CHAPTER 4

Current and Future Technical Capabilities

Stephen J. Lukasik

Information systems are vulnerable to attack, as evidenced by the rapidly growing number of system intrusions. Current protection capabilities allow for relatively limited degrees of defense or deterrence, making attack relatively easy and defense relatively difficult. This chapter outlines technical measures that could help to identify and track intruders into computer systems and networks. But in addition to deploying specific technologies, a new international regime will be required to reduce the ability of intruders to hide behind the frequently slow and incompatible processes under which defenders must operate when different sovereign nations are involved in the investigation of an intrusion.

Some of what is suggested will require explicit agreements among

nations, but there are other steps that, without government-to-government agreements, will fit naturally into the open and cooperative environment that has characterized the Internet to date. Such measures, undertaken directly by the users who, by definition, live in "Internet-time," can be accomplished more rapidly than those where formal international agreements must be agreed to and implemented. Thus, in spite of the rapid rate of growth of system intrusions, there is the prospect of being able to make significant improvements in system security on a time scale that matches that of the problem.

The focus here is on criminal violators, including terrorists who seek to attack and destroy elements of society. A different, and potentially more severe, threat is where the attacker is a sovereign state. That class of attack, constituting what is called information warfare, is beyond the scope of the discussion here and of the Draft International Convention presented in Chapter 6.

1. The Internet and Its Governance

The Internet provides the basis for the global information infrastructure, and it increasingly provides connectivity for a wide range of other infrastructures. To the extent that the Internet manifests vulnerabilities that cannot be addressed through unilateral measures, changes in its governance may be desirable. This requires examining existing mechanisms that can effect change. To the extent that these change mechanisms prove to be incapable of protecting the Internet, other approaches will be needed to protect this global resource.

Internet Structure and Function Are Determined by its Developers

The Internet is governed through the voluntary activities of the technical people who develop and extend its functionality.[1] Internet users

1. See Michael A. Erlinger, "Internet Protocols for Protection Against Cyber Crime," presentation at the Conference on International Cooperation to Combat

adopt its protocols to be able to communicate with others; vendors implement its protocols in their hardware and software products because the user market requires it; technical experts extend the Internet based on their research and analyses. Thus, although the provision of Internet services has become a major business, the Internet's core specifications and their implementations are maintained and developed by its users acting through a governance structure that is little changed in concept from the way its founding academic researchers proceeded thirty years ago. Even the businesses that use the Internet or provide its communication facilities are not, as organizations, part of the group that develops and maintains the Internet's technical specifications, although those specifications must change to keep up with new technologies that are deployed. This may be changing, as more constituencies demand to be heard, but the fundamental processes through which the Internet operates have stood the test of time and are not likely to be abandoned lightly.

The Internet runs on the basis of network protocols, agreements on how information should appear in a message, how that information is to be interpreted, and the format of that message. The Internet is managed, developed, and operated by a hierarchy of volunteer organizations. The Internet Architecture Board (IAB), the Internet Engineering Task Force (IETF), the Internet Engineering Steering Group (IESG), and the various IETF Working Groups are the primary groups responsible for the development of the Internet.

The IETF, the engineering, development and standardization arm of the Internet Architecture Board (IAB), is a self-organized group dedicated to the continued development of the protocols and standards that form the basis of the Internet. The IETF's mission includes: identifying and proposing solutions to operational and technical problems in the Internet, specifying protocols and architecture to solve such technical problems, making recommendations to the Internet Engi-

---

Cyber Crime and Terrorism, Hoover Institution, Stanford University, Stanford, California, December 6–7, 1999.

neering Steering Group (IESG) regarding the standardization of protocols and protocol usage in the Internet, and providing a forum for the exchange of information within the Internet community.

The IETF is organized into eight areas of interest, one of which is security. Each area is headed by an Area Director, who oversees the various working groups—the creators of Internet specifications—that consist of all parties having a stake in the protocol under consideration. A charter sets the agenda and timetable for the group's work. Working groups focus on developing a protocol specification, but sometimes they serve to delineate some current practice. Working groups are established based on a recognized need and they terminate when the work called for by their charter is completed.

The IETF standardization process is well defined.[2] The process starts with a working group producing a series of documents called Internet Drafts that are posted on the WorldWideWeb[3] with notification to the IETF mailing list. Based on comments received, the working group can decide that an Internet Draft reflects a consensus view. At this point the working group asks the IESG, through its Area Director, to make the Internet Draft a Proposed Standard RFC.[4] A Proposed Standard specification is advanced only if it is stable, represents the resolution of known design choices, is well understood, has received significant community review, and appears to enjoy enough community interest to be considered valuable. Thus a Proposed Standard RFC is viewed as the initial specification from which implementations should be developed.

After a minimum of six months the working group can ask that the Proposed Standard RFC be moved to the level of Draft Standard RFC. The requirement for such a change is based on the existence of two independent and interoperable implementations from different code bases and on sufficient successful operational experience. Such

2. See S. Bradner, "The Internet Standards Process—Revision 3," RFC 2026 (October 1996), available at ⟨*http://www.ietf.org/rfc/rfc2026.txt*⟩.

3. See ⟨*http://www.ietf.org/ID.html*⟩.

4. See ⟨*http://www.ietf.org/rfc.html*⟩.

implementation experience may well uncover the need for changes to the specification. It is the implementation and experience requirement that differentiates the Internet standardization process from other standards activities. Implementations are used to prove the utility and precise operation of the specification.

After a minimum of four months a Draft Standard can be moved to an Internet Standard. Internet Standards are characterized by a high degree of technical maturity and by the understanding among its developers that the specified protocol or service provides significant benefit to the Internet community. This process is shown in Figure 1. Security issues were not of concern for the original Internet specifications. Cyber crime per se is not currently an IETF agenda item. Addressing cyber crime would have to be stated as an issue in protocol design.

The Internet Involves Both Physical and Information Architectures

There are two views of Internet architecture. The most common is that the Internet is a set of linked computer hardware, software, and communication facilities.[5] In this view, security is a matter of protecting that hardware and software against theft, damage, or denial of service. A second view is that the Internet links a collection of information. Though not as easily understandable as that of physical architecture, certain features of its information architecture are becoming clearer. First, just as the physical Internet embodies a meta-architecture that embraces the interconnection of heterogeneous networks through the processes described above, so does the information system embody a meta-architecture that embraces heterogeneous information systems. Robert Kahn and Robert Wilensky in their introduction to this general notion describe the digital object as the basic architectural element in

5. See Robert E. Kahn and Stephen J. Lukasik, "Fighting Cyber Crime and Terrorism: The Role of Technology," presentation at the Stanford Conference, December 6–7, 1999.

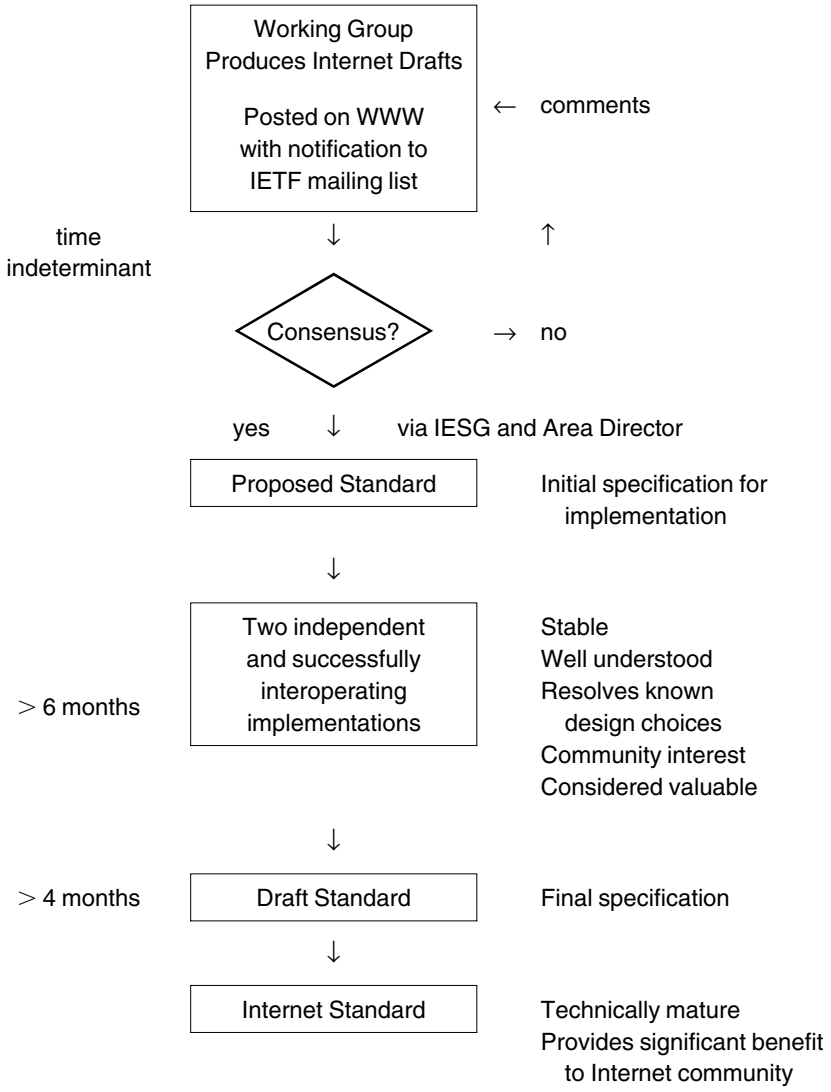

Fig. 1. The Internet standardization process.

the global information system, with each object having a unique and persistent identifier.[6]

Finally, the architecture assumes the existence of certain systems that can speak with authority on specific information issues. For example, is the following digital object subject to a claim of copyright in the United States? Or, regarding authenticity, what is the following organization? Or, are the bona fides of this organization accurate? While some entities may have their authority conferred by law, others will develop it by superior performance in the marketplace.

Given the existence of a meta-architecture for information systems, the Internet will embrace many of them, all different in some aspects. Today, the most widely used information systems are those that rely on access to files on specific machines. The older File Transfer Protocols (FTP) and the current Web protocols are examples of information system protocols, but because neither one allows for persistent access over time through unique identifiers, one can expect the spectrum of information systems in the future to range from those that are strictly informal, with possibly transitory information, to those that have more formality.

Defining the Internet

The broad term "information infrastructure" covers a range of private, public, and national capabilities. I focus here on the information systems that constitute the Internet and exclude separate national systems used by sovereign nations for the operation of their government, and facilities owned or operated by private entities that are separate from public facilities. I do, however, include both physical aspects and the information and service aspects that are not necessarily physical.

6. See Robert Kahn and Robert Wilensky, "A Framework for Distributed Digital Object Services," available at ⟨*http://www.cnri.reston.va/cstr/arch.html*⟩; see also "The Handle System," available at ⟨*http://www.cnri.reston.va/cstr.html#architecture*⟩.

Within the United States, the Federal Networking Council has adopted the following definition of the term "Internet":[7]

> Internet refers to the global information system that:
>
> (i) is logically linked together by a globally unique address space based on the Internet Protocol (IP) or its subsequent extensions/follow-ons;
>
> (ii) is able to support communications using the Transmission Control Protocol/Internet Protocol (TCP/IP) suite or its subsequent extensions/follow-ons, and/or other IP-compatible protocols; and
>
> (iii) provides, uses, or makes accessible, either publicly or privately, high level services layered on the communications and related infrastructure described herein.

An important element of this definition is the reference to "subsequent extensions/follow-ons." The dynamic character of networking technology suggests the need for a continuing review mechanism to ensure that basic conceptual underpinnings are not inadvertently changed with the passage of time. A second important element is the inclusion of high-level services within the definition. More than simply the physical devices and telecommunications capabilities that constitute the underlying network, one must view the global information system in its totality.

Abusing, Misusing, and Attacking the Internet

The Internet, after undergoing a twenty-year period of development and use by academic researchers, took on a different aspect when it was opened up to commercial users in the late 1980s. The environment quickly changed from one directed to the open and cooperative use of its information resources and services by a narrowly defined set of users to one where the new users sought privacy in their communications, and protection of their intellectual property, and expected to receive a choice of reliable services at costs set competitively.

7. See Federal Networking Council, October 24, 1995, available at ⟨*http://www.fnc.gov*⟩.

The new users display a varied set of behaviors. Some are business competitors or bargain-seeking consumers who are not above stealing intellectual property or services. Others are irresponsibly playful, and they set loose viruses and other malicious code or they deface web sites or destroy data files. Still others are criminals who use the Internet to perpetrate fraud, theft, and extortion. Terrorists, noting the increasing dependence of many societies on the Internet, may use it as a target or channel for the expression of their views, or to harass, coerce, or destroy social institutions. National security analysts translate these same concerns into the domain of strategic attacks on sovereign states. Such abuses and attacks may use the Internet to reach other targets, or they may target the Internet itself. The first category uses the Internet to attack defined entities, such as a particular computer, the information contained within it, or the service it provides, thereby causing harm to those dependent on those entities. The second category, attacking the Internet itself, intends to degrade or deny parts of its contents and capabilities to large numbers of users.

Compounding these difficulties are several related factors. As information technology becomes increasingly powerful, so do the attack tools that become available. These tools are distributed to would-be attackers using the Internet itself, much in the way that poisons spread through an organism in its bloodstream. By increasing the power of their users, the tools allow a larger number of individuals, less skilled than the tool creators, to damage information systems. This progression of increasingly sophisticated attacks is illustrated in Table 1.

Benign Causes of Disruption and Lack of Robustness Further Complicate Protection

Excluded from the concept of network abuse are failures that inadvertently result from acts of nature, wear and tear or other usage, approved maintenance and operations status monitoring, and diagnostic activities of network operators. These are considered part of the normal operational environment and part of the terms and con-

TABLE 1
The Increasing Sophistication of Computer and Network Attacks

Year	*Attack mechanism*
1982	Password guessing
1984	Self-replicating code
1985	Password cracking
1986	Exploiting known vulnerabilities
1988	Disabling audit mechanisms
1989	Use of back doors in programs
1990	Hijacking sessions
1991	Sweepers
1992	Packet sniffers
1993	Stealth diagnostics
1994	Packet spoofing
1995	Graphic user interfaces for attack tools
1996	Automated probes and scans
1997	Denial of service
1998	Web attacks
1999	Macro viruses
2000	Distributed attacks

Source: Thomas A. Longstaff, "International Coordination for Cyber Crime and Terrorism in the 21st Century," presentation at the Conference on International Cooperation to Combat Cyber Crime and Terrorism, Hoover Institution, Stanford University, Stanford, California, December 6–7, 1999.

ditions for use of the network that are made explicit in either employment contracts or service agreements. The focus of this discussion is on deliberate acts, not accidents, although the possibility of civil or criminal negligence cannot be excluded.

At the current level of networking technology numerous problems will arise from the inherent lack of robustness of the system and even of incompetence on the part of its users and operators. These unfortunate realities can have consequences as harmful and costly as those caused by deliberate conduct, but they are not addressed.[8]

8. See K. C. Claffy, "Traffic Observation in a Stateless Data Networking Environment," presentation at the Stanford Conference, December 6–7, 1999.

Protecting Individual Computers

As with any form of private property, protecting a single computer is primarily the responsibility of its owner. An owner may seek assistance from appropriate government agencies and jurisdictions, according to the national laws applying to the locations and the activities involved. The broader the impact of such attacks, the greater will be the case for public involvement. Where the information, operation, and impact of an attack involve entities from more than a single sovereign state, which is often the case, the appeal for help may extend to the international community of nations and organizations.

International actions can be as simple as supporting the resolution of civil contract disputes. In other cases they may involve the exchange of information relating to computer crimes, or suspected crimes, or to anomalies in computer operation. It is conceivable that personal information relating to the citizens of one nation might be provided to entities in other countries, such as alerts, warnings, modes of penetration, identities and aliases, and the like. In some cases assistance in the resolution of attacks on individual computers can involve disclosures of personal information that may raise concerns about the requirements of privacy.[9] A balance between property rights and those of personal privacy must be achieved. A possible approach to this problem, involving the use of automated techniques, is discussed later.

Protecting the Global Information Infrastructure

A very different circumstance arises when the information infrastructure itself, in whole or in part, is the target of attack. In this case, the jurisdiction in which such attacks occur is more complicated than for attacks on individual targets, and international agreements become increasingly important. Infrastructure "capabilities" are abstractions, substantially different from their physical manifestations; they depend

9. See Ekaterina A. Drozdova, "Civil Liberties and Security in Cyberspace," Chap. 5 of this volume, for a fuller discussion of privacy rights.

on multiple locations and jurisdictions (including the oceans and space) and, most important, have properties and capabilities that, rather than being simply resident in any of its parts, grow out of their collective existence. Without explicit agreements dealing with the protection of the network, global protective actions may be limited to those common to all jurisdictions taken together. This least common denominator will not, in general, provide the greatest amount, or even adequate amounts, of global protection.

For protection to be adequate, the shared global system should have the capacity to detect and identify violators with sufficient accuracy to deter them through the prospect of being indicted, prosecuted, and punished, or otherwise caused to pay a price for having initiated an attack. To achieve this degree of effectiveness will require creating appropriate tools, procedures, and organizations, not only to satisfy these aims but also to satisfy national authorities that their sovereign rights and the rights of their citizens are protected.

Although in some situations bilateral international interactions may be all that is required to deal with an attack on the Internet itself, in many others multinational action will be required. Certainly a global attack on the entire information system would require a coordinated worldwide response, and such a prospect is properly a subject for political consensus on how best to respond. At a minimum, the governing policies and processes will need to be established in advance.

2. Defending Information Systems Against Cyber Attack

For a realistic assessment of future prospects for the defense of information infrastructures against cyber attack, the current state of practice provides a starting point. In addition to the art and science of system defense, we should examine those areas that security professionals feel could provide the greatest leverage in blunting future attacks.

The View from the Defender: Prepared Defense

In the case of a prepared defense there is a secured network with a perimeter involving, for example, firewalls, virtual private networks, and/or challenge response systems.[10] This perimeter is monitored by an information protection staff using a combination of automated and semiautomated tools. An effective perimeter defense implies that the modes of possible attack are known a priori and that methods are in place to detect the signatures of these attacks. Early warning of an attack might, for example, come from monitoring the TCP, UDP, and ICMP activities occurring on in-bound packets.[11] Early warning could also come from automated searching for key words in the internals of network traffic that would generally indicate an attempt to gain access to systems or privileges not permitted from outside the secured perimeter.[12] This could in addition include the routine scanning of attached files in e-mail, which can contain known viruses or suspicious executable programs.

A potential attack may be indicated to the information protection staff via e-mail, pager, or other means. At this point, the information protection staff must ascertain whether the alert is real or a false alarm. This is done in several ways, and is somewhat dependent on what information is captured or buffered at the firewall. Essentially, the team must look manually at the information collected by the automated monitoring system in order to make an initial assessment. Adequate evidence may or may not be available, but in either case, in

10. See Steven D. Rizzi, "Is Technology the Answer to Infrastructure Protection?" presentation at the Stanford Conference, December 6–7, 1999.

11. For a discussion of TCP, UDP, ICMP, and other protocols, see Douglas Comer, 1 *Internetworking with TCP/IP* (Upper Saddle River, N.J.: Prentice-Hall, 1995).

12. "Internals" refers to the content of networked information. This could be the actual text of an e-mail, the graphics in a computer file, the audio or video of a teleconference, or the audio of a telephone call routed over a computer network. By contrast, "externals" refers to the routing information that is contained in the message that identifies the source and destination without providing detailed information of the message content.

order to support both evidence collection and better situation assessment, the security staff will increase the amount of auditing and data collection being done at the perimeter.

If the increased monitoring and auditing process confirms a suspected attack on the network, rapid engagement of additional processes is necessary. Engagement typically will require disaster preparation and warning, contingency plan execution, system isolation, and ultimately denial of access to the attacker. The security staff must move quickly to determine the nature of the intrusion, as well as to intervene to deny access to the attacker. During this period of the engagement, security personnel will trace the origin of the attack. In most cases, a serious attack will not have originated from the source that has been monitored; the real attacker will be located one or more "hops" away. That means that the defenders must attempt to contact security personnel at the apparent attack source to see if the attack is originating from that location or from some other "hop." This tracing period may take time, and prove difficult, since the only contact information that is usually available is from the personnel at the point of initial contact and the attack may not occur during business hours.

For these reasons, it is usually difficult for security staffs to trace a multiple "hop" attack, and the urge to deny the attacker access, either through the operating system or the firewall/router, is hard to resist. Having an automated way of tracing such attacks would be of great use, for it would increase the number of successful traces and would thus enhance deterrence.

Denial of access to attackers results in the immediate ending of any intrusion, but attackers frequently eliminate the evidence of their activities on some number of the intervening "hops," or they may place confusing or contradictory evidence in log files that would suggest to investigators that the attack came from somewhere other than the true origin. Sometimes the various "hops" along the way are outside the country boundaries of the computers under attack, creating the further difficulty of additional barriers such as different laws, languages, and time zones.

For most practical purposes, the average security staff that is well prepared for an attack of this nature will be in the position to detect, investigate, and terminate the intrusion, yet because of the "network of networks" nature of the Internet, detailed investigation leading to the location of an attacker is not likely to take place without the involvement and cooperation of law enforcement, Internet Service Providers (ISPs), and telecommunications providers. Unfortunately, this seldom occurs simply because the amount of work necessary to coordinate and carry out an investigation outweighs the value of finding the attacker—or at least it would require the group attacked to admit publicly that it had been attacked. These issues could be practically solved if we had an automated method for tracing attacks, particularly one that would protect the identity of the institution under attack.

Hasty Defense Is More Difficult

In many cases, the attack as launched has not been prepared for, and so the defenders must mount a hasty defense—not an easy task. Though the defenders know that something has happened, they may not be sure what, how, when, who, or why, but in all likelihood, speed is of the essence because evidence may be lost at any time. Usually the first step in such situations is to analyze what systems and/or networks may have been affected so that computer backups for those systems can be immediately secured and duplicated. A team may set to work analyzing the data contained on those backups, looking for clues to how the attack occurred. Additionally, the security staff may make duplicates of audit and transactions logs, as well as network traffic and/or firewall logs. At this point, the security staff will increase the level of auditing and monitoring conducted on the network, in order to collect improved evidence in the event of further attack. This may best be done with computers that are dedicated to "sniffing" packets off the network and storing the packets on removable media. Unfortunately, in the case of the hasty defense, it may or may not be clear

whether the attack occurred from outside or inside—that is, the attacker may possibly be an authorized user, operating inside the network.

In the process of analyzing the attack, defenders may discover some unique aspect that is characteristic to the attack—for example, certain accounts are used, a particular time of day, and so on. That may be difficult, however, since many computer networks are complicated interconnections of heterogeneous equipment with numerous ways in which they can be accessed. In addition, it is necessary to characterize the attack, and this characterization will form the basis by which investigators will review data, looking for possible intrusions or unauthorized accesses. The process of analyzing the network topology (how computers interact or interconnect with others on the network) may also be of assistance.

This correlation must lead to the development of a profile for the intruder that can be used to search through the large volumes of historical as well as current data, looking for further evidence of intrusion. Ultimately, the endgame in the hasty defense is locating the attacker through a more and more detailed approximation of the "attack signature." Continued refinement of this signature through the collection of a significant number of intrusions helps to identify the attacker, even when normal means to gain access to the computer systems in question are used—such as signing into the system using a stolen user name and password. Once a sufficiently discriminating signature has been developed, the hasty defense takes on the nature of the prepared defense, and defenders can continue to investigate and/or terminate the attacker's access.

Automated search tools are important to detect intrusions in the large volumes of data involved in normal network operation. Since this produces a number of sessions that contain possible events that must be reviewed individually by investigators, review at this level requires the detailed reading of interactions between users, computers, and other users. Most of these sessions, as one would expect, are those of legitimate users so technology could be of enormous value here, if

it could provide a means to detect and track unauthorized users without requiring investigators to read message internals. The response process, as seen by the defender, is shown in Figure 2.

Tracking Packets

One mechanism to identify tracks through the network could be facilitated by routers and the Internet service providers who operate them.[13] As indicated in Figure 2, if every time a packet left a router, the packet contents were fingerprinted, for example, by a checksumming or other mechanism implemented in hardware, one network could verify with its neighboring network that a packet with a specific fingerprint passed through the network at a given time.[14] This fingerprint could also be passed to the next network or destination system, which could then use the fingerprint to identify a specific packet, and each network on a hypothetical reverse path could then be queried to determine if it carried that particular packet. This, of course, could generate a very large amount of auxiliary data, which would, depending on how long it would be retained, have to be logged and stored, either by the routers or in repositories associated with them. (See Figure 3.)

A variant on this approach would keep statistical track of packet fingerprints over designated time frames, for example, over ten-minute periods, so that, with relatively much less storage, it could be determined with certainty that a given packet fingerprint did not pass through that network during that time frame. Although such a system would not be able to determine whether a given packet actually did pass through a network, it could be used to establish the negative.

If it were established that a given packet fingerprint was seen by

13. See Kahn and Lukasik, "Fighting Cyber Crime and Terrorism."

14. Such a fingerprint, also called a "Message Digest" by Rizzi ("Is Technology the Answer?") could be a cryptographic operation that takes an arbitrarily large input string and produces a relatively unique numeric representation. In other words, content information, that is, internals, are reduced to a number. The algorithm would be chosen so that it would be difficult to infer the input string from the fingerprint.

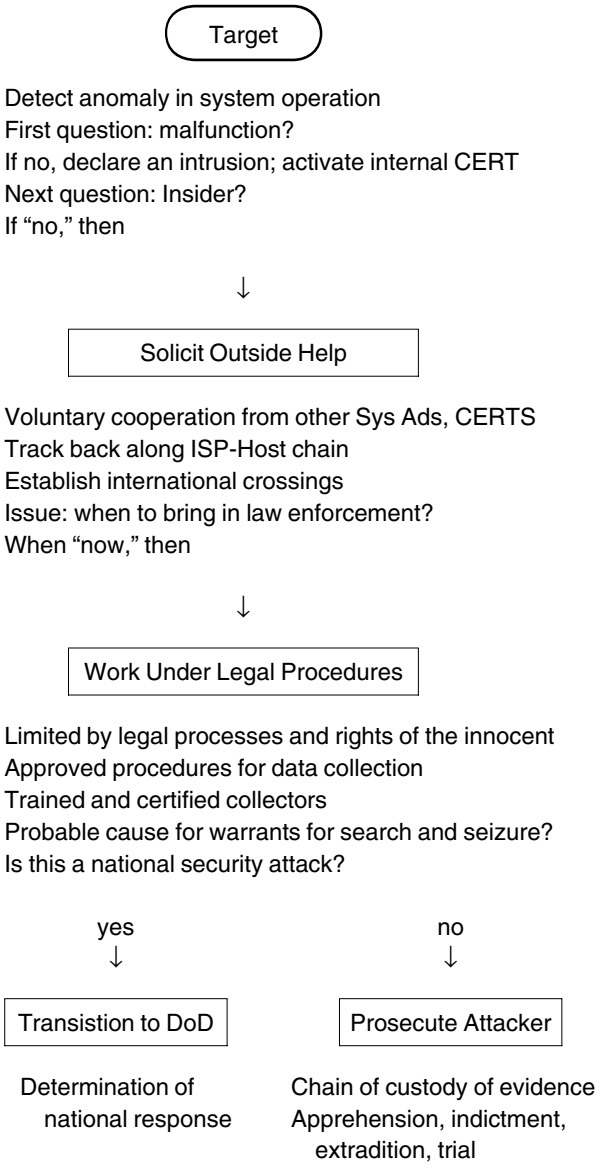

Fig. 2. The viewpoint from the defender’s position.

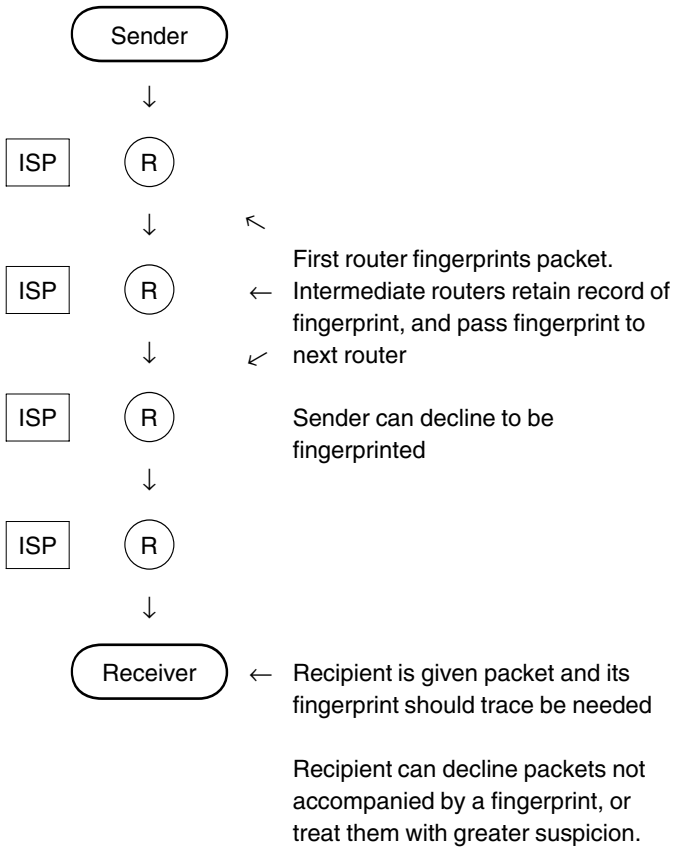

Fig. 3. Recovering the track of a packet.

N networks en route from host A to host B at a given time, that might be sufficient evidence to require disclosure of service-level information retained by host B that is correlated with the specific packet fingerprint, assuming it still existed. This information could then be used to correlate with the corresponding information in host A, which never had access to the packet fingerprint in the first place and therefore could not have initiated the inquiry. However, host A could initiate a reverse query based on received packets from host B, whose packet fingerprint host B was not privy to. This would let host A take the lead in identifying the reverse channel, and let host B take the lead in identifying the forward channel. If executed effectively, this would have an impact similar to the impact that accessing billing records has in tracing use of the telephone system and would be a useful tool for enabling cyber attackers to be identified and located.

A similar suggestion is made by Stephen Rizzi, who notes that technology for privacy-protecting packet tracing would also be desirable.[15] All too often an intrusion event is not investigated extensively because current manual methods do not allow for timely or privacy-protected tracing of multihop attacks. Technology is needed to support near-realtime automated tracing of multihop attacks. Such technology should protect the identity of the institution requesting the trace, to avoid undesired publicity, and in addition, the tracing algorithm should not directly use internals of the message to trace the origin. Such an architecture could be accomplished by the installation of a "trace server" on each registered domain subnet.

With time, such a server could be as important as a firewall. To be useful, all networks willing to support a trace capability would have such a server. The trace server would keep track of all incoming and outgoing traffic, and reduce those exchanges to time-stamped records with origin, destination, and a message digest, such as the fingerprint mentioned above. All this information would be encrypted using the

15. Ibid.

public key of a clearinghouse.[16] A tracing request would originate from a subscriber to the clearinghouse, again, encrypted using the public key of the clearinghouse with a query stating the perceived origin of the attack, the date/time range, and a message digest of suspect communications. The automated system at the clearinghouse would then begin a series of queries to trace servers of networks implicated in the attack. The automated clearinghouse matches up the outgoing traffic of one network with incoming traffic of another, tracing the communications until the point of origin is reached.

The communication channels and information resources used for coordinating investigation of attacks must be separate from the information resources that are the targets of attack, and they must receive special protection. This can be accomplished through an overlay on the network, but its functionality and points of origin must be limited to avoid compromise of the overlay itself. Clearly, the design and operation of such channels is a matter for international cooperation. As noted earlier, a need exists for anonymous communications between incident responders under some conditions, which suggests that care be taken in implementing such "back-channel" facilities.

Integration of Defensive Technologies

Current defenses against a cyber attack include prevention mechanisms such as firewalls, intrusion detection and response components, and security management applications, but a lack of communication and coordination between vendors' security components limits their effectiveness in large heterogeneous environments. Key technical and organizational issues limiting coordinated cyber defense across administrative and national boundaries can be identified, and challenges

16. A clearinghouse is an objective third-party organization that is considered a trusted recipient of information from member parties. A clearinghouse makes relevant information available to all parties without divulging source information that could violate the privacy of participating members.

in achieving agreements between international organizations on how these technologies can be integrated are substantial.

Automated response to intrusions is a major need for defending critical systems. Vendors have developed products that support intrusion response.[17] These products use proprietary protocols and are limited by an architecture that requires all response decisions to be made at a central controller. Because an adversary can take actions at computer speeds, systems must react at comparable speeds, implying the absence of human intervention.

Current tracking mechanisms have significant limitations, especially when applied to large heterogeneous environments such as the information infrastructure.[18] First, intrusion detection systems detect local intrusion symptoms and can only react locally, for example, by reconfiguring local boundary controllers and hosts. Because an attacker may cross many network boundaries, a local response by the target cannot identify or mitigate the true source of the attack. Second, even if intrusion detection systems were capable of communicating with boundary controllers near the attacker, there is no common language for remotely instructing them to handle selected traffic. It is also unlikely that intrusion detection systems would know enough about all such devices to be able to reconfigure them remotely using low-level, device-specific commands. Nor is it likely that the owners of such devices would allow it. Third, if intrusion "symptoms" are detected in different areas of an internetworked environment by different intrusion detection systems, current technology lacks the infrastructure and protocols for pooling this information to allow intrusion correlation and to develop and promulgate a coordinated response.

Current research is providing a framework that allows the inte-

17. See: Network Associates, "Active Security," available at ⟨*http://www.nai.com/asp_set/products/tns/activesecurity/acts_intro.asp*⟩; Internet Security Systems, "RealSecure," available at ⟨*http://www.iss.net/prod/*⟩; AXENT Technologies, "Intruder Alert," available at ⟨*http://www.axent.com/product/smsbu/ITA/*⟩.

18. See Randall Smith, "Coordinated Cyber Defense," presentation at the Stanford Conference, December 6–7, 1999.

gration of detection and response components, thereby enabling experimentation with automated response strategies.[19] The Intruder Detection and Isolation Protocol (IDIP) has been shown to be capable of providing cooperative tracing of intrusions across network boundaries, blocking intrusions at boundary controllers near attack sources, using device-independent tracing and blocking directives, and centralizing reporting and coordination of intrusion responses.

Figure 4 shows system architecture that incorporates the IDIP. Each network in an administrative domain—for example, a company intranet—has an intrusion detector. Networks are connected through boundary controllers. In the example shown, an attacker, having legitimate access to his own network, intrudes upon two more remote networks where he is noted as an intruder. Intrusion reports are forwarded to a central discovery coordinator who is able to discern the attack path. The discovery coordinator issues a response instruction to the boundary controller closest to the attacker. The attacker is thereby identified and either an automated or a manual response can be initiated. Software components that have been successfully integrated using the IDIP are shown in Table 2.

To support communication between the varied IDIP components requires a flexible and extensible language. IDIP uses the Common Intrusion Specification Language (CISL) developed using the Common Intrusion Detection Framework (CIDF) as the language for describing attacks and responses.[20] This language includes terms for describing the blocking actions used in the current IDIP implementation, and it can be extended to support additional responses as they are developed. Currently, IDIP uses only two actions: block and allow. These can be

19. See *Protocol Definition Intruder Detection and Isolation Protocol Definition*, Interim Technical Report (CDRL A005), Boeing Document No. D658-10732-1, January 1997; *Dynamic Cooperating Boundary Controllers*, Final Technical Report (CDRL A003), Boeing Document No. D658-10822-1, February 1998; *Adaptive System Security Policies Preliminary Assessment* (CDRL A005), Boeing Document No. D658-10821-1, February 1998.

20. See Rich Feiertag et al., "A Common Intrusion Specification Language," June 1999, available at ⟨*http://www.gidos.org/*⟩.

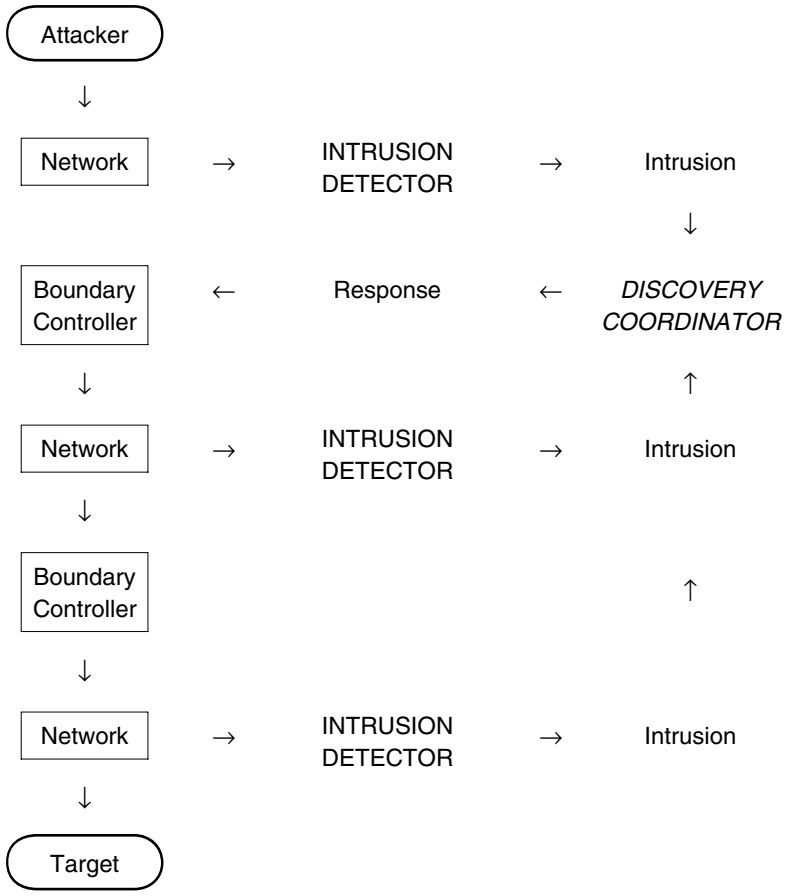

Fig. 4. Intrusion Detection and Isolation Protocol (IDIP) system architecture.

used with various objects (for example, users, processes, messages, or connections) to cause a number of different responses. A "block user" message, for example, is interpreted as a request to stop that user from doing anything; a "block user and connection" message is interpreted as a request that the user be prevented from using the specified connection. Connection information includes protocol, source address, source port, destination address, and destination port. Response messages can also include a specification of when to start and stop such actions.

An IETF working group is currently investigating standards for communications between intrusion detection components. One of the proposed standards is CIDF. The major modifications needed are a limitation of the type of information communicated, definition of data formats and exchange procedures for sharing information of interest with intrusion detection and response systems and the management systems that may need to interact with them, and the integration of the protocol into the TCP/IP suite of protocols. The requirements specification is currently an Internet Draft and has been forwarded to the IESG for publication. Other documents to be produced relate directly to the protocol: a definition of the data items desired in the messages to be exchanged; a definition of the message format; and a protocol for communicating the messages.

Coordinated Response to an Attack

Responses should be in proportion either to the damage already done or to the potential for future damage. Thus damages need to be assessed. When the attacker is inside, no external aid need be solicited, unless domestic law enforcement or other investigatory organizations are brought in. If the attacker is outside, the first question is whether he is located in the target's country; for if the attack is from or through another country, international cooperation will be needed. Furthermore, if the foreign country is the point of origin of the attack rather than a pass-through country, response will be different, because the

TABLE 2
Software Components That Have Been Integrated Through the Intruder Detection and Isolation Protocol (IDIP)

Boundary controllers	*Intrusion detection systems*	*Host-based responders*
NAI Gauntlet™ Internet Firewall [a]	Net Squared Network Radar [b]	NAI Labs Generic Software Wrappers Prototype [c]
Secure Computing Corporation Sidewinder™ Firewall [d]	SRI EMERALD BSM and EMERALD FTP Monitors Prototype [e]	TCP Wrappers [f]
Linux Router [g]	UC Davis Graphical Intrusion Detection System (GRIDS) Prototype [h]	IP Filter [i]

[a] See ⟨*http://www.nai.com/asp_set/products/tns/intro.asp*⟩.

[b] See ⟨*http://www.NetSQ.com/Radar/*⟩.

[c] See T. Fraser et al., "Hardening COTS Software with Generic Software Wrappers," Proceedings of the 1999 IEEE Symposium on Security and Privacy, Oakland, Calif., May 1999.

[d] See ⟨*http://www.securecomputing.com/*⟩.

[e] See Ulf Lundqvist and Phillip A. Porras, "Detecting Computer and Network Misuse Through the Production-Based Expert System Toolset (P-BEST)," Proceedings of the 1999 IEEE Symposium on Security and Privacy, Oakland, Calif., May 1999.

[f] See ⟨*http://www.nai.com/asp_set/products/tns/intro.asp*⟩.

[g] See ⟨*http://www.linux.org/*⟩.

[h] See S. Saniford-Chen et al., "GRIDs—A Graph-Based Intrusion Detection System for Large Networks," Proceedings of the 19th National Information Systems Security Conference, October 1996.

[i] See ⟨*http://coombs.anu.edu.au/~avalon/ip-filter.html*⟩.

required investigation can impinge on its sovereignty. Greater cooperation can be expected if the foreign state is itself a victim or is an uninvolved transit country.

Need for a Global Incident Response Capability

Clearly, information exchange and interaction among many parties is necessary for producing comprehensive approaches and solutions to

TABLE 2
(*continued*)

Boundary controllers	*Intrusion detection systems*	*Host-based responders*
NAI Labs ARGuE Prototype [j]	Oregon Graduate Institute StackGuard [k]	
NAI Labs Multiprotocol Object Gateway Prototype [l]	Odyssey Research Associates CORBA Immune System Prototype [m]	
	NAI CyberCop™ Server and CyberCop Monitor [n]	
	Internet Security Systems RealSecure™ [o]	

[j] J. Epstein, "Architecture and Concepts of the ARGuE Guard," Proceedings of the 15th Annual Computer Security Applications Conference, Phoenix, Ariz., December 1999.

[k] C. Cowan et al., "StackGuard: Automatic Adaptive Detection and Prevention of Buffer-Overflow Attacks," Proceedings of the 7th USENIX Security Conference, San Antonio, Tex., January 1998.

[l] See G. Lamperillo, "Architecture and Concepts of the MPOG," NAI Labs reference no. 0768, June 1999.

[m] See "Computational Immunology for Distributed Large Scale Systems," available at ⟨*http://www.oracorp.com/Projects/Current/CompImm.htm*⟩.

[n] See ⟨*http://www.nai.com/asp_set/products/tns/intro.asp*⟩.

[o] See ⟨*http://www.iss.net./prod/*⟩.

system intrusions. The need is to support a global incident response effort and thereby to reduce the number and extent of computer security incidents. The Computer Emergency Response Team Coordination Center (CERT/CC) at Carnegie-Mellon University provides one basis for such a capability.[21] This Center, established by the U.S. Department of Defense in the late 1980s, has extensive practical experience in the conduct of violations of computer security. Incident response and computer security teams consist of practitioners and technologists who have operational experience but may lack authority to make policy and security decisions for their organizations. A re-

21. See Longstaff, "International Coordination for Cyber Crime and Terrorism."

sponse team may not have sufficient staff to respond effectively to all security incidents. At this time there is no infrastructure to support a coordinated global incident response effort, although there are some components that could form the basis of such an infrastructure.

A variety of issues must be addressed when considering how to promote an effective global incident response infrastructure. These include which organizations will coordinate and participate in the development effort, how current groups and forums can fit their missions and objectives into an agenda to create a global infrastructure, and what possible structures and mechanisms might be required in the future.

The 1999 Melissa virus attack underscores the lack of such a global response structure for incident response. Because individual teams focused on individual or national response needs, there was no operational global response effort. Although the Forum of Incident Response and Security Teams (FIRST) played an essential role in the early identification of the problem through reports shared among its member teams and was therefore able to notify others, it lacks the operational mission and funding necessary to facilitate further responses; almost four days elapsed from the initial activity report to solicitation and receipt of status reports and generation of the global activity summary. Even so, the resulting summary provided the needed global perspective on impacts and spread of remedial activity.

A global response capability can be achieved by building on existing incident response and security teams. Successful resolution of international incidents has been possible when the following elements were in place:

- A common terminology between parties involved in the incident to include identification of the intruder's modus operandi, the technical attack details, and the identification of the targets
- Knowledge of the technical skills of all parties involved in resolving the incident

- Existing agreements on how incidents of a variety of types will be handled
- An understanding of the common and conflicting societal issues surrounding the incidents

Such an approach is not necessarily dependent on international agreements or treaties between governments, but government-to-government agreements can significantly improve the effectiveness with which incidents are resolved.

FIRST is a possible basis for such an expanded international technical cooperation. Organized in the early 1990's, FIRST consists of more than eighty incident response and security teams from nineteen countries. It provides a closed forum for these teams to share experiences, exchange information related to incidents, and promote preventive activities. Although other teams exist that are not yet FIRST members, and new teams are constantly being established, the labor-intensive nature of incident response and the growing number of incidents leave the world with a dearth of capability. FIRST is a voluntary organization that provides an introduction service and meeting place for teams to establish trusted interactions, but since it lacks operational elements it cannot provide the necessary coordinated global effort or meet other needs, such as a more open flow of sensitive information and close collaboration to respond to widespread events. Overcoming these shortcomings will require appropriate policies and procedures, formal contractual agreements among its member organizations, and documented procedures to serve as guidance for new entrants.

Beyond formal structure, what is needed is a way to build on personal trust relationships to achieve organizational trust. Gaining entry to the incident response community can be a difficult and lengthy process; the community is ready to embrace new members, but it is wary of interacting with new teams until an existing member of the trusted community can vouch for them. A global incident response capability is difficult to build rapidly, but national boundaries, which

provide a demarcation for policies, procedures, and jurisdiction for information exchange, are a natural starting point. Response teams that cross national boundaries, such as incident response teams for multinational corporations, are another useful basis for international cooperation. Above all, there must be participation and cooperation among governments, law enforcement agencies, commercial organizations, the research community, and practitioners who have experience in responding to computer security incidents.

Different Viewpoints of Victims and Law Enforcers

Rapid collection of forensic evidence is needed by both victims and police.[22] Police seek to identify the attacker; the victim has the task of cleaning the attacked hosts and getting them back into operation. With limited technical resources, the defender's efforts must be divided among learning the extent of the invasion, reconfiguring hosts to be resistant to future attacks, getting the hosts back on-line, and helping law enforcement agents track down the attackers. Even though the evidence gathered at these first steps is often too vague to prove a defendant's guilt, it can provide probable cause for further investigation. Rarely does an attacker explicitly give away his or her identity. Either the "smoking gun" evidence is found on the attacker's own computer, or is observed through interception of a data stream while a crime is being committed.

Once law enforcement agencies have collected evidence from the scene of the crime, the evidence must then be combined with the evidence collected by the defender. This should show that the evidence collected from the crime scene is directly tied to evidence collected from the intrusion site. Items found at the scene such as lists of usernames and passwords, computer and network addresses, help screens from attacked applications, and so on, can be correlated with evidence stored by the defender that documents the intrusions. To be usable in

22. See William Cheswick, "Internet Forensics and Cyber Crime in Court," prepared for the Stanford Conference, December 6–7, 1999.

court, evidence collected by the defender must be properly collected and stored.

Usually law enforcement and the victim will keep a write-locked copy of each disk image dump. A growing number of tools are available for examining and processing image dumps. Speed is essential because ISPs generally keep their logs for only a few days. In the U.S., ISPs generally require a subpoena before supplying log data to law enforcement agencies. But they will preserve log data, which may normally be kept for only two or three days, in anticipation of a subpoena. Evidence collected at the scene of the crime will provide additional clues for what to look for in the defender's historical data, so that eventually, a comprehensive profile can be developed to reanalyze the data, documenting information to support prosecution.

Computer Forensic Issues in Law Enforcement

Log-keeping is an important part of dealing with the Internet. Logs help identify usage patterns, administrative and configuration errors, misuse, and attacks. Mailers keep logs to help identify sources of spam mail. Firewalls log rejected packets. Authentication servers record account usage, and DHCP servers record caller ID information, accounts, and IP addresses assigned. ISP records of this sort are particularly important in tracing attacks back to their source. Such logs, kept in the ordinary course of business, are admissible in court. However, since nearly all computer forensic evidence is machine-readable, it is subject to easy and undetectable editing. Governments have to deal with this obvious possibility.

Providing access to logs is a source of tension for ISPs, particularly those who do not wish to become involved in legal actions. If logs are discarded routinely, without backup, the investigatory process will have less information to utilize. Logs to handle routine problems such as mailing errors are seldom needed for more than a week. On the other hand, firewall logs of suspicious activity are preserved on a WORM drive, where they remain available indefinitely.

Bulk backups and disk image copies may provide usable and admissible evidence if the chain of evidence is preserved. CD-ROMs are useful for preserving evidence, though they lack the capacity to deal with current online storage technology. Newer technology such as write-once DVD disks should help. Image backup tapes need a write-protect switch to prevent inadvertent overwrites that can be sealed at the time of the dump. Although a switch can be defeated with a modified tape reader, there are cryptographic solutions to this problem as well, since cryptographic checksums verify that data have not been tampered with. There are also time-stamping services that can verifiably time-stamp a checksum without revealing the actual data.[23]

Possession of username and password files is illegal in the U.S.[24] Password files are access devices, and the mere presence of several of these files on a defendant's computer is illegal, even if there is no evidence that they have been cracked. However, the use of a username/password pair is not proof that the owner is at fault, since accounts are easy to steal and many sites offer free e-mail accounts with user-selectable account names. Nor is possession of code evidence that it has been used. Although idiosyncrasies of code may be suspicious, it has been difficult for law enforcement agencies to prove that particular code was actually used. This has been a crucial problem in prosecutions. Code idiosyncrasies may strengthen a case, but software is often widely known. Further complicating prosecution is the fact that cyber crime usually involves innocent third parties.

Issues Surrounding ISPs

Internet Service Providers are the entry point to cyberspace. On one side of the ISP is the "user," the arena of private property, civil rights against unreasonable search and seizure, rights to privacy, and due process. The other side of the ISP can be characterized as "commons,"

23. See S. Haber and W. S. Stornetta, "How to Time-Stamp a Digital Document," *Journal of Cryptology* 3 (1990/91): 99–112.

24. 18 U.S.C. § 1029.

something shared, and hence something where the rights of various entities are not absolute but must be balanced against the common good. Damage to the commons affects all who use it. This balancing of rights and responsibilities is a matter of process, which can be voluntary or may be subject to various domestic and international laws and agreements.

Law enforcement agencies need help from ISPs, regardless of their location. They will want real-time access to packet streams and authentication to tap specific sessions, giving stronger links between the user and criminal activities. Some ISPs assist in these matters when they can, but it is a difficult job. The growth of the Internet leaves hardware running at full speed, with few spare facilities for this activity. For a busy router, some kind of hardware assist will be necessary, and this can only be provided by the router manufacturers, and only in response to ISP or legal requirements. Since this would increase the costs of the router, it may take legislation similar to the CALEA requirements for the telephone system.[25] Such requirements would have to be international to be effective, and in the long run they will probably not work, for the ubiquitous encryption that is coming will frustrate many of these efforts. The new generation of CPUs, driven by such needs as voice recognition and game graphics, have adequate power to apply strong encryption to network traffic streams, and there is little hope that even a government will have the resources to penetrate these sessions directly. Even weakened or broken cryptography presents a large economic obstacle to real-time wiretaps; 40-bit encryption is considered weak, but it is not easily amenable to real-time cracking. High-performance hardware is required to extract even plaintext packets from packet streams.

The Fifth Amendment to the U.S. Constitution complicates the question of whether a defendant can be forced to reveal passwords and unlock cryptographic keys. Other complications arise when an

25. See Commission on Accreditation for Law Enforcement Agencies ⟨*http://www.calea.org*⟩.

ISP is itself under investigation. Specific circuit identifications can be obtained from telephone companies to determine connectivity, but the extent to which one can trust the logs of an ISP that may itself be compromised, possibly inadvertently, is unclear. Given an IP address, there are still questions about the actual location of the computer at the time of the alleged crime, and the identity of the actual user.

The Wide Range of Responsibilities of ISPs

From the standpoint of the ISP, it must, as a first priority, protect its own hardware, software, and databases from compromise, meet contractual commitments to its customers, maintain the continuity of its business, and guard against liabilities arising from allegations of negligence.[26] Attacking routers and switches can compromise the entire network infrastructure, and such attacks are heavily defended against, though how effectively remains to be established.[27] In addition, ISPs must help protect their customers from accidental and malicious actions on the Internet. Finally, they may be seen to have some sort of responsibility, at least implied, to the global community to protect it from the accidental or malicious actions of their customers. These responsibilities are difficult to fulfill in the face of rapidly changing technologies—which imply frequent upgrades in systems; rapidly growing market demands that require frequent upgrades in capacity, rapid changes in the ability of hackers and criminals to compromise networks and computers, and rapid changes in security technology that must be assessed and in which investments must be made.

ISPs cannot ignore security issues, but selecting and implementing appropriate security measures in a timely manner while maintaining high traffic throughput to the Internet nevertheless requires a high degree of cooperation among ISPs and communication providers. The

26. See Barry R. Greene, "ISP Security Issues in Today's Internet," presentation at the Stanford Conference, December 6–7, 1999.

27. See "Improving Security on Cisco Routers," available at ⟨*http://www.cisco.com/warp/public/707/21.html*⟩.

open management environment in which the Internet operates and the dedication of its vendors and operators to meeting the needs of its users requires the balancing of these opposing tendencies of cooperation in protecting the commons and in competition among themselves. The maintenance of this environment is under severe pressure, however, as the Internet and the number of its users expands. A central consideration, as we move to protect the Internet and its users, is that while doing good, we should also do nothing that will limit its potential for continued growth.

A Proactive Program for Internet Security

Given these considerations, a statement of best common practices for ISPs is needed.[28] Such a document, or family of documents addressing recommended practices to various degrees of depth, should be prepared and updated to reflect current business and technical trends. This is the proper task of an industry or trade association. Adherence would be voluntary, although in view of the tradition of service to Internet users, one might expect it to be adopted for reasons of efficiency and economy. Should risk management through insurance become widespread, such best common practices could naturally assume the role of minimum standards for insurability and protection against allegations of negligence.

Security research and product development is undertaken by ISP hardware and software vendors and this can be expected to increase the level of protection in deployed information networks. A desirable result would be for differences in security to become a market differentiator for ISPs, much as price, quality and reliability of service, and ease of use are today. There are a number of vendor roles including close interaction of router vendors' operations staff with those of ISPs; providing personnel for product support emergency reaction teams; having product development staff working with customers on new

28. See "BCPs for ISPs—Essential IOS Features Every ISP Should Consider," available at ⟨*http://www.cisco.com/public/cons/isp/documents/*⟩.

features; providing security consultants for assistance with countering attacks, undertaking audits, and prosecuting intruders; and staff who track hacker communities.[29]

3. Automation of Computer and Network Protection

Labor-intensive approaches to computer and network security can use automated tools and techniques to improve their speed and efficiency. Timeliness will minimize losses due to attacks; efficiency will enable more protection to be provided for a given level of resources applied; scaling to accommodate the growing number of attacks can, in principle, be achieved; and privacy for innocent users can be enhanced to a degree through automated rather than manual screening of traffic. All too frequently, the requirement for privacy conflicts with the need for protection; once people have exhausted their ability to protect themselves, they must appeal for assistance, and this inevitably involves some sacrifice of privacy. Nevertheless, it is reasonable for victims of cyber intrusion to expect their protectors to tread as softly as possible, and that those seeking protection have options as to how much protection they will receive and what price, in terms of loss of privacy, they are willing to pay.

Tools to Automate Protection

Automated tools can improve protection in various ways. For example, if information is only accessible by reference to its unique identifier or handle, the ability to collect and analyze data on system use can be greatly augmented. At a minimum, the ability to establish the presence of a user at a given place in the global information system will be critical for evidentiary purposes, recognizing also that the ability to

29. See "Product Security Incident Response Teams (PSIRT)," ⟨*http://www.cisco.com/warp/public/707/sec_incident_response.shtml*⟩.

challenge such evidence and contest differing points of view or interpretations will also be needed.[30]

One class of tools, noted earlier, are those that detect intrusion and other unauthorized uses of network and computer systems and high-level services. These will range from the more obvious tasks of checking log-ins to be sure they match authorized users, to running software agents on the machine to detect other software agents that may arrive without authorization. Agents that are able to detect the presence of, and oppose, other agents may be a possible countermeasure.

Another class of tools for protection would monitor current usage of all relevant machine resources and look for unusual patterns. Periodic checks of user identity might be warranted. Compared with traditional password protection schemes, cryptographic log-in systems provide considerable increases in protection against unauthorized access. Unlike the traditional password systems, cryptographic log-ins are not vulnerable to playback attacks and other attacks that involve stealing passwords. Alternate systems based on public key encryption can be used to authenticate users.

Automation of protection could be implemented in the network, tracking patterns of usage in real-time and alerting system operators to unusual conditions for manual or semiautomated review. Another automated approach to preserving privacy would be to package suspected sessions and e-mail them to the user to verify that it was indeed that user who was actually at the keyboard for that session.

Timely Tracking

Pursuing an attacker can reduce future exploitation of system vulnerabilities. Rapid response will minimize compromise and contain damage that may have occurred. This would require the cooperation of

30. See Joseph Betser. "Tracking Cyber Attacks," presentation at the Stanford Conference, December 6–7, 1999. See also the contributions by Rizzi and by Kahn and Lukasik.

automated software modules en route from and at the attack source. It would also require trust among the cooperating organizations and technical activities that are in place to facilitate such automated communication among software modules.[31]

There are several meanings of "timely." One is "session time," the time during which the intruder is logged-in. Information collected during this time will enable tracking most easily since all the links in the attacker's path are open, but this sort of tracking requires not only an unusual degree of readiness but also technical capability. Once the intruder is no longer on-line, traces of the surreptitious activity, other than any changes made or code purposely left behind, are—if the intruder is skilled—likely to have been erased.

A second time period of importance is the transaction clearing time, that is, the time between the on-line action by the intruder and when the intruder's desired goal is achieved. This will depend on such things as the organizations and business processes involved, the calendar date, and the objectives of the attack. In some cases, the intruder can achieve the goal while still on-line; in others, actions will be required by other organizations to bring the act to fruition.

Another time period is the "revisit" interval of the attacker. From penetration experiments, we recognize that attacks are not single isolated events but frequently consist of multiple intrusions to collect information about the system, to undertake various test and practice actions, or to exploit a vulnerability repeatedly.[32] For the intruder,

31. See M. Wood, "Intrusion Detection Message Exchange," IETF Draft, October 1999; D. Schnackenberg, K. Djahandari, and D. Sterne, "Infrastructure for Intrusion Detection and Response" (forthcoming: DARPA Information Survivability Conference and Exposition (DISCEX), Hilton Head Island, S.C., 2000); R. Smith et al., "Multi Community Cyber Defense," DARPA Information Assurance and Survivability Principal Investigator Meeting, Phoenix, Ariz., August 1999; Smith, "Coordinated Cyber Defense."

32. See Raymond Parks, G. Schudel, and Bradley Wood, "Modeling Behavior of the Cyber-Terrorist," presentation at the Stanford Conference, December 6–7, 1999; B. Wood and G. Schudel, "Red Team Experiments 9901 and 9907," DARPA Information Assurance and Survivability Principal Investigator Meeting, Phoenix, Ariz., August 1999.

success means achieving the objectives of the attack and remaining undetected. Though one may have multiple opportunities to detect an intruder, one may have to sustain repeated successful intrusions before succeeding in plugging the hole.

A fourth measure is the time needed by the attacker to evade pursuit, such as to shift from one location or jurisdiction to another. There might also be an explicit threat time, such as occurs with an announced deadline.

Next Generation IP Protocols

The output of the IETF working groups in the security area may have positive effects on protecting information systems.[33] A major development that will help is the implementation of IPv6, the next generation of the Internet Protocol.[34] IPv6 addresses several issues that can have a significant impact on network security. First, IPv6 expands the address space for network device addresses from 32 bits to 128 bits. A problem with the depletion of addresses in IPv4, the current version of IP, is that in order to support new network hosts, various workarounds have been required, resulting in a situation where not all hosts have unique IP addresses.

IPv6 has another feature likely to prove powerful in tracing the origin of a message: the "hop-by-hop" header, which allows each of the routers along the delivery path to exercise certain options. An obvious option that would enhance tracing activities is to have each router that has forwarded the message record its address in the message header. The problem with using "hop-by-hop" in this manner is that each packet will be modified by a number of devices enroute, with the distinct possibility that one of those devices, for example, controlled by an attacker, could change the routing history.

Another approach that could be used is controlling the intercon-

33. See Erlinger, "Internet Protocols for Protection Against Cyber Crime."

34. See S. Deering and R. Hinden, "Internet Protocol, Version 6 (IPv6 Specification)," RFC 2460, available at ⟨*http://www.ietf.org/rfc/rfc2460.txt*⟩.

nection of routers to the point where routers will only accept traffic from certain other "trusted" routers or users. Although this changes the way routing is currently done, it could provide enhanced security over the current practice.

Encryption is another method of securing network activities from cyber crime. In an environment where each transaction between hosts is encrypted, there is a guarantee that the source and destination are known. This allows traceback, but only if we are willing to support strong encryption and incur the overhead of encryption on network interactions.

Facilities for Internet Monitoring

There is also a need to enhance security by monitoring for persistent but marginal internal network problems from locations outside the network. For example, a network that consistently loses one percent of all packets sent over it may appear to be working well, but the small loss can be important. The loss may be due to a failing component in the network, but it could also be caused by an insider who has altered the network without authorization. Collectively, other networks connected to such a network could federate to detect and report the problem if they had access to all its inputs and outputs, or the network itself could detect the problem if it insisted on the equivalent of double-entry bookkeeping. Some of these problems can be alleviated with more effective cooperation among ISPs, as is happening within the Internet operators (IOPS) organization.[35] IOPS consists of many of the largest national and international ISPs who collaborate to prevent or alleviate problems in the Internet and routinely share information on a confidential basis. In addition, they seek to improve performance and efficiency where such improvements require collaboration.

Whether part of tagging and tracking or as passive observation in a statistical sense, monitoring Internet traffic statistically will require

35. See ⟨*http://www.iops.org*⟩.

an embedded data collection infrastructure that does not currently exist.[36] Modification of router and switch hardware and software is technically possible, although it is unclear whether vendors would be interested. Because of concerns over privacy, it will be even more difficult to secure user buy-in. Regulation is one way of addressing user reluctance; market incentives, such as insurance, are another.

Even granted that Internet monitoring is politically and economically feasible, much detailed technical work must be undertaken. Requirements must be defined, monitoring facilities must be designed and deployed, and an organized data collection operation must be managed.[37]

Scaling and the Need to Provide Information to Aid in Establishing Priorities

Providing international assistance involves more than a potential incursion on a nation's sovereignty and on its citizens' privacy. It will require human and technical resources that are often in short supply. If cyber attacks continue to increase, and if the number of affected countries increases, and as all equip themselves with intrusion detectors, the number of requests for assistance will grow substantially. Furthermore, owing to the presumed time urgency of the requests, responding to such requests will have disruptive effects on the nations involved.

Consider, for example, data reported in recent U.S. GAO reports.[38]

36. Claffy, "Traffic Observation in a Stateless Data Networking Environment."

37. See "High Performance Networks: Measurement and Analysis Collaborations," Workshop, June 29–30, 1999, and "Challenges and Opportunities for Measurement and Analysis in a High Performance Computing Environment," Workshop, July 1, 1999. Both workshops were sponsored by the National Science Foundation and hosted by the San Diego Supercomputer Center, San Diego, Calif. See also ⟨*http://www.caida.org*⟩ and ⟨*http://www.nlanr.net*⟩.

38. See *Information Security Computer Attacks at Department of Defense Pose Increasing Risks*, GAO/AIMD-96-84, May 22, 1996; *Information Security Opportunities for Improved OMB Oversight of Agency Practices*, GAO/AIMD-96-110, September 24, 1996.

It is believed that in 1995 there were roughly 250,000 penetrations of computer systems owned by the U.S. federal government alone. Based on controlled penetration testing, it is estimated that 64 percent (160,000) of these were successful. The GAO also estimated that only 1–4 percent of these attacks were detected, and only a quarter of those detected were reported. Based on a database of 30,000 incidents, Thomas Longstaff notes that 40 percent have a foreign component.[39] This would suggest that 600 attacks per year might qualify for requests for foreign assistance in tracking intruders. If, as estimated by the GAO, the number of such attacks is doubling annually, that many more of those detected could be reported, that many more penetrations could be detected with the deployment of new intrusion detection technology, and that these numbers only cover the intrusions into U.S. government computers, the obvious conclusion is that the number of requests in the future for assistance from other countries could easily exceed the ability of those called upon to respond. Therefore nations will need some kind of criteria for assessing the impact of attacks in order to evaluate the seriousness of intrusion events in the light of available investigatory resources.

Reconciling the positions of those who see governments as part of the solution and those who see governments as part of the problem will not be easy. Making Internet protection architectures and operations public will reduce the anxiety of some, but possibly at a certain cost in the effectiveness of the protection offered, and it may complicate the collection of forensic evidence. Since much of the burden will fall on ISPs, we have to evaluate both the economic cost-benefits and the social cost-benefits in deciding how best to strengthen the protection of the commons.

39. See Longstaff, "International Coordination for Cyber Crime and Terrorism."

4. The Need for Cooperative Action

The common situation, where the victim and the attacker are in different sovereign jurisdictions, or where the attacker has transited other sovereign jurisdictions, will require agreements on a variety of subjects. In light of the speed with which computer-mediated events unfold, the processes employed by those seeking to understand the nature and source of the attack will require prior agreements if the operation of systems is to be restored in a timely manner and damage minimized.

An International Forum in Which Diverse Stakeholders Can Interact

Since, unlike jurisdictions based on national boundaries, the digital information infrastructure does not have a central location in the physical world, responding to attacks not only is difficult technically, it also limits the use of accepted methods for practicing law enforcement. Recent G-8 and OECD activities are examples of increasing recognition of this international problem. Improving critical information infrastructures requires involvement of diverse parties, including governments, policy- and lawmakers, law enforcement, software vendors, the research community, and practitioners, such as FIRST members, who have experience responding to computer security incidents.[40] Attempting to address the problems in one group without input and feedback from the others can result in incomplete solutions. For example, recent U.S. legislation, the Digital Millennium Copyright Act, resulting from the World Intellectual Property Organization (WIPO) treaty, generated concern within the Internet security community. Practitioners, researchers, software vendors, and incident response teams noted the legislation could limit some aspects of their efforts to address security flaws and reduce risk to critical infrastructures. Though this was clearly not the drafters' intent, it is an example

40. Ibid.

of the need for ongoing communication among policymakers, technologists, and others to ensure that future policies and agreements on national and international scales are practical and effective.

Frequently, information is reported to incident response teams that does not involve a specific victim or computer security incident but does indicate that some activity may be ongoing in some part of the community. Currently, there is no standard way of sharing that information, although sometimes this information is posted to a shared list of FIRST members. The sharing of information relating to emerging technical threats is increasing and this trend is to be encouraged. The technical forums that FIRST sponsors on a periodic basis to discuss recent developments in technical threats and vulnerabilities have been effective for the incident response community, but they are not open to other international experts to relate the technical trends to broader international concerns.

The tracking of sources of even small amounts of traffic is likely to be important in locating perpetrators of crimes. Hints of activity in the form of programs left behind, usernames assumed, methods of operation, and so on are all likely to be of significance in determining identities and locations of attackers. Systematically checking for such information, as well as other best practices is required. Exploiting these understandings of attacker behavior can assist in defending against them, a technique Raymond Parks refers to as "dynamic defense."[41] Therefore it is important to exchange both general and specific information on attacker modus operandi to provide the greatest degree of protection for the greatest part of the global information infrastructure.

A best-practices document prepared by an international group of experts and updated periodically would assist in establishing de facto standards. Conforming to best practices should be a part of justifying international cooperation in obtaining redress.

A database of attacks and known viruses has been compiled by

41. Parks, Schudel, and Wood, "Modeling Behavior of the Cyber Terrorist."

organizations such as the CERT/CC at Carnegie-Mellon University's Software Engineering Institute. It would be helpful to restructure this information so that it can be used directly by automated tools to detect patterns of criminal activity.

Prioritizing Requests for Assistance

Not every case of cyber crime will require international cooperation, nor will every case be equally deserving of such cooperation; it will be necessary, as Kahn and Lukasik have suggested, to establish priorities, with each request that is made being subject to some form of evaluation.[42] As a minimum, and assuming assistance is merited on the basis of the magnitude of the attack or the extent of the loss, the following questions are likely to be raised:

1. Due Diligence. Has the requesting organization conformed to best practices as formulated for its industry?
2. Rapid Response. Has the requesting organization implemented near real-time monitoring and auditing of its information systems?
3. Potential Impact. Has the potential impact of the intrusion been evaluated and ranked in terms of importance to enable an assessment of the degree of international cooperation that is justified?
4. Probable Cause. How has the requesting organization established that the intruder used the facilities of the country whose assistance is sought?

Although all signatories to an international agreement will have a right to assistance under its terms, there are practical limits on what can be reasonably provided. If information systems continue to remain poorly protected and if their vulnerabilities are increasingly exploited,

42. Kahn and Lukasik, "Fighting Cyber Crime and Terrorism."

the need for assistance will greatly exceed what can be made available, and the necessary expertise will become a rate-limiting factor in the resolution of intrusions. In such a case, some form of rationing or prioritizing of assistance can be expected. It will be useful to factor such prioritizing criteria into an agreement to encourage improving the state of self-protection throughout the world. Thus, apart from the direct assistance that can arise from an international agreement, the long-term systemic improvement that can be thereby facilitated is an important goal.

Internationally Agreed-Upon Means to Validate Information

Certificates can be used to authenticate information as well as users. A piece of information in digital form can be cryptographically "fingerprinted" and the result attached to the information or stored separately from it. The certificates can be used to verify packet fingerprints, which in turn will verify the underlying information.

Another method of validation could rely on encrypted archived "snapshots" of critical information provided by or taken from key locations in the net or even from user systems. The archive could be run by a trusted third party, who would warehouse the information for whatever period was deemed appropriate. This information could be retrieved and decoded after the fact to provide insight into problems and to corroborate other evidence.

Automated Cooperation Beyond Local Administrative Domains

Global cyber defense will involve the sharing of cross-organizational intrusion information and arriving at cooperative responses. The mechanisms for doing this must be capable of being tailored to protect sensitive information and must allow organizations to manage their trust relationships. An essential part of such cooperation is the ability

to recognize when multiple parts of the global infrastructure are simultaneously under attack.

The Defense Advanced Research Projects Agency (DARPA) program in Multi-Community Cyber Defense (MCCD) is directed to identifying the primary barriers that limit effective sharing of attack-related information with neighboring organizations and mounting a coordinated defense against detected attacks. This work seeks to extend the IDIP by focusing on three key areas: (1) providing local administrative control over the release of their internal attack-related information, including sanitizing data prior to release; (2) establishing and maintaining trust relationships between organizations; and (3) developing higher-level capabilities for conducting attack analyses and data fusion.[43] These capabilities will be integrated into the IDIP framework by adding additional functionality to the IDIP components and by extending the message language used to communicate between components.

To provide strategic defense of critical infrastructure requires that organizations that do not normally share information are able to cooperate in responding to attacks. In the DoD, problems arise when communication occurs between classification domains or across coalition force domains. In the commercial arena, organizations are mutually suspicious. Even when organizations work together as partners, they must protect various types of proprietary information. In either case, there are problems in both releasing data to, and accepting response directives from, remote domains.

In a cooperative environment, the following potential types of information could be sent between administrative domains to improve intrusion detection, correlation, and response:

(a) Near real-time information and requests, including (1) attack notification, and (2) response recommendation

(b) Slower, but still immediate information and requests, includ-

43. Smith, "Coordinated Cyber Defense."

ing (1) correlation results concerning an immediate problem, and (2) discovery coordinator requests for immediate action

(c) Human-speed information and requests, including (1) policy information, (2) correlation results, and (3) other information to be used by higher-level correlators, and alerts from higher-level strategic warning systems

All but the human-speed actions are intended to be sent automatically between administrative domains. The last type would instead be handled by a "trusted" third party that would provide a global situation awareness and response capabilities. In addition to enhancing the IDIP framework to include this type of message, each domain must have the ability to establish a policy for information sharing and for verifying that the policy is implemented.

The likelihood that remote organizations will not completely trust each other, or will not be able to share information fully because of policy, requires more constrained information flows at organizational boundaries—such as, for example, the IDIP requests for severity and certainty fields. The severity field indicates the degree of potential damage the requester might suffer if the attack continues, the certainty field indicates the degree of confidence the requester has that its detection mechanisms have detected a bona fide attack. Together, these fields may reveal the power of the domain's detection mechanisms and the extent to which it can protect itself against various kinds of attacks. That information may need to be sanitized at a domain boundary. Because the remote domain may have fewer safeguards than the local domain, releasing this information can result in giving attackers additional data to be used in an attack. Executing remotely generated response directives requires trusting the originator, or at least establishing controls that limit the damage from untrusted originators. For example, one might take action only if attack information can be corroborated locally.

Blind trust in the results of intrusion-detection algorithms could enable a serious attacker to cause the detection and response system

to misbehave by first penetrating the detection component. Unquestioned erroneous reports from a penetrated detection component could have a number of effects, including shutdown of critical subsystems in response to nonexistent attacks, claims that other detectors have been penetrated and therefore should be ignored, or changes to the state of other detectors, causing them to reduce their warning levels. These issues are being addressed by investigating mechanisms for determining: (1) the current trust in a remote domain, (2) changes in detection and response policy based on the current trust relationships between domains, (3) changes in trustworthiness, and (4) the point when trustworthiness has been reinstated.

Research on cryptographic trust models and fault-tolerant systems can be applied to this area. Results from the cryptographic trust model community can be applied to establishing authenticated identity, but the trust computed for what appears to be authentic data must be modified based on intrusion-related data from other sources. Techniques used in fault-tolerant systems for voting, diagnosing components, and redundancy are directly applicable.

The MCCD architecture, based on an enhanced IDIP framework and integrated with high-level analysis and correlation techniques, should provide the capabilities necessary to enable organizations to implement cooperative agreements on sharing attack-related information, analyzing and identifying serious threats, and executing coordinated responses to detected attacks against global information systems.

Cooperative International R&D to Meet Evolving Threats

The global research and development community must be heavily involved in efforts aimed at protecting information systems and their users from cyber attack. As information technology rapidly evolves, threats, vulnerabilities, and protective measures are also changing. Hence continuing R&D activities are required in order to be able to meet the still evolving threats.

In addition to currently available technology, a number of areas of industrial and academic research that might yield novel paradigms for addressing the rising challenges of information assurance and survivability have been noted:[44]

- Economic, financial, and market-based paradigms. Tapping the checks-and-balances, which are used by the financial community in order to reduce unauthorized activities, and using market incentives to promote proper use of cyberspace.
- Biological immunology paradigms. Adopting some of the paradigms that make the human body successful in identifying and fighting invading bodies, in spite of the complex nature of the biological systems involved.
- Public health paradigms. Computer viruses, and perhaps even the computer-based attack, can be considered an invading disease in the multinational body of the Internet. In public health practice, nations participate collaboratively in the exchange of infectious disease information, and concepts such as quarantine, immunization, and treatment may apply.
- Reliability paradigms. Using experience gained in complex process systems, such as chemical and nuclear plants, should be considered. The large number of dynamic variables and the delicate interplay among them might provide insight for dealing with complex cyber scenarios.
- Correlation paradigms. The defense against coordinated attacks that use large enterprises to execute multiphase complex attacks requires correlation of different attack components in order to detect and defend successfully. The development of fast correlation algorithms and their implementation in systems will be important.
- Expert systems paradigms. Learning algorithms are very useful

44. Betser, "Tracking Cyber Attacks."

in studying the normal behavior of a system and achieving the ability to detect abnormal and anomalous behavior that can occur during a cyber attack.

- Data mining paradigms. Obtaining useful information from voluminous audit logs and event records requires expertise in data mining.
- Control science paradigms. The ability to generate an effective adaptive response to combat the attack in progress. Feedback to the response policy could stabilize and move the system to a healthier state.

A distributed international facility for experiments, tests, and demonstrations of security products and services could speed the transfer of R&D advances. The IDIP technology developed to date can be inserted in a number of commercial off-the-shelf (COTS) components, including intrusion-detection products, firewalls and filtering routers, security management components, and clients and servers. The IDIP software was designed for portability and is currently executing on Solaris, BSDI, Linux, and Windows NT platforms. Operating system dependencies were minimized during the development and have been encapsulated. This design provides an easy, low-cost integration path into COTS products, enabling vendors to adopt the technology with minimal investment. Both widespread integration and acceptance of this framework, and agreements on the international use of this technology, are needed to protect global information systems.

Facilitating Trust

Agreements are on paper, and they are necessary for the reasons suggested. But agreements are implemented by people and people tend not to accept matters at face value; they do not trust people with whom they have not previously interacted with satisfactory outcomes. Thus agreements, while necessary, are not sufficient. Trust is another necessary dimension. One can expect that self-protection technologies,

agreements, and trust together will provide both necessary and sufficient conditions for global security. Ultimately, some trust must be placed in parts of the system—for example in the authentication systems or encryption systems. But because, as with any human system, even those parts could be compromised, we shall still need trust at the individual level as well as at the organizational level. An international treaty, by facilitating increased interactions among all stakeholders can be expected to help in this process.

Clearinghouses, or other impartial third-party nongovernmental organizations, will be important as security risks become more pervasive and more complex. Having a mechanism whereby anonymous interactions and cooperation can take place through a trusted third party would help in these circumstances. This trusted-third-party construct would be part of a larger "trust network" in which communications can be passed without conveying identity. Institution of a "hot line" concept for computer attack might also reduce risk. Such a secure and redundant system could provide signatory nations in a multinational regime with the means to communicate assessments of the intent of detected activities. It could also provide states the confidence needed to collaborate against a common threat.

Reducing "Safe Harbors" for Criminals Through Adherence to International Agreements

Those nations that do not recognize the realities of cyber vulnerabilities—or wish to exploit them—will ultimately become safe harbors for criminals and cyber terrorists. The international community should seek incentives for all nations to participate in international conventions to combat these threats. Such a convention should put "teeth" into what is expected of signatory nations: signatory nations would put pressure on other nations to meet minimum standards for the deterrence of cyber crime and for investigation, prosecution, and extradition, so that, ultimately, nations that refuse to sign the convention could face sanctions, possibly extending to "disconnection" from in-

ternational networks—a sort of cyber isolation. The international convention proposed in this volume would satisfy these requirements.

Implementing Cooperative Actions

Most of the above suggestions, such as exchanging intrusion data and attack profiles, undertaking cooperative R&D, and establishing clearinghouses for anonymous communication, do not require broad international agreements; specific actions could be implemented on bilateral or multilateral bases. But if global changes in the security of information infrastructures are to be achieved, some larger international framework can assist in facilitating cooperation. Elements of such a framework for international cooperation, drawn from various international contexts, are:

- Broad membership, consisting of both the world's most technologically advanced nations as well as developing nations, all of whom share the benefits and the risks of global information architectures
- A voluntary and noncoercive environment based on concepts of consensus and practical experience
- Open technical standards that prevent the manipulation of information technology for unilateral gain
- An open organizational structure that provides opportunities for all constituencies to express their concerns
- A mechanism for providing continuous monitoring of actions that can adversely impact privacy
- Mechanisms for reviewing the state of information technology and its practical implementations to enable the international framework to remain relevant in the light of changing capabilities and requirements
- Mechanisms that can assist in building trust relationships globally

- Funding arrangements that can assist less developed nations in meeting their responsibilities to protect the information commons

A specific proposal that incorporates such features is discussed in Chapter 6 of this volume.

5. Looking Ahead

The problems addressed here derive from a confluence of factors: an increasing social dependence on information-based infrastructures, an increasing complexity of those infrastructures that makes it difficult to anticipate all their failure modes, a growing number of people versed in information technology who can harm information systems, and tools readily available to assist in carrying out malicious acts. These have resulted in alarming rates of growth of system malfunction and system intrusion.

Conclusions

1. The Internet operates by means of a voluntary but structured process that provides the capability, in principle, to respond to changing technology and user needs, including enhanced security. But since the process is driven by its developers and users, the incorporation of security features in the Internet is not assured, however valuable that may be from a public policy standpoint.

2. There are a number of existing techniques computer and information system owners and operators can utilize for their protection. These include firewalls, virus protection software, one-time passwords, encryption, and virtual private networks. They also include instituting and enforcing security policies, real-time and off-line auditing of system operation and use, penetration testing, and implementing data and system backup practices at all levels. But enhanced security brings associated financial and operational costs that can result in lesser levels of security than are technically feasible.

3. Intrusion detection systems are available today and are increasingly being deployed, but at the same time the number of system attacks, already large, continues to grow. Furthermore, tracking intruders is difficult, slow, and uncertain. Hence the current state of affairs, where attack is relatively easy and defense is difficult or absent, leaves the balance very much on the side of the attacker.

4. Ongoing security R&D is pointing to ways of protecting information systems such that detecting, tracking, and identifying intruders will be possible with greater ease and certainty through collective actions by users and service providers. These include deploying new Internet protocols, level-of-service agreements making security a contractual requirement, the automation of advanced intrusion-detection systems for warning and tracking, the integration of security tools to provide more complete capabilities to meet wider ranges of needs, the creation of global incident response capabilities, third-party clearinghouses for secure and anonymous communications among incident responders, the use of digital objects to better define ownership of and appropriate uses of information, and increased capabilities for network traffic analysis.

5. Advances in intruder detection and tracking will aid in deterring attacks by increasing the risk of being caught. They can also be expected to reduce intrusions on user privacy implicit in current tracking techniques.

6. International agreements, both informal and formal, will be needed if information infrastructure users are to receive greater protection than they can reasonably provide for themselves. These include extending intrusion detection to operate across larger domains, development of new Internet protocols, coordinating international responses to global incidents, shared R&D to keep pace with evolving international threats, and collecting and providing attack information to users in a timely manner to allow them to provide for adaptive defenses. The Draft Convention presented in this volume illustrates the kinds of steps that can assist in achieving these capabilities.

7. National policies that encourage the introduction of informa-

tion technology into critical infrastructures, thereby allowing systems of unlimited degrees of complexity and vulnerability to be constructed without corresponding increases in system security, should be examined. It is possible that nations can encourage the evolution of their infrastructure systems in ways that will make them more robust as well as more capable.

8. Today's information infrastructure, which has provided such dramatic improvements in access to information, can usefully be reexamined from the point of view of system architecture. What is needed is to overlay on the current information transfer network an assurance network that makes possible the definition and enforcement of standards of behavior among its users. This assurance network will involve both technical facilities to assist in protecting user rights as well as provisions for allowing operators of the assurance network to establish and maintain concomitant trust relationships that will be necessary for international cooperation.

Short-Term Prospects for Enhanced Security Are Encouraging

Protective actions will take time to implement. A central question then becomes one of relative rates. Since available defensive technologies have not been universally deployed, users can do a great deal in the short term to reduce their vulnerabilities. The pace at which short-term enhancements in system security can be made will depend on several elements: the acceptance by users and system operators that increased spending on security is needed, the deployment of available technologies by both individuals and organizations, improvements in current security products to make them easier to integrate, and the availability of new and more powerful security products and services. Looked at from this perspective, the picture over the next several years merits some degree of optimism because there is so much "low-hanging fruit." The combination of defensive technology and operational

process redesign can accomplish a great deal in comparatively short times.

Long-Term Prospects Are Less Certain

More difficult to assess is how much society as a whole is willing to change in more fundamental ways. Will infrastructure operators realize that their rush to adopt information technology risks system failures that can only be addressed at the level of system architecture? Will utility regulators recognize that security must be on their agenda and that, without private sector initiatives, a more aggressive public posture may be called for? Will we recognize that deregulation without consideration of the architectural issues can have severe unintended consequences? Will law enforcement agencies increase their levels of investigatory and enforcement capabilities, and will legislators appropriate the required resources? And will the nations of the world agree that the protection of the information commons is a shared responsibility?

The highly dynamic nature of information technology is a further complication in the long-term outlook for protection of infrastructures. New technology creates new vulnerabilities, and it increases the power of attackers as well as that of defenders. System and network security is not a problem that can be solved once and for all; it has the measure-countermeasure and offense-defense nature of military competition. From this perspective there is less reason to be sanguine.

A threshold issue for considering fundamental long-term changes in information systems will be that of weighing the cost of ignoring cyber attacks against the cost of actions to reduce the frequency and severity of their failures. There is no simple or obvious answer to this question. Because no fatal infrastructure failures have so far been induced by cyber attack, our only evidence of catastrophic failure is indirect. The rates of attack, and computer crime of many forms, are increasing, in some cases doubling annually. Such strong exponential

increases can rapidly dominate the balance. Unless it can be shown that these exponential growth rates will saturate at some comfortably low level, policymakers in both public and private sectors would be well advised to adopt conservative positions. It would seem prudent to invest now to hedge future downside risks.