-card-sized device—called a "smart card"—uses cryptographic technology to control access to computers and computer networks. A smart card contains one or more integrated circuit chips, constituting a microprocessor, memory, and input/output interface. The card manages, stores, receives, and transmits information.

Each smart card has its own personal identifier known only to the user and its own stored and encrypted password. When the user inserts the

[5]Also Cerberos—in Greek mythology, the name of the three-headed dog who guarded the entrance to the underworld.

smart card into the reader/writer device, the terminal displays a message that identifies the smart card's owner. The user then enters the personal identifier. Once the identifier is authenticated, the host computer allows the user access. The smart card contains information that identifies what level of access the user is allowed. The smart card also maintains its own user audit trail.

According to a NIST official, smart cards are not currently in widespread use. This official stated, however, that a major credit card company is currently testing smart cards. In addition, the Belgian banking industry is testing smart card technology for use in electronic funds transfers, and NIST is testing smart card technology for the U.S. Department of the Treasury. Potential applications of smart card technology for the Treasury Department include authenticating disbursement requests from other federal agencies.

According to researchers, other portable access control devices are currently available. For example, one device—also a small-sized card—periodically displays changing encrypted values based on the time of day. A user enters the value displayed by the card to gain access to the host computer. Each card contains a unique encryption key. Because the host computer knows the time of day and can decipher the value displayed on the card, the host computer can authenticate a user.

Another small authentication device is available that contains a display screen and a small keyboard. When a user requests access to a host computer system, the host computer sends an encrypted challenge to the remote terminal. The user enters the challenge in the portable device and obtains an encrypted response to send to the host computer. If the user's response is correct, the host computer allows the user access. The advantage of these devices over smart cards is that no reader/writer device is required.

Improved Controls in Message Routing

Messages exchanged on the Internet travel through a series of networks connected by electronic switching units or "gateways." Messages are transmitted piecemeal in separate data groupings or "packets." Each packet contains address information, which a gateway reads to route the packet to its destination. Gateways also decide which paths to use. For example, a gateway can decide which path can route the data packet to its destination most quickly.

The message-switching technology incorporated on the Internet is very sophisticated. Although Internet uses advanced technology, Internet users have limited control over message routing. Data may travel through several different networks on the way to their ultimate destination. However, users cannot easily indicate their routing preferences to the Internet. For example, they cannot practically specify that their packets not be routed over a particular network, nor can a network sponsor practically specify that only packets of certain Internet users be allowed to traverse that network.

Research into a method called policy-based routing is currently underway that would allow Internet users the option of selecting their own communications paths by specifying certain parameters. Network sponsors could enforce their own individual network policies, perhaps by restricting their network resources to a certain class of users. Policy-based routing gives network users and owners some control over the particular routes data may take. For example, data packets that belong to the Defense Department could be routed using its network resources.

According to researchers, some of the technology needed for policy-based routing is not very complicated. Technology exists that can sort traffic into categories and route it through selected networks. However, labeling individual data packets with the necessary policy-based routing information is difficult. In particular, it is difficult to determine what information should be included on labels.

Improvements in Operating System Quality

Other researchers are attempting to improve operating system quality by decreasing program flaws and other security vulnerabilities. For example, DARPA is sponsoring formal methods projects for the development of high-quality assurance software systems. These techniques will be applied to operating systems. The formal methods techniques involve using mathematically precise specifications statements for critical program properties, such as safety and security. Using these specifications, it may be possible to ensure, by using a chain of mathematical proofs, that a program will operate as intended, and not in any other way. According to a DARPA official, unlike past approaches, current efforts focus on achieving assurance of quality during the design stage rather than attempting to apply techniques to already existing systems. The official noted that although the formal methods project is in the relatively early stages of research, the techniques are already being applied on a small scale in applications where very high levels of assurance are

required. The official said that there is significant progress in Europe in this area, particularly in the United Kingdom.

Major Contributors to This Report

Information Management and Technology Division, Washington, D.C.	Jack L. Brock, Jr., Director of Government Information and Financial Management, (202) 275-3195 Glen Trochelman, Assistant Director Jerilynn B. Hoy, Evaluator-in-Charge Mary T. Brewer, Evaluator Beverly A. Peterson, Evaluator Gwendolyn Dittmer, Evaluator
Office of the General Counsel, Washington, D.C.	John Carter, Attorney/Advisor
Boston Regional Office	Jeffrey Appel, Site Senior Debra Braskett, Evaluator
San Francisco Regional Office	Don Porteous, Evaluator