Stanford University

CISAC

Center for International Security and Arms Control

The Center for International Security and Arms Control, part of Stanford University's Institute for International Studies, is a multidisciplinary community dedicated to research and training in the field of international security. The Center brings together scholars, policymakers, scientists, area specialists, members of the business community, and other experts to examine a wide range of international security issues. CISAC publishes its own series of working papers and reports on its work and also sponsors a series, *Studies in International Security and Arms Control*, through Stanford University Press.

Center for International Security and Arms Control
Stanford University
320 Galvez Street
Stanford, California 94305-6165

(415) 723-9625

http://www-leland.stanford.edu/group/CISAC/

Public and Private Roles in the Protection of Critical Information-Dependent Infrastructure

Stephen J. Lukasik

May 1997

Prepared for presentation at the Workshop on Protecting and Assuring Critical National Infrastructure, Stanford University Center for International Security and Arms Control, Stanford, California, March 10–11, 1997.

The Center is grateful to the Center for Global Security Research of the Lawrence Livermore National Laboratory for support of this project. The opinions expressed here are those of the author and do not represent positions of the Center, its supporters, or Stanford University.

Public and Private Roles in the Protection of Critical Information-Dependent Infrastructure

Michael M. May, co-director
Center for International Security and Arms Control

Seymour E. Goodman, director
Project on Information Technology and National Security,
Center for International Security and Arms Control

In July 1996 President Clinton established the Commission on Critical Infrastructure Protection, with a charter to designate critical infrastructures and assess their vulnerabilities, to recommend a comprehensive national policy and implementation strategy for protecting those infrastructures from physical and cyber threats, and to propose statutory or regulatory actions to effect the recommended remedies. The charter gives examples of critical infrastructures (telecommunications, electrical power systems, gas and oil storage and distribution, banking and finance, transportation, water supply systems, emergency services, and continuity of government), and also notes the types of cyber threats of concern (electronic, radio-frequency, or computer-based attacks on the information or communications components that control critical infrastructures).

Some of the critical infrastructures are owned or controlled by the government, and hence the government can, in principle, harden and restructure these systems and control access to achieve a greater degree of robustness. However, the president’s executive order recognizes that many of the critical infrastructures are developed, owned, operated, or used by the private sector and that government and private sector cooperation will be required to define acceptable measures for the adequate protection and assurance of continued operation of these infrastructures.

The Stanford Center for International Security and Arms Control (CISAC), as part of its ongoing program on Information Technology and National Security, and the Center for Global Security Research (CGSR) of the Lawrence Livermore National Laboratory (LLNL) are conducting workshops to examine many of the issues connected with the work of the Commission. In addition to the questions of vulnerabilities, threats, and possible remedies, we will discuss the impact on the marketplace of possible protective actions, cost in terms of capital and functionality, legal constraints, and the probable need for international cooperation.

The first of these jointly sponsored workshops was held March 10–11, 1997, and included extensive participation by members and staff of the Presidential Commission. As part of the preparation for the workshop, Dr. Stephen Lukasik took on the task of describing the wide-ranging landscape of this complex problem. Dr. Lukasik, now retired, is a former director of the Department of Defense's Advanced Research Projects Agency (ARPA), a former chief scientist of the Federal Communications Commission (FCC), and has served in various capacities as vice president of TRW, Inc., the Xerox Corp., and the Northrup Corp., and was thus eminently qualified for the task. We believe he has contributed significantly to the study of this problem by providing an overview with emphasis on describing a logical structure that can be used to help understand public roles in the assurance and protection of critical infrastructure systems and hence what options are, at least in principle, available to public policy-makers.

A more comprehensive report covering all the presentations and discussions during the workshop will be issued under a separate cover.

Abstract

The discussion begins with a conceptual framework for addressing the protection of infrastructure systems subject to attacks on their information subsystems. This includes treating the types of infrastructure systems, possible strategies for their protection, and the nature and scale of the attack. Three components of a protection strategy are identified: preventing attacks, limiting the damage in an attack, and ensuring rapid reconstitution of the target system following an attack. The paper concludes with a discussion of public and private responsibilities for infrastructure protection and the identification of a number of areas where public initiatives might be effective. These are ordered roughly in terms of the cost and difficulty of implementation. In addressing the subject, the analysis is from the perspective of minimizing government intervention in privately owned infrastructure systems.

Public and Private Roles in the Protection of Critical Information–Dependent Infrastructure

Stephen J. Lukasik

Conceptual Framework

The subject of this discussion is the protection of U.S. infrastructure systems against attack. Attacks can take many forms, from physical assaults to financial attacks, public relations attacks, and electronic attacks. Electronic attacks can range from direct electromagnetic attack on the sensitive mechanisms that regulate and control the system, as well as what have come to be known as cyber-attacks, where the software vulnerabilities of their distributed information technology-based control components are exploited. The infrastructure systems are those that make up the basic support of society: telecommunication; electrical power generation and transport; natural gas and oil storage and transport; the transportation of people, raw materials, and finished commodities; water supply and distribution; banking and finance; and emergency services such as medical, police, fire, and rescue.

The nature of the protection problem changes with the type of infrastructure system involved. Some, like telecommunications and finance, are almost "pure" information systems where the computing and digital information content is high and where the information subsystem and the global information infrastructure are at the heart of the enterprise. Others, like gas and oil storage and transmission, transportation, and water supply, are organized around the physical handling of material, and the information systems that control their operations are more deeply embedded in the organization. In such infrastructure systems the information and control subsystem may not be as visible to top management as activities regarded as more central to the organization. Electric power generation and distribution is an intermediate case. Generation involves the physical handling of fuel, but the good that is distributed more closely resembles information. Emergency service systems pose different challenges. They are characterized by their highly decentralized nature, down to counties, towns, and precincts where sophisticated technical management is less likely to be available, where investment resources are constrained, and where, from a national perspective, no central management structure exists.

Before addressing the subject, it is useful to have a conceptual framework. How does one go about "protecting" anything from damage or loss? Figure 1 suggests that there are five coupled issues. First, one naturally attempts to *prevent* damage or loss. But since absolute perfection cannot reasonably be expected, one must recognize that defenses will at some point be breached and undesired outcomes will be realized, the familiar strategic nuclear "When deterrence fails …" thought process. There are two courses open: to *limit damage* and to *reconstitute* the pre-attack state of affairs. There may be trade-offs between these two courses, as there is between how much one invests in air bags and how much one invests in insurance. Finally, damage limitation can be viewed two ways. One can take a short-range view and seek to limit the damage in a single attack, or one can take a long-range view and seek to minimize losses over a period of time. Again, there are trade-offs, between detailed and potentially costly scrutiny of individual transactions and aggressive prosecution and punishment of attackers.

Figure 1. Strategies for Protecting Infrastructure Systems

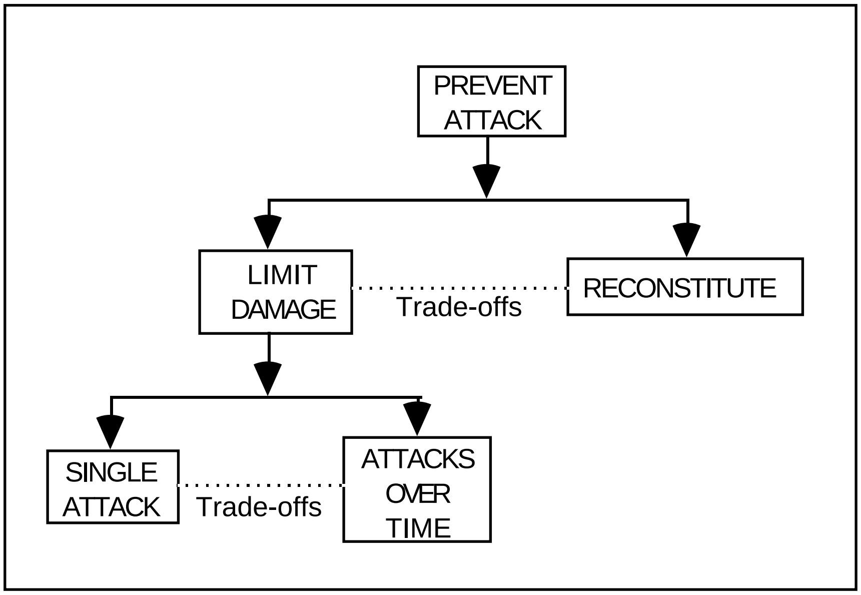

Infrastructure systems are mixes of public and private ownership, and each type of owner will have different approaches to investment in protection. Private owners, faced with loss of revenue and loss of confidence by their customers, regulators, investors, and insurers, will seek to restore revenue and customer confidence, satisfy regulators, document losses, and avoid liability. Governments will focus on protecting national security, preventing future attacks, and identifying and punishing attackers. Industry and trade groups are likely to act in some intermediate way, taking the long view for the benefit of all whom they represent, but reflecting very directly the individual financial and competitive concerns of their members.

Finally, it is useful when thinking about attacks to appreciate the possible motives of the attacker and the scale of attack anticipated. Figure 2 suggests both a continuum of scale of attack as well as a continuum of motivations of the attacker. This could also represent a training sequence for attackers. At the lowest level of violence are the innocent hackers, students in high school and college, who are doing it for the thrill and intellectual challenge. Their efforts to date have helped to educate us in the technical possibilities while still at a modest level of destructiveness. But they can also be viewed as at the "entry level" positions in the "attacker profession." Innocent hackers can develop into malicious hackers with a personal or social agenda.

Figure 2. Continuum of Motivation of Attackers and the Scale of the Attack

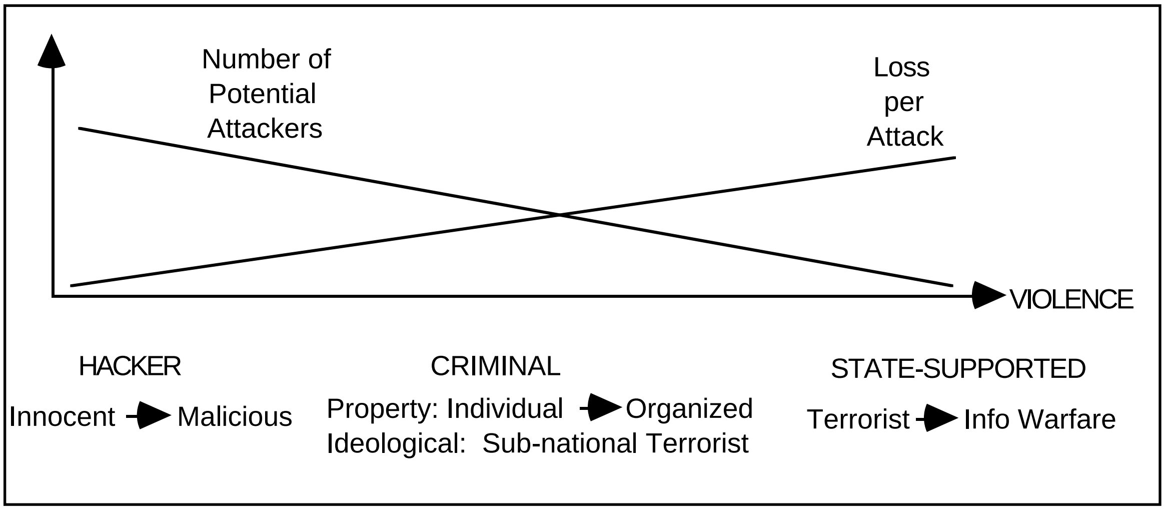
Number of Potential Attackers
Loss per Attack
VIOLENCE
HACKER
Innocent → Malicious
CRIMINAL
Property: Individual → Organized
Ideological: Sub-national Terrorist
STATE-SUPPORTED
Terrorist → Info Warfare

Criminals, attacking for personal gain, are the next level of seriousness, and they, in turn, provide the skilled labor utilized by organized criminal groups. An ideological branch is represented by sub-national terrorists, whose interest is their social rather than personal agenda. The highest level of concern in the present context is the state-supported terrorist and the state-on-state attack. One can expect that the number of potential attackers is large at the left side of Figure 2 and decreases as one moves to the right. Private responsibility is dominant at the left side of this spectrum of threats while government responsibility is dominant at the right side. They share responsibilities in the middle of the threat spectrum.

In the analysis that follows, military and national security metaphors are used extensively. There are two reasons for this. First, the subject *is* one of offense and defense, action, counteraction, and counter-counteraction, strategy and tactics. And second, many of the concepts, their nuances, their historical development, and the dominance of technology in warfare are familiar and thus provide a common set of concepts. At the same time, one must caution that these are metaphors and their application to the subjects discussed here will require reinterpretation in the light of new technology and new domains of application.

Preventing Attack

There are several approaches that can be taken to prevent attack:

- Deter Attacker
- Ban Attack
- Take Preemptive Action Against Potential Attackers
- Detect and Block Attack
- Erect Barriers Around Systems
- Inform Operators of Protection Possibilities
- Harden Systems

These are illustrated in Figure 3. For the first four, the actions taken are external to the infrastructure system and are directed against known or potential attackers. For the last three the actions are internal to the system, examples of the classical fortification approach. The former are "active" approaches while the latter are more "passive" in nature. All are candidates for implementation, depending on the organizational and physical characteristics of the infrastructure system being defended.

Figure 3. Preventing Attack—Possible Approaches

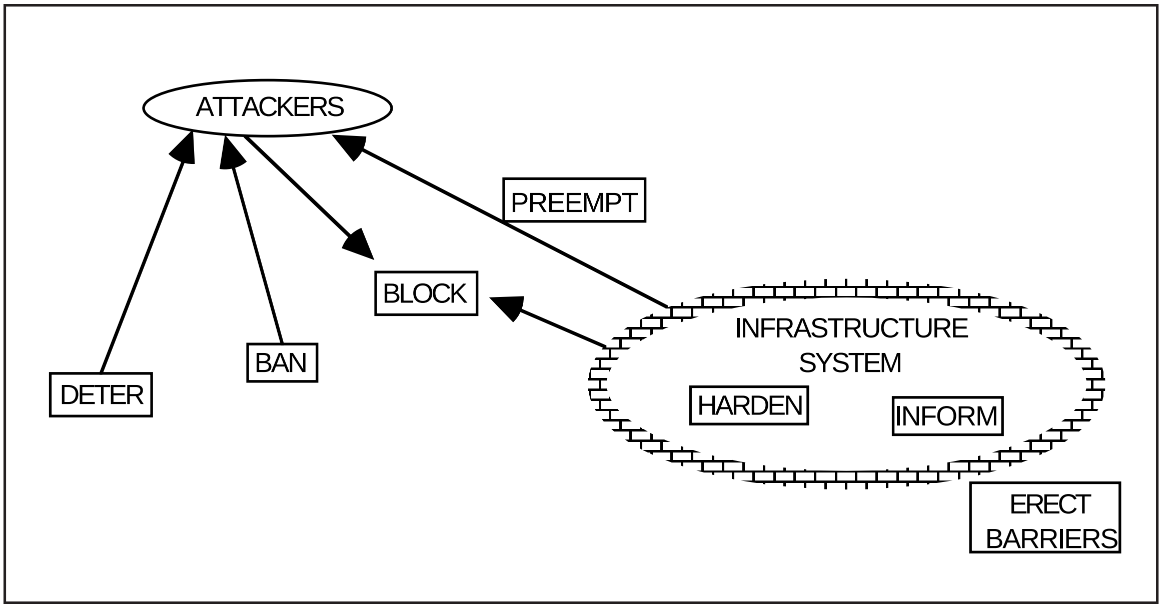

Deter Attacker

Deterrence, in principle, provides an approach to preventing attack. One relies on establishing a credible capability to identify the attackers and punish them for their acts. At low levels of violence, however, deterrence is less useful, because the attacks are likely to be too numerous to prosecute, especially when attackers can hide behind a multiplicity of political and bureaucratic jurisdictions. Even at higher levels of attack, by criminal and sub-national terrorist organizations, the utility of relying on deterrence is questionable. The threat of punishment does not deter criminals today, and ideological terrorists welcome the attention their acts draw regardless of their possible fate.

Thus only at the state level of violence might deterrence be a feasible strategy, although several questions must first be addressed. Can the perpetrator of an infrastructure attack be identified sufficiently unambiguously that an international response can be justified? The ability to mount attacks on infrastructure information systems from many locations means that responsibility for supporting the attack may be widely shared. Tit-for-tat responses to infrastructure attacks require retribution of like kind directed to civilian populations. While economic sanctions are applied today, they are of mixed effectiveness and sustained international support is difficult to achieve. Targeting civilian populations may be neither politically effective nor desirable.

A theory of deterrence of infrastructure attacks, absent concurrent conventional military attacks, is needed. Infrastructure attacks may have some of the character of the Cold War: continual probes small and large, with occasional accidents and losses, played out on a global stage and related to other global issues such as terrorism, criminal conspiracies, economic competition, and international trade.

Ban Attacks

The use of international arms control agreements to limit actions mutually undesirable to their signatories has achieved considerable success. Such agreements have, for example, reduced the threat of nuclear war by limiting the testing of nuclear weapons and weapon systems, the numbers of nuclear delivery vehicles, and the deployment of nuclear weapons in various locations. Such agreements require for their success the ability of signatories to monitor compliance with the agreement and to consult when actions that appear incompatible with the agreement are documented. In other cases, such as laser and particle beam missile defense, such agreements have limited the engineering development of the technology to some extent.

It is questionable whether an arms control approach to lessening the threat of cyberattacks on civilian infrastructure systems is even feasible. The basic “technology” has already been developed and successfully reduced to practice. The barrier to entry is low; techniques for identifying “illegal” transactions in information systems must be developed; and detecting, analyzing, and documenting such transactions could create severe problems for the management of infrastructure systems.

Practical concerns notwithstanding, international agreements that would outlaw infrastructure attacks through communication networks have some merit and should not be dismissed from the set of possible responses to the threat. There are two reasons for this. Internationally agreed-upon statements of principle have an important moral force. Furthermore, such agreements could add pressure from the world community in support of domestic legislation and enforcement, much as agreements under the International Telecommunications Union (ITU) constrain domestic communication regulation, define compliance requirements, and provide mechanisms through world, regional, and special radio conferences to highlight areas of disagreement and to discuss remedial measures.

This suggests that the International Telecommunications Union, and other international organizations, can assist in exploring this avenue for reducing the threat of infrastructure attacks.

Take Action against Potential Attackers

One can seek to identify the attacker or attacker group. This is impractical to do for the large population of innocent hackers, and given their location worldwide made possible by computer and communication networks, will generally be beyond the resources and jurisdictions of both private and public entities. But as attacks become increasingly violent, the practicality of government action increases.

Just as threats against the president are recorded and assessed, so also might malicious threats against infrastructure be reported by system operators and reviewed for seriousness by appropriate federal agencies. These might include federal system operators if they are the implied target, state or federal regulatory agencies, and federal law enforcement agencies. The point is to collect specific system threats at a central point so that patterns can be sought and clues provided as to the time and location of attacks. As such early-warning information is collected and assessed, response teams consisting of system operators, information security specialists, and law enforcement officials can be assembled on a case-by-case basis. Should the level of malicious hacker threat increase in number, seriousness, or effectiveness, ad hoc response teams can be replaced by standing teams managed either by industry, regulatory agencies, or law enforcement agencies or by some combination.

Since hacker threats will presumably be worldwide, both in terms of perpetrators and targets, an important role of government is to create an international environment that encourages cooperation. This can be done through existing international organizations and agreements or, if those fail to provide adequate jurisdiction, new agreements for sharing threat information and providing for coordinated pursuit of suspected attackers can be established. International cooperation could be encouraged through the provision of reciprocity.

As one moves up the scale of violence, the identification of criminal and terrorist acts against critical infrastructure fits readily into currently accepted models of public and private roles. Domestically, federal jurisdiction over threats to interstate systems provides a basis for centralizing threat information, and international law provides a basis for cooperation to anticipate criminal and terrorist acts. Nonetheless, it would be useful to review existing laws, treaties, and protocols to establish how coordinated action against specific, potentially identifiable threats can be enhanced.

At the highest threat level, the cooperation of the attacker nation will not, of course, be forthcoming, but defensive alliances are possible. Anticipatory actions are the concern of the federal level and most naturally the domain of the national security agencies. But the private ownership of the targets, the understanding of their vulnerabilities, and the desirability of enlisting the R&D community to react to technically leading-edge threats complicate organizing a government response. Early action to address this jurisdictionally complex question is called for. In effect, do we need a "CINC" for information infrastructure defense, and where in the U.S. government should such a point of responsibility be located?

Apart from the organizational question, there is need for R&D on information infrastructure defense. Some of this development will be done by private industry, either infrastructure system owners and operators or by the computer, communication, and software industry, and their solutions will be put into deployed systems through the normal processes of technology transfer. Nevertheless, there is an important role for government-funded R&D at the system level, where individual companies are less able to carry out system simulation and testing. These could be jointly funded through industry-wide consortia with both public and private resources.

Detect and Block Attack

Identifying attackers will be more effective the more “sensors” one can deploy for that purpose. Sensors can be of two sorts, the eyeballs and brains of people and automated computer misuse detection systems. At lower levels of attack the attackers may have interacted with the system, as a customer or employee. And operators, if properly trained and tested, can serve usefully. Operators of a system understand what is “normal” and what is out of the ordinary. However, they need to be encouraged to think defensively and rewarded for alerting and thwarting attacks. Planning for attacks on a system and designing in safeguards has been demonstrated for national security systems such as those involved in the production, deployment, and maintenance of nuclear weapons. The protection of civilian infrastructure systems will require similar dedication to detail. It would be useful to collect the experiences of complex military and other government systems to extract from them lessons that are applicable to critical infrastructure systems.

Such consciousness-raising is needed at management and regulatory levels as well as at operational levels, for these are where security requirements are established and validated, and where resources to provide enhanced protection are approved. The threat to infrastructure systems must be more widely appreciated. This could be enhanced through government-chartered simulation and carefully controlled mock attacks. Turning real systems, upon which society depends, into experimental test beds is potentially dangerous, but responsible teams of testers can usefully be employed to collect system vulnerability information and to assess potential approaches to their protection. There may, for example, be “attack signatures” that can be recognized and responded to in sufficiently near real time to identify the target of the attack, the location or identity of the attacker, and to activate defenses. There is, unfortunately, a kind of “uncertainty principle” operating. The more one works to convince the defenders of the need for defense, the more one highlights the vulnerability of systems to attackers. And the more detail one circulates about countermeasures, the more one compromises what should be protected. These concerns notwithstanding, the protection of critical infrastructure must be transformed from a matter of study to one of practical system engineering and testing.

There are a number of areas where national security and law enforcement agencies can apply their expertise. Profiles of attackers, ranging from the malicious attacker and extending to potential state-supported groups, can be constructed. State-supported attack planning may leave a trail that could be followed by domestic and foreign intelligence agencies working together. Attacks intended to disrupt an adversary’s infrastructure system will require recruiting, training, exercising, and managing skilled people, and these are likely to provide further indicators and signatures. Attack planning will be required, and this could involve probes and rehearsals, information on which can be sensed and collated.

Erect Barriers around Systems

The most obvious approach to preventing attack is to raise the threshold for successful attack by building “walls.” This term is intended to cover a variety of technical measures: passwords, with varying degrees of password discipline; firewalls; call-back procedures;

link, session, and file encryption; etc. The design and implementation of walls is the most natural technical approach to the problem and one that most naturally appeals to a technological society. But this approach shares the dilemma of military electronic warfare: sophisticated attacks generate more sophisticated countermeasures, which then are met with even more sophisticated attacks. The same information technology supports both, and access to that information technology is widely available.

Nevertheless, the engineering of walls is something that we know how to do, and it does not make sense to fail to respond technologically to technological threats. It is the responsibility of a system operator to incorporate protection, though at some point it becomes more effective to take community action rather than individual action. Thus, national air defense makes more sense than every facility hardening itself. Where this crossover point is for information attacks remains to be established.

There is a role for combined public-private action to define standards for security so that users can be aware of what protection is offered and where it fails. Regulators of infrastructure systems have a primary responsibility to encourage their industries to take cooperative action through existing administrative procedures. Absent public and private initiatives to provide protection, customers will eventually seek legal protection and compensation for losses. Thus private system owners and operators need to exercise leadership lest they find themselves the recipients of engineering solutions imposed by regulators, legislators, or the courts.

The development of protective technology is something that public agencies and private parties can undertake. Public agencies have a role to develop protection for their own systems as well as to stimulate the development and transfer of such advanced technology to the private sector. The software, computer, and communication industries have their own incentives to develop protective measures, both as products and to embed in their own systems to enhance their utility. They also have a need to protect their internal production processes to assure the “purity” of their products.

Inform Operators of Protection Possibilities

In parallel with increasing public appreciation of threats to enhance the detection of attackers is increasing public awareness of the type and cost of system protection options. Protection will cost money and this must be paid by users, either as a personal expenditure, e.g., encryption software for their personal computer; or as a ratepayer; or as a business expense akin to insurance, which the manager and investor must recognize as a prudent expenditure.

Increasing safety regulation for the protection of infrastructure systems will inevitably increase public awareness of the costs and benefits of system protection, and this should be encouraged. To be sure, regulatory proceedings and the debate over the length of exportable encryption keys are not the stuff of popular culture. But the technical sophistication of people increases with time. The number of people who have been forced to become knowledgeable about on-line services, modems, the Internet, operating systems, and application software is astonishingly greater than one would have guessed in 1970 or even in 1980.

There are two ways to help people appreciate the magnitude of electronic and cyber-threats. One learns by being burned, and inevitably much public appreciation will come the hard way. The other way is to learn through information and warnings. Regulatory proceedings will provide some part of the public education role. While the details are arcane, they nevertheless can seize the public consciousness, as we have seen recently over the matter of V-chips as protective technology against socially undesirable entertainment content and filters to block Web sites.

An important role for government in this regard is to help increase awareness of the losses due to infrastructure attacks through the collection and publication of statistics. Government reporting of airline schedule performance has made this a useful statistic in planning flights. The current debate over airline safety metrics suggests how potent a tool publishing the facts can be. Managers search for the most sensible metrics, operators are aware of them in their daily decisions, and customers vote with their feet, which, in turn, impacts profitability and market performance. And where regulators are timid, legislation can help. This will, of course, run counter to the current concerns with big government, over-regulation, and excessive requirements for government reporting. But legislation will be useful in resolving what are basically political issues.

Harden Systems

Related to the building of walls, but a separate system design issue, is the "hardening" of systems against attack. Hardening systems against electromagnetic attack is a well-developed technology, originally intended to prevent damage from the pulse of electromagnetic radiation accompanying a nuclear detonation. Similar engineering approaches are used to protect space systems against damage from cosmic radiation, and consumer systems against electromagnetic interference. Similarly in automobile design, the structure is hardened to increase the survivability of occupants for specified levels of impact.

Such approaches can both limit damage in the event of attack and prevent attack by discouraging the attacker or by raising the ante for a successful attack. There are two ways of viewing the development of such approaches. The private role is to develop and incorporate hardening into infrastructure systems. The cost of this can be recovered through ratepayers in the case of regulated industries, through reduced product and service warranty and product support costs in market-driven industries, or through lower liability insurance premiums. The choice between regulatory and market mechanisms to fund infrastructure protection is a public policy issue to be settled through political processes. The answer is likely to be different for different types of infrastructure systems.

Limiting Damage

As with preventing attack, there are a number of ways of limiting the damage sustained in an attack:

- Audit System Operation
- Analyze Attacks
- Identify Attackers
- Warn of Attacks
- Counterattack
- Reallocate System Load

- Provide Subsystem Redundancy

Figure 4 shows internal and external system options for limiting damage. A new aspect here is an external “defense center” where damage control measures are coordinated. Both strategic as well as near real-time tactical responses are possible.

Figure 4. Limiting Damage from an Attack—Coordination of the Defense

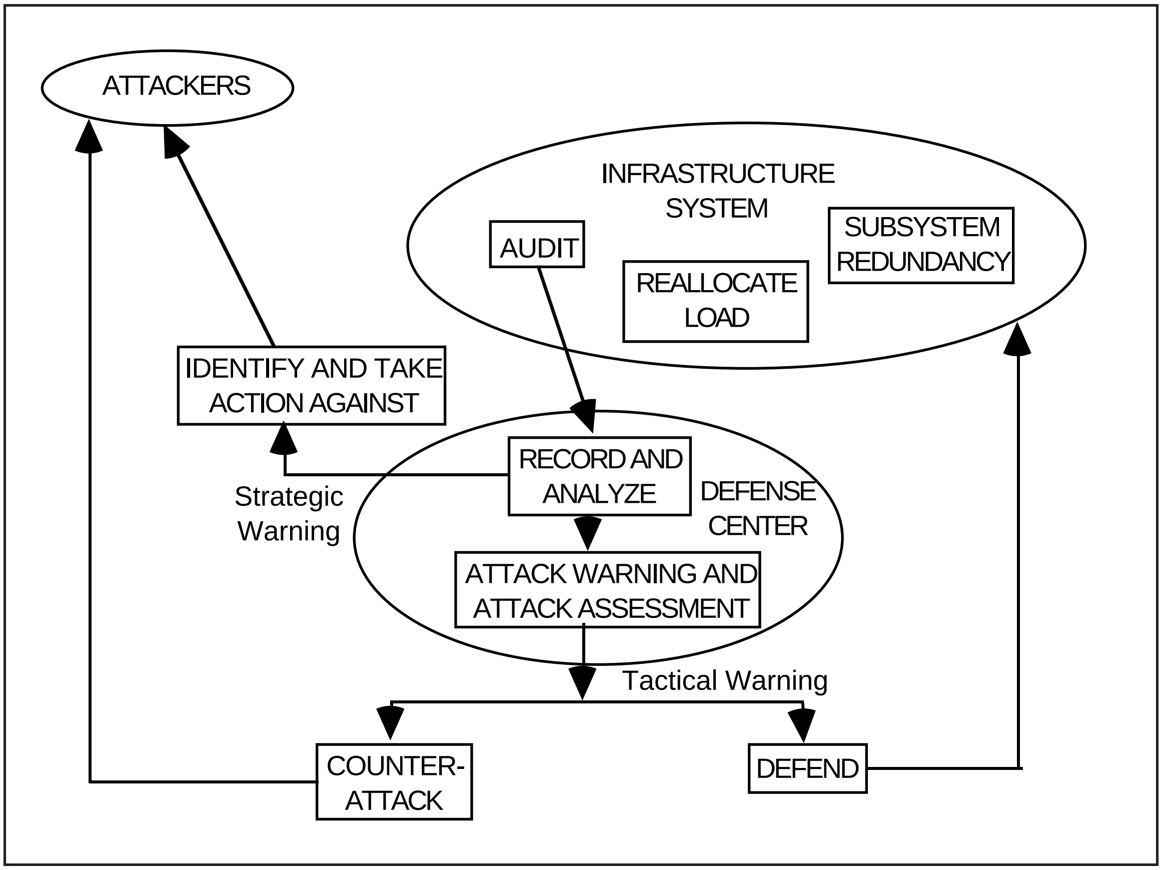

Audit System Operation

Taking as a premise that despite pre-attack intelligence, early warning systems, and a defensive perimeter, the infrastructure control system will be penetrated, the next line of defense is to limit the loss resulting from the penetration. One question is, losses to *whom*? It is possible that national leaderships, infrastructure owners, and clients served will have different perspectives on this. An example is where small losses are allowed to occur to enable the defender to prevent larger losses later. Limiting damage thus has at least two different thrusts as noted in Figure 1, a short-term view where the object is to limit the damage from an individual attack, or a long-term view where smaller losses are sustained if larger losses can be avoided later.

Limiting damage also involves considering the scale of the attack. At the hacker end, the losses may be annoying, but of a magnitude, even over some period of time, that could be less than the cost of system countermeasures, limitations on service, and possible reduction of market size. Losses due to system penetration for criminal purposes may also fit into auditing and law enforcement paradigms, more a cost of doing business than a major factor driving the restructuring of enterprise systems.

There are, at this lower end of the scale of violence, a number of measures that can be undertaken to control, though not eliminate, the threat. A first line of defense is provided by the audit process. The same information technology that allows sophisticated attacks on information systems can be used to implement countermeasures, much as increasing computer capacity makes feasible both more sophisticated cryptographic algorithms as well as tools for breaking into encrypted information.

In the case of information systems that track financial and other similar transactions, normal audit techniques can be applied, as they are today: multiple copies of records, exhaustive cross-checking, encryption of contents to make modification more difficult, passwords and other forms of access control, audits of entries, use of behavior styles and other types of biometric approaches to the verification of users, expert systems for flagging unusual patterns in data, etc. The development of such defenses is a viable commercial business and it will grow as losses to penetration not only become more widespread but receive greater publicity. Government roles in this could take the form of testing of commercial products, with disclosure of results, for the purpose of adopting them for the government's information systems and to encourage or approve their use by government contractors. Government can also support the development of such defensive systems but allowing their full commercial exploitation by the developers without interposing lengthy approval processes between developer and market.

Audit procedures are critical not only for tracking transactions but also for monitoring the software resident in or otherwise controlling information systems. The software content of any computer changes continually as new programs or program updates are loaded, as programs are downloaded through network connections, and as various "fixes" are introduced by system operators, maintenance personnel, and users in response to discovered flaws or changed needs. Any of the fixes can themselves have flaws, from logic errors, viruses, and "time bombs," all ranging from trivial to fatal. Tools to better understand software resident in a machine and on a network are needed. Mandating the use of such tools and establishing the standards for assuring system integrity is an area where government leadership can be exercised using government owned or operated systems as both test beds and as early markets.

While NSA has superb capabilities for developing and validating encryption systems, the lack of public confidence in government security agencies is substantial. There is a conflict of interest between a government that wants its own information systems to be as secure as possible, yet wants to be able to penetrate others' systems for national security and law enforcement purposes. It would seem best to separate the two communities, and let each work separately to achieve its desired ends. What may be appropriate, however, would be for the government to cease attempting to regulate private cryptography and to adopt a more hands-off posture.

For the protection of information systems that support the transportation of goods, the same audit techniques as used for financial data are applicable, but system protection also involves a greater measure of physical access control. Nevertheless, the same technology that can be applied to digital information can be used for more sophisticated lock-and-key and area surveillance systems. What is different about these systems is that, dealing with "analog" materials, they allow additional analog metrics. Thus location, mass, velocity, temperature, and the like provide further system control parameters and status indicators.

It is at the state-supported levels of violence, where high levels of loss are the adversary's objective, that high cost protection is warranted. Some of these high-value targets will be government controlled, such as military systems and the preservation of high-value, high-risk items such as weapons of mass destruction and their constituent materials, gold, and the like. Civil systems such as the air traffic control system and customs monitoring at ports of entry are similar in character.

Government owned and operated systems have interfaces with comparable systems of other sovereign states. It is here that there is an important government role, since the system engineering of interfaces is, by its nature, a matter of state-to-state negotiation. Even where private entities, such as airlines, undertake discussions with foreign entities, they often take place under government cognizance. Such situations pose extremely complex problems. Between private entities, mutually satisfactory arrangements can be made and enforced. But when the arrangements must fit into a broader context of international relations and where agreements in one can have ripple effects into quite separate areas, system engineering meets practical limits. What is needed are techniques to understand the dynamics of systems of systems, much as we have in the Internet.

Analyze Attacks

Part of the problem in tightening the security of information systems derives from the nature of the underlying technology. Contrast computing and air transport. Aircraft accidents occur quite infrequently, and when one does occur it is a major event that sets in motion detailed investigations to determine the cause and to formulate steps to prevent a recurrence. Computer systems, on the other hand, do strange things many times a day; diagnostic evidence relating to them is often difficult or impossible to collect; the problem is often not reproducible; the logic of computer programs defies thorough understanding, with the result that the potential for unexpected or undesirable behavior of software can lie undiscovered for years; and changes in requirements and technology result in frequent changes to systems, further impeding the analysis of anomalies. While aircraft crashes are treated like cases of murder, information system problems are like colds. With this attitude, opportunities for abuse abound.

Private actions to improve the state of the art in software and to build a software engineering discipline are proceeding through the efforts of software industry researchers and professional educators. Government action in the past has been most helpful in catalyzing these actions. In the 1980s, the DoD supported the development of software productivity and quality metrics and imposed software maturity level certification on its contractors, aided by the Software Engineering Institute it established and funded. These metrics and standards have been adopted by other government agencies and, more recently, by private industry. The success of this approach derives from its addressing a root cause of software quality problems, the *process* by which software is created. It is possible that this successful strategy could be applied to the problem at hand through the creation of a comparable System Security Institute. International quality standards as embodied is ISO 9000 specifications have also provided an impetus for improvements in software quality and reliability, and could similarly be extended to address system security.

Government leadership, especially in encouraging private industry to establish security standards for infrastructure systems, would be valuable. Such standards should be voluntary, but their existence would serve the important purpose of allowing wide adoption without requiring each system operator to research and develop its own security standards. In particular, security breaches in infrastructure systems should become matters for some form of central analysis so that quantitative vulnerability and threat data can be collected and used to improve the national level of system security. What can be seen as private problems need to have more resources applied to their analysis and correction than private owners may be able to provide. Such procedures should protect the privacy and property of innocents as well as the integrity of legal evidence and of national security information.

Identify Attackers

Audit processes, if properly constructed, can not only reveal the presence of "anomalies," but can provide clues to or a trail back to the source. This is, however, an area that poses significant practical problems. Anomalies that may indicate the presence of unauthorized activity in an information system are rarely unambiguous with respect to either the identity of the person responsible or his intent.

In pursuing peculiar system behavior, there are extraordinary demands placed on those involved. Investigations into intricate technical penetrations of information systems require skills in system programming and communication management. Such people tend not to be thoroughly acquainted with the legal ramifications of privacy and liability, nor do they have the perspective of law enforcers for whom conditions for search and seizure and chain of custody of evidence are matters of course.

The intent of system security personnel can also be critical in monitoring and controlling intrusion. Private owners can choose between investigating intrusions with the intent only of eliminating them, or they may focus on the prosecution of offenders. Their position will be influenced by the rights of their employees, liability to their customers, legal conditions in contracts, and responsibilities to shareholders. Public entities will operate under different legal and administrative requirements, and they may have fewer degrees of freedom.

All this is to say that "rummaging around" to solve a "technical" problem, exercising system-level privileges that in hindsight can be seen as unreasonably intrusive, as resulting in potential damage to employees through the leakage of information, and as compromising prosecution of offenders, requires management attention. These considerations notwithstanding, audit systems should facilitate locating the sources of intrusion. Early action can result in halting intrusion before damage is sustained, or it may serve as an early warning alert for the deployment of more comprehensive monitoring to secure legally valid evidence or to assist in understanding more thoroughly the attacker organization.

The choice of approach to "active auditing" will depend on the scale and target of the attack. Hacker attacks on less critical systems are less likely to justify the expenditure of countermeasure resources that state-supported attacks on critical systems will. This suggests that the need for a formal "attack assessment" function be recognized. This is a national-level need where government leadership can be effective. The suggestion is not for continuous detailed government monitoring of all infrastructure information systems. Instead, what is envisaged is a federally funded educational program regarding threats, sources of expertise, a national-level center for reporting suspicious events, databases of comparable events, a directory of available technical expertise, and assistance in contacting appropriate government offices.

Warn of Attack

A benefit of such central reporting and analysis of real or suspected attacks on infrastructure information systems is that timely warning of other parts of the affected infrastructure becomes possible. The analogy is to an air defense control center. Defensive measures that would ordinarily be too costly to implement regularly can be temporarily adopted; and recording of or longer retention of system information relating to anomalies and operating parameters can be instituted to assist in identifying and locating attackers or in retaining and protecting legal evidence.

Such warnings pose the problem of how and to whom they should be distributed. Too wide a distribution increases the likelihood that the warnings will become known to attackers, thus enabling them to avoid interception. Too narrow a distribution and they will not reach operators who need them. The issue is best dealt with by working with the large corporations that manage infrastructure, with professional and industry organizations, and with government regulators to establish procedures appropriate to individual circumstances. The international distribution of attack warnings requires consideration in view of the varying degrees to which cooperation occurs between sovereign nations and various legal jurisdictions involved.

Counterattack

Attack warning and attack assessment lead to the idea of active defense by near real-time reaction. The premise is that upon sensing an attack the system operator has options for restraining or disabling the attacker. Private entities are unlikely to risk the liability for responding to false alarms and possibly alienating customers, but such an approach may make sense at the national level in the case of attacks launched by adversary states directed at security systems operated by the federal government. Here again is an area where government initiatives to examine the protection of their own systems can create capabilities that could be modified for possible transfer to other organizations.

Current paradigms in information systems are inconsistent. On the one hand, access to information and wide connectivity is good, the more the better. Information and resources are to be shared, yet privacy and intellectual property rights are to be protected. As information systems evolve from free and public systems into those structured to enable interaction between identifiable entities in “known” locations, e.g. for billing purposes, the information system analog of counter-battery fire becomes conceivable.

Reallocate System Load

For infrastructure systems where the information-dependent parts of the system are supporting rather than primary, the analyses of failure modes can tend to emphasize the primary system functions and place inadequate emphasis on the supporting information subsystems. Thus a highly reliable control system, with built-in redundancy, can easily be taken for granted and relied on to detect, diagnose, and manage service interruptions. But as information technology is used to achieve enhanced operating efficiency and as utility systems are geographically extended as the result of growth of the economy, industry consolidation, deregulation, and the separation of production and distribution functions, there are increasing opportunities for dynamic load management. There are, however, also opportunities to over-centralize system control, thereby concentrating system vulnerabilities in a few locations that are increasingly dependent on links to remote sensors and to adjacent control centers that are potentially penetrable.

One response to cyber-attack can be to shed priority load to those parts of the system operating normally, assuming that the primary infrastructure is internetted; that their respective control systems are internetted; and that the control system has adequate information to dynamically redistribute load. As geographical monopolies are reduced and as deregulation encourages local competition, dynamic load shedding becomes increasingly feasible. Factors operating to discourage the internetting of utilities' control systems are competitive pressures and the desire to avoid the costs of designing, installing, and maintaining control system interconnections. Nor are regional plans for response to cyber-attacks by dynamic redistribution of priority load likely to be high on system operators' agendas. But the technology that produces these system vulnerabilities has within it the power to reduce the problems. It is important that full advantage be taken of the positive aspects of information technology. Priority load shedding will, of course, be useful in maintaining infrastructure operation through a variety of types of system failure and for response to natural disasters as well.

Government action can encourage dynamic load shedding as a protection of essential infrastructure services through public awareness campaigns that can stimulate customer interest; through joint actions coordinated through existing federal-state regulatory channels; through recognition of necessary expenditures for favorable tax treatment or for incorporation into tariffs; and through sympathetic interpretation of existing anti-trust statutes if appropriate.

Provide Subsystem Redundancy

Finally, consider the role of redundancy in limiting loss during attack. Information attacks are likely to be directed at the heart of a control system. By design one can imagine, for example, partitioning a system in a way that keeps parts of the system at each level or in each region insulated from the others through internal walls. Within each partition, sensing of attack, either manually or automatically, can provide an opportunity to switch to a redundant element. Since redundancy is costly, the most effective approach is to make redundant elements addressable from all parts of the system. The feasibility of providing redundancy will depend to a great extent on the details of the infrastructure system in question as well as on the implementation of the control subsystem. Most important is the issue of legacy software. It is quite possible that system software will have its roots so far in the past as to render such approaches infeasible or unreasonably costly.

Possible government actions could be to consider incorporating such redundancy-based system architectures into future procurements. To be sure, governments suffer at least as much as do private entities from the limitations imposed by legacy designs. Nevertheless, it is not unreasonable to expect that new DoD system procurements will offer opportunities for the incorporation of novel approaches to system protection in the years ahead. Thus new approaches to infrastructure system design may be feasible.

There is a substantial built-in delay in this proposal, unfortunately. There will be a reduced number of DoD procurements of new systems in the years ahead. There will be a need to minimize risk by sticking with proven designs, and the long period between the system concept stage and operational deployment makes the use of military infrastructure systems as test beds for new protection technologies difficult. Nevertheless, it should be considered as part of a government-wide procurement strategy.

Reconstitution

Reconstituting both the infrastructure system and the users it supports can take several forms:

- Backup System Status
- Assess Damage
- Insure against Loss
- Ration Residual Resources
- Stockpile Long-lead Items
- Provide Standby Capacity

Figure 5 suggests that the central reconstitution issues, absent physical damage to the infrastructure system requiring repair, are related to the responses of the users.

Figure 5. Reconstitution after an Attack

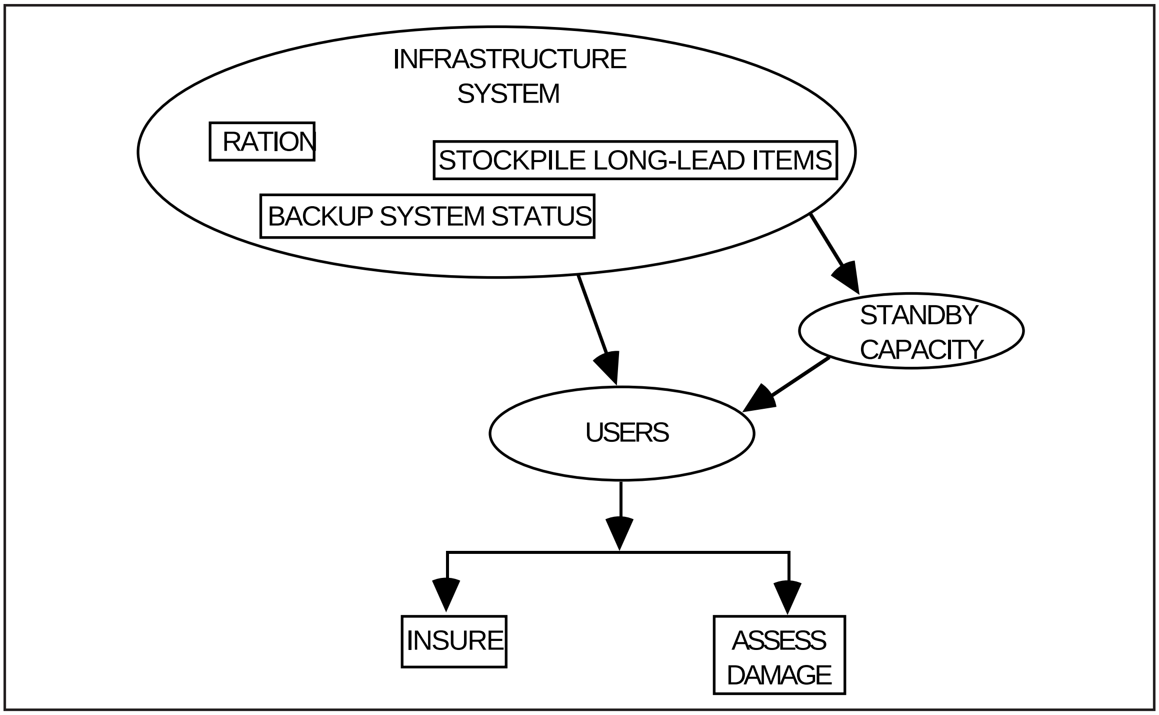

Backup System Status

In addition to limiting damage from an information system penetration is the need to plan for the reconstitution of the system and the restoration of service. As Figure 1 suggests, there is a trade-off to be made between investing resources into limiting damage and investing to make reconstitution of the pre-attack state swifter and more complete.

The most direct approach to reconstitution is to backup system state information from which one can restore the system to its pre-attack condition. In the case of information-intensive infrastructure systems, such as financial services, this can be effective though still costly, considering the time and effort to reconstruct transactions that occurred during the shutdown period and the longer term loss of confidence in the system by its users.

The capability to restore information systems is well developed in light of the current state of computer and communication system reliability. It is also addressed in disaster recovery plans that every prudent organization, public and private, prepares. While these more often address physical security issues such as fire, flood, earthquake, civil disturbance, and the like, they have as their premise loss of service and the need for its restoration.

At low levels of violence, such plans provide a useful starting point, although having been prepared in the past they may require updating to address current and future threats that have been exacerbated by the connection of computers to public and private information networks. For example, disaster recovery plans will involve shifting work to other facilities able to assume the additional load. But at higher levels of information attack, where the loss of service could extend to large numbers of facilities, the ability of the residual facilities to absorb load will saturate at some point. Thus straightforward reconstitution of service has natural limits.

The more difficult problem is the restoration of service in infrastructure systems that control systems of physical objects. Not only can bringing down such systems involve irreversible actions, such as physical damage and loss of life, but even for less severe outcomes the reconstitution will have physical requirements and constraints to contend with. Engines and pumps must be inspected and restarted, maintenance and operating personnel must be deployed, etc. Putting a railroad system back into operation is quite different from rebooting a computer and reloading files.

The responsibility for restoring government-owned systems is clear. Government roles in restoring private systems are less clear. At low levels of attack, one can take the position that private system owners bear the responsibility driven by their contractual responsibilities to their customers and the potential loss of customers to their competitors.

But as the scale of the attack increases, government is likely to play a larger role. A parallel is the role played by the Federal Emergency Management Agency in coordinating responses in regional disasters, where federal aid programs may apply, and where the ability to organize region-wide information and deploy military manpower may be crucial.

Governments may reasonably inquire into disaster planning in privately owned and operated systems that have relevance to national security and thus can usefully “seed” interest in the issue. Similarly, regulators are concerned with the degree to which regulated industries fulfill the public interest they were chartered to serve. But with the shift in regulatory policy toward deregulation, the degree to which regulators have the administrative and statutory tools to do this has been reduced.

Government at the national level can usefully explore the issues of threats, vulnerabilities, and responses through support of research, development, studies, and analyses. As part of such initiatives, governments can undertake simulations of information system attacks and support games and exercises to assist in bringing public and private parties into better contact. Such efforts would have benefits on both sides. Public agencies can better understand real-world constraints and limitations, and private parties can get insight into the public interest issues from a national viewpoint. Reconstitution exercises are likely to suggest where industry standards are needed, where gaps in essential knowledge exist, how public agencies and private organizations can work together, and where legislation may be required.

Assess Damage

Following an attack on information infrastructures, it will be important to assess the damage and to understand the entry mechanisms and the details of the failure modes. There are several reasons for this. Early on it will be important for coordinating relief and reconstitution efforts. Later there will be lessons to be learned, as regulators seek to understand system failure, as insured and insurers reach settlements, and as operators work to prevent recurrences.

Assessment will be necessary because the first response can be to attribute the failure to normal system unreliability. Even if some malicious intent is indicated, there will be a need to establish if deployment of countermeasures is in order; and, human nature being what it is, efforts will be made to cash in on the disaster.

Society suffers regularly from disasters and has developed procedures for handling them. Private organizations seek to return to normal operation as soon as possible, claims to insurers are prepared and substantiated, customers are assuaged, victims helped, and employees cared for. Public agencies including federal, state, and local emergency response organizations, military and national guard organizations, and public relief organizations deploy their assets. Hurricanes, earthquakes, and floods have served to hone these response skills. What will be new in a high intensity information attack will be its scale.

To this end, coordinated federal-state-local emergency planning should be reexamined and updated. The last time the nation addressed large-scale disaster was in the civil defense efforts in the early days of the Cold War. What is different now is that the scale of automated systems and the richness of communication facilities suggest that entirely new defensive approaches may be possible. The information systems whose potential vulnerabilities concern us here also offer important capabilities for responding. If one believes total system collapse will not occur, then the planning problem is to determine how to use the surviving infrastructure to enable the nation to lift itself by its bootstraps to achieve the most effective recovery possible.

Insure against Loss

Insurance has been mentioned already, but always in the context of current insurance industry products and procedures. Can one go further? The answer depends on the scale of attack. Losses from low level information attacks are, or can be, covered now. At higher levels of information attack, one could imagine insuring against loss, although at some point governments naturally enter as an underwriter.

At some level of attack, putting resources into prevention and damage limitation makes more sense than insurance. Nevertheless, governments do, in effect, underwrite catastrophe recovery now, through tax revenues and loans. As frequency of attacks and experience with losses grows, the allocation of recovery funding between public and private sources will become an increasingly important public policy issue.

Ration Residual Resources

When capacity in any system is destroyed, the remaining parts knit themselves together in a variety of ways. Undamaged competitors assume part of the load; others loan equipment and personnel; elasticity in labor markets provides additional capacity, as do the efforts of the affected; private funds and public emergency appropriations contribute to the reconstitution effort.

In the small to medium sized disasters that have been experienced, both pre-planned and ad hoc responses have been successfully employed. Absent a clearer understanding of the threat and expected loss, perhaps that is all that is required at this point. But information technology is increasing society’s capacity to assess damage and to organize and deliver relief. This must be considered as well as the potentially increasing vulnerability of society to information attack. There is a fundamental stability to society that has not been recognized, however. Prudent design of information systems can reduce the risk of attack while at the same time increasing the stability of social systems and the ability of society to respond to threats and crises.

This optimistic speculation notwithstanding, what more can be done? Just as national resources have been centrally managed in major wars, so also could government step in to manage residual capacity until recovery was under way. Mobilization efforts during World War II provide numerous examples of the rationing of limited resources and their allocation to overriding public needs.

Stockpile Long-lead Items

In a cleverly planned information attack, the damage can be substantial and long-lasting. An attack will be planned to make reconstitution lengthy and difficult. One way to do this, beyond the scale of the attack, will be to cause damage to long-lead parts of systems. The response to this should be to stockpile long-lead items. For mechanical systems these are easy to identify, and system operators, both as part of their disaster planning and for regular maintenance management, maintain stocks of critical items. The issue is, are such stocks adequate?

Clearly the answer depends on the number, type, and scale of attack envisaged. Large-scale attack scenario development is important, and this is a role that government can assume since private operators normally do not build such major contingencies into their thinking. It is important, however, that such scenario development, as well as games, exercises, and simulations mentioned earlier, involve system operators. This will not only add realistic constraints to the planning, but it will serve as a conduit for transferring such thinking to corporate organizations.

Provide Standby Capacity

Finally, one can provide standby capacity to replace failed system elements. This is much like redundancy, discussed earlier, but it refers to redundancy at the system level rather than its application at the subsystem level to limit damage.

Some “redundant” capacity always exists, because no system is ever expected to operate at full capacity. Information system attackers will seek to achieve sufficient damage to drive the residual infrastructure beyond its ability to cope. What this means is that the first goal of reconstitution should be to bring the capacity of the residual system to within the limits of safe and prudent operation. This can be done by load shedding at the local level, by regional management of system resources by system operators, and by rationing at the national level.

Additional standby system capacity can be provided by other nations. In some cases this will mean neighboring states, but in other areas it could come from more distant allies. This suggests that a possible government role would be to assist in organizing such alliances or international provider-consortia.

Public and Private Responsibilities

The preceding discussion has called attention to various actions that both public and private organizations can take to protect critical infrastructure from cyber and other electronic attacks. Three major components of a protection strategy have been identified, providing a range of choices. Protection can be expected to put heavy emphasis on prevention of attacks. But in recognition of the fact that prevention cannot be counted on to be completely effective, some investment must be made in damage limitation during attack and reconstitution of the attacked infrastructure system or systems after an attack.

Twenty classes of responses have been presented, and for each of these specific public and private actions have been indicated. The response alternatives are, of course, not exhaustive, but they provide a starting point, a menu, as it were, to flesh out at the next level down the implementation of whatever protection strategy might be selected. The approach selected is likely to differ according to the various types of infrastructure, the assessments of threats, the seriousness of vulnerabilities, and the resources that can be committed.

The complexity of addressing infrastructure protection arises from the intricate mix of public and private assets to be protected and the difficulty of implementing coordinated actions in the private sector due to the multiplicity of managements, shareholders, and customers that interact in a market economy. Further, a prevailing theme of governance in the United States is that corporations should be kept separate for a number of important reasons: the maintenance of price competition; ensuring the widest possible set of consumer options; the protection of companies from predatory pricing; and the protection of workers and consumers from monopolistic power. From these considerations comes a considerable body of antitrust legislation and legal precedents. Thus, the need for service providers to work together to take joint and coordinated action against strategic information attacks runs counter to prevailing orthodoxy. Compounding the problem of analysis and decision making is the fact that these infrastructure systems, providing power, communication, transportation, health and safety services, and capital, interact with each other to support the national economy. Whatever solutions are proposed will have to cope with these factors.

Some infrastructure is federally owned and operated, such as military transportation and communication systems, although strategic mobility and command and control depend heavily on private sector assets as well. Military capabilities, including the National Guard components that are important for disaster relief, also are public entities. The emergency service infrastructure is, for the most part, publicly owned, and, while managerially fragmented, works together remarkably well in times of crisis on the basis of local leadership. Transportation, banking, finance, telecommunication, and power generation and distribution are almost totally privately owned, although all are subject to varying degrees of federal and state regulation.

Thus, an important set of questions is:

Who *decides* on the national infrastructure protection strategy?

Who *selects* among the various response alternatives, industry by industry?

Who *monitors* compliance and effectiveness?

Who *manages* protection implementation programs?

Who *pays* for them?

The answers to the last two questions are the same: the *owners* of the enterprise in question. In the cases of public systems, or components of public systems, the owners are the taxpayers of the jurisdiction involved. Those investments are financed either by taxes or by the ratepayer. In the case of private systems, the shareholders make the required investments with the expectation of future return from customers.

The first three questions are more difficult to answer. When regulation was more heavily relied upon, the regulators would decide, guided by legislated processes reflecting the political will of the nation. Currently we do not depend as heavily on regulatory mechanisms, substituting instead market forces and private decision making. Thus in the face of needs to address system issues, we find ourselves having dismantled much of the regulatory machinery that could otherwise have been employed.

Answering these questions requires merging the different views of public and private entities to make top level strategic decisions in ways that make economic sense, technical sense, and business sense while meeting essential societal needs for reliable and cost-effective infrastructure systems. Each sector brings particular strengths and capabilities to this process.

The suggestions in this paper were crafted to rely on what each participant does best. The federal government has the lead responsibility in national security, and this is central to our concerns here, particularly those attacks that could result in severe damage to the nation's industrial economy. Governments at all levels own infrastructure and have the same needs for security as private owners. In addition, publicly owned systems offer opportunities for development and test activities. On matters where domestic infrastructure interacts with factors external to the country, federal agencies can provide assistance as well as the larger framework within which international agreements are made and enforced. Finally, governments have an important responsibility to collect, analyze, and distribute information on which private entities depend.

Such rationales can, however, lead to intrusive, hyperactive, expensive government programs that may meet neither the needs of the private partners in the enterprise nor those of the nation. Therefore, in the next section, which sets forth some fifty possible public roles extracted from the preceding, they are to be understood as tasks that fall to government by default. Fundamentally, the protection of private enterprise must depend on the choices made by its owners. But this is not to deny that governments provide “bully pulpits” from which to educate, persuade, and lead.

The question of dealing with monitoring compliance and determining the effectiveness of system security measures is perhaps the most difficult, because it brings to a focus the choice between intrusive external regulation and potentially inadequate self-regulation. Given experiences with system penetrations and costly losses and failures, public demands to “do something” are likely to be met eventually. Perhaps pragmatic nations can only learn by experience. Thus in the security strategy that is ultimately adopted, seeking cost-effective approaches to damage limitation in order to minimize the pain of the learning process may be the most practical approach at this time.

Possible Public Initiatives

The following possible public initiatives are intended to be a first step at defining options for action. They are by no means exhaustive and are offered to stimulate further analysis and understanding of the technical and policy issues involved.

Figure 1 presents the main elements of strategy: first *prevent* attack, second *limit* damage from what is not prevented, and third, *reconstitute* both infrastructure and users to as near the pre-attack condition as feasible. These are the cornerstones of the analysis, and the initiatives suggested in the earlier discussion are organized in this way.

The formulation of public programs to implement infrastructure protection initiatives are, then, the next step. What criteria should be used to establish program goals, time schedules, and levels of expenditure? It is suggested here that the guiding principle should be the perceived seriousness of the threat. If the threat is seen to be serious and imminent, a responsible policy approach would be to implement strong measures with great urgency. At the opposite extreme, if the threat is not widely perceived as serious, or as potentially serious but not imminent, then a more cautious and graduated approach is in order. An advantage of the cautious approach is that as understanding of the problem increases, more strenuous measures can be adopted, while starting with major efforts runs the risk of wasting both financial resources and institutional credibility.

It is for this reason, therefore, that the possible initiatives identified earlier are grouped in terms of the ease with which they can be implemented. To illustrate, consider the following nine “rungs” on an escalation ladder of possible public responses.

1. *Undertake policy studies within government.* While chartering and undertaking policy studies are not without disruptive effects, they can be performed through existing agencies, using people currently employed and within existing budgets, and are relatively inexpensive.

2. *Jawbone*. Undertaking to persuade both public agencies and private parties to do something requires going public, and this implies the expenditure of personal and political capital.

3. *Fund contractual studies.* This requires the appropriation of public funds, or the reallocation of already appropriated funds, and consigns the effort to the complexities of government procurement processes.

4. *Develop software to help in protecting systems.* To the problems inherent in spending public funds are added issues of establishing "requirements," and completing software development programs on time, within budget, and delivering planned performance.

5. *Make organizational and policy changes.* Changing organizations and policies is more difficult than implementing actions under existing organizations and policies.

6. *Set standards.* Standards, regardless of their immediate utility, are often seen as inhibiting innovation, an especially serious concern when the nature of the threat being addressed and the most effective approaches to protection are not yet clear.

7. *Mandate change through regulation.* In an era where market forces are seen as more efficient than regulation for achieving public goals, this approach carries a lot of baggage with it.

8. *Undertake operational activities.* Injecting public agencies into the operation of private systems is especially difficult because of the substantially different styles of planning, staffing, and financing of each, their different types of public accountability, and the already complex web of statutory and regulatory requirements on each.

9. *Emergency responses.* Crisis responses, even when coordinated and effective, are by their nature too little, too late. Aside from their cost, in both human and financial terms, they are "non-standard" operations and are intrinsically inefficient.

This, then, is the structure used to organize the possible public initiatives that have been identified.

Preventing Attack

Internal Studies

The federal government should *review* existing laws, treaties, and executive agreements to establish how international cooperation against potentially identifiable threats can be effected. The federal government could then examine the possibility of proposing formal defensive *alliances* to protect infrastructure systems. These could be new organizational initiatives or they could be based on existing organizations such as the ITU or NATO.

Jawbone

The national security and law enforcement arms of the federal government should assist in creating an *international environment* that encourages cooperation across the multiple jurisdictions that are involved in identifying the source of infrastructure threats.

Consciousness-raising concerning system threats and the nature of an attack is needed at infrastructure management and regulatory levels. There is a need to add system protection against malicious and state-supported system attacks to the personal and organizational agendas of managers and regulators.

Regulatory proceedings can also provide a mechanism to educate the public regarding infrastructure security threats and requirements. Available administrative procedures provide opportunities, through notices of inquiry, notices of rule making, and the establishing of industry advisory panels, to make recommendations for further action. Executive branch commissions, task forces, and advisory groups can provide similar opportunities.

Congressional hearings to explore the understanding of these issues and planned responses on the part of regulatory agencies can provide an additional way to formulate policy guidance.

If “walls” are to be built, an *architecture* is required so we can understand where they should be built and how high they should be. Does the nation erect a single fortification like a medieval castle, or does it have a layered structure like ballistic missile defense systems? Or does every citizen don personal “armor”? Issues of architectural alternatives, cost-effectiveness, and risks must be understood from a national standpoint. Both government-sponsored and private studies are needed.

In the case of electronic attacks on infrastructure components, including their information processing elements, the costs and benefits of incorporating *physical hardening* should be understood. This can be done through government-sponsored R&D, employing the extensive technology of electromagnetic hardening developed to protect military systems in nuclear environments. Commercial experience in reducing electromagnetic interference and in ensuring electromagnetic compatibility can also be applied to the problem. Regulatory proceedings can be used to elicit public response to proposed hardening measures for those infrastructure systems where they could be significant. In other cases, liability for endangering public safety can have an important influence on system owners and managers, especially as a body of supporting case law develops.

External Studies

A *theory* of the deterrence of large-scale infrastructure attacks is needed, just as in the early 1950s a theory of nuclear deterrence was formulated. Infrastructure attacks must be understood as economic attacks as well as precursors to or accompanied by conventional military action. The theory of nuclear deterrence rested on a secure second-strike capability and involved relatively clear and simple metrics of national strength. A dysfunctional national infrastructure may not provide an assured second-strike capability, nor even unambiguous evidence of responsibility for the attack. What will deter and what will not, and under what conditions, must be understood. Studies sponsored by and encouraged by the federal government are needed, joined by as many developed nations as choose to participate.

Construct *profiles* of attack sources to better understand attacker motivations and goals and the tools and experience they bring to an attack to help narrow the search for them.

Measures of effectiveness will have to be established to assist in formulating public policy and for guiding public and private investment in infrastructure protection.

Collect *past experiences* with complex systems, those that have been compromised and those that have not, to extract from them lessons that may be more broadly applicable to critical infrastructure systems. Particular care will have to be taken in studying past system intrusions to avoid the implication of mismanagement or failure and consequent potential liabilities.

Organizational Changes

Establish a point within the federal government for system owners and operators to *report* malicious threats against infrastructure under their jurisdiction. This is not to preclude industry initiatives of the same type, nor would reporting be mandatory. But the value of being able to collect threats for correlation and analysis would, if effective, be seen as mutually beneficial. For threats against federally owned or operated systems, reporting should be mandatory.

Establish threat *response teams* to respond to requests for assistance from infrastructure system operators. Such teams would be composed of technical experts, system operators, and law enforcement officials. They can be ad hoc, but a small permanent organization should be established to develop and refine the response process, advise system operators of its availability, and coordinate the organization and reporting of response teams.

The federal government should address the jurisdictionally complex question of its *own organization* to respond to infrastructure threats. When is a threat of a magnitude that it should be handled by a “CINC” for information infrastructure defense, when is it a domestic law enforcement matter, and when it is a private responsibility?

Set Standards

Standards for infrastructure security are needed so that users can be aware, in quantitative terms, of what protection is offered and where it fails, and of the relations between cost of protection and risk of its failure. Regulators of infrastructure systems have a primary responsibility to encourage their industries to take cooperative action through existing administrative procedures. Where such considerations are not on regulatory agendas, efforts are needed to raise the issues so that they can be put into perspective with other regulatory goals.

Cyber-attacks will require planning, and this can involve probes and exercises, all of which are susceptible to *sensing and analysis*. These will help establish both the identity of the attacker and the target of the attack. The reporting of anomalies in system behavior can provide information possibly bearing on attack planning.

Operational Activities

The federal government should support *R&D* on information infrastructure defense. Joint R&D activities with infrastructure operators and industry, including the computing, software, and communications industries and infrastructure equipment providers, should be encouraged.

Public agencies that operate infrastructure systems have a responsibility both to develop protection for those systems and to stimulate the *transfer* of such technology and their experiences to the private sector.

The federal government, in its role as collector and distributor of *national statistics*, can help increase public awareness of the losses from infrastructure attacks. There is a natural tendency for victims to suppress such information in order to avoid revealing system vulnerabilities and competition-sensitive information.

The federal government should participate with public and private operators in planning and executing *simulations, tests, and exercises* at the system level to help identify vulnerabilities and attack signatures and to test defensive concepts. Funding for such activities should be jointly operator and government funded. This will assist in achieving relevance and realism and will aid in transferring lessons learned to system operators.

Limiting Damage

Internal Studies

The launching of a counterattack to a cyber-attack requires *policy analysis* to understand scenarios and responses, both conventional and information-based. Attacks against U.S. military systems are at one extreme. Attacks from unknown or uncertain sources against privately owned civil infrastructure are at the opposite extreme.

Jawbone

A response to an attack on an infrastructure system can be to sense the target and the scope of the attack and to shed critical load to surviving parts of the infrastructure. This will require near real-time awareness of system load and available capacity and understandings between customers and operators of relative priorities. These can be reached through quality of service agreements, agreements between providers of infrastructure services, or through *state-federal regulatory processes*. To the extent that load-shedding will require capital expenditures to implement, such expenditures can be provided for either through *tax incentives* or by incorporation into *tariffs*.

External Studies

The Defense Center shown in Figure 4 would, in addition to its defense coordination role, serve as a way of constructing and preserving a *record* of real and suspected infrastructure attacks. The Center could also coordinate security experiments on infrastructure systems and perform detailed analysis of accidental system failures and suspected attacks. In the case of suspected attacks on private systems, the use of the Center’s services would be at the discretion of the system operator. But if the Center is of high quality, able to perform a unique function, and possessing data relating to and experience derived from system attacks, the demand for its services should be high. It would probably be most efficient if there were multiple Defense Centers specializing in the major types of infrastructure systems. But the set of Defense Centers would probably find it helpful to address some R&D, threat, and analysis issues jointly.

Joint government–industry working groups could be chartered to develop *future requirements* for system audit procedures and to provide information and education on the state of the technology.

In the military metaphor used here, counterattack has the sense of attack in kind (warhead against warhead or warhead against attacker). Cyber responses potentially offer richer sets of options; e.g., responding with a “marker” to label the attacker or attacker site. Developing *response options* should be a goal of public and private *R&D*.

Develop Software

Government infrastructure system operators could undertake *testing of commercial audit system products*, possibly under the supervision of the National Institute of Science and

Technology, with publication of the results available to both public agencies and private operators.

Government owned or operated infrastructure systems can provide an *early market* for commercial products. They can also, in principle, provide test beds for the development of security products, although the potential conflict between the reliable provision of service and a development environment is recognized. In the case of the development of the ARPANET the development community was also the user community. Some comparable solution will be needed in developing new approaches to the protection of infrastructure systems. An inventory of public infrastructure systems, or the public components of national systems, should be compiled and reviewed to see the extent to which they provide useful test beds and early market opportunities.

Where government funding of infrastructure protection systems is undertaken, the terms of the procurements should be such as to leave *commercial market rights* to the contractor in order to provide the greatest incentive to private developers and the least impediment to the transfer of new technology to the private sector.

Organizational and Policy Changes

The federal government might establish an Attack Assessment Center having as a long-term focus an educational program regarding threats; providing a directory of available technical expertise; a national-level center for reporting suspicious events; databases of real, suspected, and simulated events; and assistance in contacting government agencies having a charter in or jurisdiction over infrastructure systems.

Recognizing the effectiveness of the Software Engineering Institute in improving the software creation process, a similar approach is suggested for infrastructure systems by establishing a *System Security Institute*. It would, like the SEI, collect comparative information on current practices, develop measures of effectiveness of system security planning, run professional and educational programs to improve the state of the art, and advise policymakers as required on technical aspects of government initiatives. Like the SEI, it could be funded initially by the government but it should also have an industry membership program to deepen its detailed understanding of infrastructure system needs and to provide a distribution channel for its results. The SEI could assist in an SSI start-up, or such a mission could be added to the current organization. Other joint industry activities, such as the Electric Power Research Institute, could play a similar role.

The strongest possible commercially available encryption could be encouraged, including its export, to enable the most secure connectivity between global infrastructure operators and users. To the extent that this conflicts with other national security and law enforcement requirements, creating or licensing special classes of secure service for users requiring access to infrastructure control systems could be considered.

Regulation

Regulatory agencies could consider *mandating* the use of audit systems and establishing standards for assuring the integrity of the infrastructure systems under their jurisdiction. Public input to this process should be sought to balance the competing views of shareholders, ratepayers, and national leaders.

Regulatory agencies could require “autopsies” of accidental system failures in order to identify failure modes and to understand the details of system failure. In this way, rare system failures can be turned into experimental opportunities, to enable the collection of information not possible under even the most carefully controlled experimental conditions.

The incorporation of redundancy-based system architectures can be made part of future *federal procurements*. In cases where there is *federal sharing* of private infrastructure costs, there is also a possibility of introducing similar concepts into system design. Infrastructure investments can require cyber-attack *vulnerability analysis* for their approval.

OperationalActivities

An *assessment of attacks*, their source, their target, and their intent is needed. The Attack Assessment Center that has been suggested could provide the needed capability. While initially it would serve as an R&D organization tightly linked to operational facilities, as its capabilities expanded it could begin to provide attack warnings.

Strategic warning would deal with an impending attack, such as would come from the sensing of probes on an infrastructure system. It would allow defenses to be moved to a higher state of alert, requiring more thorough verification of individual system transactions. *Tactical warning* would allow near real-time blocking of attack transactions, or limiting the scope of possibly malicious transactions. Given the state of our understanding, tactical warning is probably not possible at this time, but, should such an approach to damage limitation be of interest, R&D to establish the basic parameters and to examine possible operational capabilities would be required.

Related to the provision of warning is the question of the *distribution* of warning information, since it will be required by private entities and, potentially, foreign entities. Public policy must be established and implementation options examined in order to do this.

Both public and private infrastructure systems have interfaces with comparable systems of other nations. Government participation in the engineering of *international interfaces* between national systems is of major importance. The International Telecommunications Union provides an excellent example of how worldwide public and private goals can be achieved. Comparable organizations exist for other infrastructure systems. The issue of system protection should be on their agendas. As the overseer of such relationships, the Department of State is a central agent to assist in this. Even where the international interactions are between private entities, governments play an important facilitating role.

Reconstitution

Jawbone

Government actions at all levels, especially by regulatory agencies, should encourage infrastructure managers to develop awareness and, where appropriate, to add to the robustness of systems under their jurisdiction.

Industry trade associations and professional societies can, through sponsorship of technical sessions at professional meetings, examine normally occurring emergencies in detail for lessons relevant to system damage. Government support for such activities, by direct funding, by providing expert participants, and through encouraging the release of information, can help direct the attention of engineering professionals to protection issues.

Public agencies including federal, state, and local emergency response organizations, military and national guard organizations, and public and private relief organizations are activated in regional and national crises. The effective application of relief resources and the allocation of surviving infrastructure to support *relief operations* is thus a high priority. This requires rapid assessment of damage: what still works and how well, assessments to support triage decisions, coordination of efforts and the allocation of responsibilities by agency and region, and the like. Governments have the organization and procedures needed to accomplish these tasks. So do system operators. It will be important to assure that these emergency response and system control functions are adequately connected, in terms of capacity and responsiveness. Emergencies of various kinds occur sufficiently often that response systems do not go untested. Nevertheless, such plans should be updated in light of the large-scale attacks possible by cyber-attacks on national infrastructures and in light of new technology that can both improve system robustness and at the same time introduce unexpected system vulnerabilities.

In addition to damage assessment for immediate reconstitution of service, the assessment will also be needed to support *claims* on insurers. Government can assist this process by making data available to those affected in order to support their claims.

External Studies

Federal, state, and local governments can undertake, through direct funding or by in-kind support, *simulations* of information system attacks and *games and exercises* to assist in bringing public and private parties into better position to understand the issues.

Fundamental to the technical protection of systems of systems, such as infrastructure systems, is the continued evolution of the system engineering discipline and the formal training of system engineers. To this end efforts should be made to identify ways to strengthen degree and certificate programs by postsecondary educational institutions.

Regulation

To the extent that infrastructure capacity is expanded to accommodate the kind of system attacks addressed here, a federal role could be to encourage consortia or alliances of service providers. The government could *encourage infrastructure strengthening and expansion* through regulatory processes and tax incentives.

Insurers have an important role to play in reconstitution after an attack through their contracts with policyholders. Governments step in when protection of the public is important, as in the case of automobile liability insurance, where government intervention has been to establish *minimum levels of coverage*. Government can also intervene to require insurers to provide minimal *coverage for high risk pools* as, for example, in fire-prone areas. Through regulation of insurers, government plays a role in *ensuring the financial health of insurers* and assuring the public a *minimum level of service*.

Operational Activities

All organizations have *disaster recovery plans*, and those plans are reviewed by various overseers: corporate management, insurers, and regulators. Recovery plans are also exercised in such emergencies as naturally occur, some of noteworthy size such as floods, hurricanes, regional power outages, etc. One can expect that as awareness of the possibility of cyber-attacks grows, and as small intrusions are experienced, disaster recovery plans will evolve to address these challenges.

Regulators, or other executive branch agencies, should encourage, with funding where necessary, the analysis of system and organizational *responses to normally occurring emergencies* of relevance to information system damage.

The federal government can reasonably inquire into disaster planning in privately owned and operated systems that have relevance to *national security*. In this way it can usefully “seed” interest in the issue.

In the event that there is substantial loss of infrastructure capacity over an extended period of time, rationing of the residual capacity may be a necessary response. The least intrusive approach to rationing is voluntary, and takes the form of requests to limit demand while reconstitution is effected. The next level is for government to put its authority behind the request, in effect *validating the need* for moderating demand for service. The most serious situation is when government must intervene in the market to *establish priorities* among uses and users to ensure that critical public needs are met. In this event, the government would manage resource allocation until the market resumes normal operation.

While this discussion has not addressed the specifics of infrastructure systems, one can consider the case where the infrastructure attack results in the loss of physical assets. While there is usually replacement hardware available somewhere, it is possible to imagine a situation in which delay in reconstituting a system could compound the damage. Should there be such cases, *stockpiling critical items* may make sense. Both public and private infrastructure operators need to examine worst-case situations to establish whether such eventualities can reasonably be expected.

Large-scale attacks on critical infrastructure have not occurred, so practical experience is lacking. What are needed are *large-scale attack scenarios* that can serve as a basis for estimating probabilities of such events, their spatial and temporal extent, and reconstitution requirements. The scenarios should be at a sufficiently technical level that they require understanding the details of system failure. In this sense they must go beyond the political-military scenarios that are useful in addressing higher-level leadership issues. Government-supported studies, undertaken by defense planners and system operators, can provide much needed guidance on which to base system and national response requirements.

In a dynamic business environment, corporate attention to what can be perceived as exotic threats can easily be dismissed. It is important, therefore, to address risks of loss due to infrastructure attacks as a financial issue as well as a purely national security issue. Given credible scenarios, especially if coupled with loss data drawn from past crises, it will be important to *introduce this thinking into the perspectives of corporate management.*

Emergency Response

As the scale of a disaster increases, many agencies of government will play larger roles, much as the Federal Emergency Management Agency does in coordinating responses in regional disasters. Disaster victims become eligible under federal relief programs. The DoD, and National Guard units, often play critical roles in assisting local recovery actions, by deploying units and by providing critical *logistics and communication capabilities*.

Concluding Observations

While infrastructure systems are mainly privately owned, important infrastructure, particularly military systems, is publicly owned. In pursuing the objective of seeing to the protection of infrastructure systems generally, the federal government could usefully focus on those systems it is responsible for and share its lessons and experiences. A problem in doing this, one common to all having system responsibilities, is how much to reveal, lest one compromise a system by revealing its vulnerabilities and its protections.

Even after taking a position of minimal government roles in infrastructure protection, a number of possible actions present themselves. These ideas are presented here with no attempt to evaluate their advantages and disadvantages, costs, risks, and effectiveness. They are merely a set of possibilities, selections from which can be made either after a national consensus on strategy has been reached or as part of the process of achieving consensus. Both public and private implementation plans would then be required, consistent with the chosen strategy, with future system investments, and with perceptions of threat.

Federal government "handles" on the problem include performing studies and analyses; collecting national level statistics of risks; assisting in the performance or coordination of tests and exercises; responsibility for or oversight of foreign interfaces in the areas of system operation, economic impacts, and law enforcement; intelligence on threats: and the regulatory structure that includes politically expressed goals, legal protection of rights and the adjudication of disputes, and the administration of legislation, executive branch rules, and judicial decisions.

While at lower levels of violence grave danger to public order is not an issue, this is nevertheless an important part of the problem. Not only is the hacker community worldwide a training ground for more serious attackers, but by presenting us with numerous examples of cleverly executed intrusions, they provide incontrovertible evidence of vulnerabilities of systems we might otherwise overlook.

If private system owners are to change the scale of their planned investments in system protection, it will be necessary to provide them with credible scenarios describing state-supported attacks on information infrastructures, including potential attackers, their goals and objectives, and potential losses. This last is most important, yet it is the most difficult part of scenario construction. Since cyber-attacks on U.S. infrastructure have not occurred, there is no basis for making credible projections, despite ample evidence from the strategic bombing of wartime adversaries from World War II to the Gulf War and from experiences with "normal" system failures and natural disasters. Such evidence should be brought to bear to better support "consequences of execution" analyses.

Without a convincing showing to the contrary, system operators will base their security investments on their own assessments of the threat. These may be adequate, especially as they adapt to what can be expected to be a continually escalating level of sub-national attacks and losses. The central role for government, then, is to focus attention on issues and actions at the national "systems-of-systems" level, including system interactions above those addressable by any single operator or even single industry.

While the question of relative priorities is beyond the scope of this discussion, two views have been suggested in this paper. In a prevention-dominant strategy, one might, referring back to Figure 1, invest 70 prevent of the resources allocated to infrastructure protection to prevention, 20 percent to damage limitation, divided equally between short- and long-term approaches, and the remaining 10 percent to reconstitution. A different position has also been noted. The basis for this is the pessimistic view that adequate resources to prevent infrastructure attacks will not be available, so that such resources as are available had better be spent on damage limitation, especially to limit the damage from a single attack. In that case, one might allocate 60 percent to single-attack damage limitation, 20 percent to prevention, and the remaining 20 percent equally divided between long-term damage limitation and reconstitution. But the choice of specific approaches, relative resource allocations, and program design and implementation should be based on more detailed analyses than have probably been done to date.

Throughout the discussion, “government” has been used as a broad and inclusive term, with relatively few references to specific branches, departments, agencies, and programs, whether state or federal, domestic or international. There are several questions of the “Who’s in charge?” variety to address before going into details.

Is the protection of critical infrastructure one problem or several problems, depending on the infrastructure in question? The implication at several points in this analysis is that it is several problems, not a single problem. In suggesting that separate infrastructures be treated in separate ways, the essential synergy between hardware and software approaches to solutions, between threat and attack assessments for all systems, and in projecting a uniform position when dealing with international issues and interfaces must be kept in mind, however.

Does one set up a single office of responsibility, or is the problem better handled by “matrixing” the specialists in various parts of the government? In the latter case, how is a coherent long-term national program funded and managed? In the former case, at what level of government should the office of responsibility be located? Too high and it lacks effective contact with government expertise; too low and it lacks authority.

How is state-federal coordination to be accomplished? How does one relate new initiatives with what is already under way, however fragmented and inadequately funded it may be? What international posture should be adopted? Is this a domestic problem requiring some international assistance on occasion, or a problem best addressed by existing or new international bodies?

Is this a “military” threat best given to DoD, a problem better handled through foreign intelligence channels, a law enforcement problem, or one with such major technical uncertainties that it is better given to a technical agency like DARPA for a while?

At a time when new government initiatives face an uphill fight for funding, the natural tendency is to seek maximum leveraging of available resources. Private industry owns the bulk of the infrastructure in question, has natural incentives to protect its investment and its customers, and makes investments to enhance the return on its assets. It would seem, therefore, most effective to supply private infrastructure operators with the best available assessments of threats, ask them to match these to system vulnerabilities, assist in those areas proposed by private industry, and act as a balance wheel in areas where industry may be less able to act. The point is not to argue how to best set up new governmental structures but to ask how to make possibly uncoordinated private investments more effective. In this regard, current regulatory arrangements need careful review. Where regulators can help but are unaware of the issues, they should be encouraged to do so.

Center for International Security and Arms Control

Stanford University

Please send orders to: Publications, 320 Galvez Street, Stanford, California 94305-6165. Enclose check payable to Stanford University. Add $2.00 postage and handling for first item ordered ($5.00 for overseas delivery), $1.00 for each additional item. Foreign orders must be in U.S. dollars and drawn on a financial institution with branches in the United States. California residents, add appropriate sales tax.

MacArthur Consortium Working Papers in Peace and Cooperation

Tarak Barkawi. *Democracy, Foreign Forces, and War: The United States and the Cold War in the Third World*. 1996 (40 pages, $6.00).

Byron Bland. *Marching and Rising: The Rituals of Small Differences and Great Violence in Northern Ireland*. 1996 (32 pages, $6.00).

Charles T. Call. *From "Partisan Cleansing" to Power-Sharing? Lessons for Security from Colombia's National Front*. 1995 (60 pages, $7.00).

David Dessler. *Talking Across Disciplines in the Study of Peace and Security: Epistemology and Pragmatics as Sources of Division in the Social Sciences*. 1996 (40 pages, $7.00).

Lynn Eden and Daniel Pollak. *Ethnopolitics and Conflict Resolution*. 1995 (21 pages, $5.00).

Daniel T. Froats, *The Emergence and Selective Enforcement of International Minority-Rights Protections in Europe after the Cold War*. 1996 (40 pages, $7.00).

NEW Robert Hamerton-Kelly. *An Ethical Approach to the Question of Ethnic Minorities in Central Europe: The Hungarian Case*. 1997 (34 pages, $6.00).

Bruce A. Magnusson. *Domestic Insecurity in New Democratic Regimes: Sources, Locations, and Institutional Solutions in Benin*. 1996 (28 pages, $6.00).

John M. Owen. *Liberalism and War Decisions: Great Britain and the U.S. Civil War*. 1996 (22 pages, $5.00).

Center reports, working papers, and reprints

Andrei Baev, Matthew J. Von Bencke, David Bernstein, Jeffrey Lehrer, and Elaine Naugle. *American Ventures in Russia*. Report of a Workshop on March 20-21, 1995, at Stanford University. 1995 (24 pages, $7.00).

David Bernstein. *Software Projects in Russia: A Workshop Report*. 1996 (28 pages, $7.00).

David Bernstein, editor. *Defense Industry Restructuring in Russia: Case Studies and Analysis*. 1994 (244 pages, $14.00).

George Bunn. *Does the NPT require its non-nuclear-weapon members to permit inspection by the IAEA of nuclear activities that have not been reported to the IAEA?* 1992 (12 pages, $4.00).

Irina Bystrova. *The Formation of the Soviet Military-Industrial Complex*. 1996 (28 pages, $6.00).

Cooperative Security in Northeast Asia (text in English and Russian) 1993 (17 pages, $4.00).

John S. Earle and Ivan Komarov. *Measuring Defense Conversion in Russian Industry*. 1996 (40 pages, $7.00).

John S. Earle and Richard Rose. *Ownership Transformation, Economic Behavior, and Political Attitudes in Russia*. 1996 (40 pages, $7.00).

David Elliot, Lawrence Greenberg, and Kevin Soo Hoo. *Strategic Information Warfare—A New Arena for Arms Control?* 1997 (16 pages, $3.00).

Anthony Fainberg. *Strengthening IAEA Safeguards: Lessons from Iraq*. 1993 (64 pages, $6.00).

NEW James E. Goodby. *Can Strategic Partners Be Nuclear Rivals?* (First in a series of lectures on *The U.S.–Russian Strategic Partnership: Premature or Overdue?*) 1997 (26 pages, $6.00).

NEW James E. Goodby. *Loose Nukes: Security Issues on the U.S.–Russian Agenda*. (Second in a series of lectures on *The U.S.–Russian Strategic Partnership: Premature or Overdue?*) 1997 (19 pages, $6.00).

NEW Seymour Goodman. *The Information Technologies and Defense: A Demand-Pull Assessment*. 1996 (48 pages, $9.00).

NEW Lawrence T. Greenberg, Seymour E. Goodman, and Kevin J. Soo Hoo. *Old Law for a New World? The Applicability of International Law to Information Warfare*. 1997 (48 pages, $8.00).

John R. Harvey, Cameron Binkley, Adam Block, and Rick Burke. *A Common-Sense Approach to High-Technology Export Controls*. 1995 (110 pages, $15.00).

John Harvey and Stefan Michalowski. *Nuclear Weapons Safety and Trident*. 1993 (104 pages, $12.00; summary $2.00).

Ji, Guoxing. *Maritime Security Mechanisms for the Asian-Pacific Region*. 1994 (25 pages, $5.00).

Leonid Kistersky. *New Dimensions of the International Security System after the Cold War*. 1996. (34 pages, $8.00)

Amos Kovacs, *The Uses and Nonuses of Intelligence*. 1996 (68 pages, $10.00).

Allan S. Krass. *The Costs, Risks, and Benefits of Arms Control*. 1996 (85 pages, $8.00).

Gail Lapidus and Renée de Nevers, eds. *Nationalism, Ethnic Identity, and Conflict Management in Russia Today*. 1995 (106 pages, $12.00).

John J. Maresca. *The End of the Cold War Is Also Over*. With commentaries by Norman M. Naimark, Michael May, David Holloway, Arthur Khachikian, Daniel Sneider, and Renée de Nevers. 1995 (60 pages, $8.00).

Center for International Security and Arms Control

Stanford University

Please send orders to: Publications, 320 Galvez Street, Stanford, California 94305-6165. Enclose check payable to Stanford University. Add $2.00 postage and handling for first item ordered ($5.00 for overseas delivery), $1.00 for each additional item. Foreign orders must be in U.S. dollars and drawn on a financial institution with branches in the United States. California residents, add appropriate sales tax.

MacArthur Consortium Working Papers in Peace and Cooperation

Tarak Barkawi. *Democracy, Foreign Forces, and War: The United States and the Cold War in the Third World*. 1996 (40 pages, $6.00).

Byron Bland. *Marching and Rising: The Rituals of Small Differences and Great Violence in Northern Ireland*. 1996 (32 pages, $6.00).

Charles T. Call. *From "Partisan Cleansing" to Power-Sharing? Lessons for Security from Colombia's National Front*. 1995 (60 pages, $7.00).

David Dessler. *Talking Across Disciplines in the Study of Peace and Security: Epistemology and Pragmatics as Sources of Division in the Social Sciences*. 1996 (40 pages, $7.00).

Lynn Eden and Daniel Pollak. *Ethnopolitics and Conflict Resolution*. 1995 (21 pages, $5.00).

Daniel T. Froats, *The Emergence and Selective Enforcement of International Minority-Rights Protections in Europe after the Cold War*. 1996 (40 pages, $7.00).

NEW Robert Hamerton-Kelly. *An Ethical Approach to the Question of Ethnic Minorities in Central Europe: The Hungarian Case*. 1997 (34 pages, $6.00).

Bruce A. Magnusson. *Domestic Insecurity in New Democratic Regimes: Sources, Locations, and Institutional Solutions in Benin*. 1996 (28 pages, $6.00).

John M. Owen. *Liberalism and War Decisions: Great Britain and the U.S. Civil War*. 1996 (22 pages, $5.00).

Center reports, working papers, and reprints

Andrei Baev, Matthew J. Von Bencke, David Bernstein, Jeffrey Lehrer, and Elaine Naugle. *American Ventures in Russia*. Report of a Workshop on March 20-21, 1995, at Stanford University. 1995 (24 pages, $7.00).

David Bernstein. *Software Projects in Russia: A Workshop Report*. 1996 (28 pages, $7.00).

David Bernstein, editor. *Defense Industry Restructuring in Russia: Case Studies and Analysis*. 1994 (244 pages, $14.00).

George Bunn. *Does the NPT require its non-nuclear-weapon members to permit inspection by the IAEA of nuclear activities that have not been reported to the IAEA?* 1992 (12 pages, $4.00).

Irina Bystrova. *The Formation of the Soviet Military-Industrial Complex*. 1996 (28 pages, $6.00).

Cooperative Security in Northeast Asia (text in English and Russian) 1993 (17 pages, $4.00).

John S. Earle and Ivan Komarov. *Measuring Defense Conversion in Russian Industry*. 1996 (40 pages, $7.00).

John S. Earle and Richard Rose. *Ownership Transformation, Economic Behavior, and Political Attitudes in Russia*. 1996 (40 pages, $7.00).

David Elliot, Lawrence Greenberg, and Kevin Soo Hoo. *Strategic Information Warfare—A New Arena for Arms Control?* 1997 (16 pages, $3.00).

Anthony Fainberg. *Strengthening IAEA Safeguards: Lessons from Iraq*. 1993 (64 pages, $6.00).

NEW James E. Goodby. *Can Strategic Partners Be Nuclear Rivals?* (First in a series of lectures on *The U.S.–Russian Strategic Partnership: Premature or Overdue?*) 1997 (26 pages, $6.00).

NEW James E. Goodby. *Loose Nukes: Security Issues on the U.S.–Russian Agenda*. (Second in a series of lectures on *The U.S.–Russian Strategic Partnership: Premature or Overdue?*) 1997 (19 pages, $6.00).

NEW Seymour Goodman. *The Information Technologies and Defense: A Demand-Pull Assessment*. 1996 (48 pages, $9.00).

NEW Lawrence T. Greenberg, Seymour E. Goodman, and Kevin J. Soo Hoo. *Old Law for a New World? The Applicability of International Law to Information Warfare*. 1997 (48 pages, $8.00).

John R. Harvey, Cameron Binkley, Adam Block, and Rick Burke. *A Common-Sense Approach to High-Technology Export Controls*. 1995 (110 pages, $15.00).

John Harvey and Stefan Michalowski. *Nuclear Weapons Safety and Trident*. 1993 (104 pages, $12.00; summary $2.00).

Ji, Guoxing. *Maritime Security Mechanisms for the Asian-Pacific Region*. 1994 (25 pages, $5.00).

Leonid Kistersky. *New Dimensions of the International Security System after the Cold War*. 1996. (34 pages, $8.00)

Amos Kovacs, *The Uses and Nonuses of Intelligence*. 1996 (68 pages, $10.00).

Allan S. Krass. *The Costs, Risks, and Benefits of Arms Control*. 1996 (85 pages, $8.00).

Gail Lapidus and Renée de Nevers, eds. *Nationalism, Ethnic Identity, and Conflict Management in Russia Today*. 1995 (106 pages, $12.00).

John J. Maresca. *The End of the Cold War Is Also Over*. With commentaries by Norman M. Naimark, Michael May, David Holloway, Arthur Khachikian, Daniel Sneider, and Renée de Nevers. 1995 (60 pages, $8.00).

Michael May. *Rivalries Between Nuclear Power Projectors: Why the Lines Will Be Drawn Again*. 1996 (20 pages, $7.00).

Michael May and Roger Speed. *The Role of U.S. Nuclear Weapons in Regional Conflicts*. 1994 (24 pages, $5.00).

Michael McFaul, ed. *Can the Russian Military-Industrial Complex Be Privatized?* 1993 (60 pages, $6.00).

Robert F. Mozley. *Uranium Enrichment and Other Technical Problems Relating to Nuclear Weapons Proliferation*. 1994 (64 pages, $9.00)

William J. Perry. *Defense Investment: A Strategy for the 1990s.* 1989 (43 pages, $9.00).

Scott D. Sagan, ed. *Civil-Military Relations and Nuclear Weapons*. 1994 (163 pages, $12.00).

Scott D. Sagan and Benjamin A. Valentino. *Nuclear Weapons Safety after the Cold War: Technical and Organizational Opportunities for Improvement* (text in English and Russian). 1994 (25 pages, $5.00).

Capt. Alexander Skaridov, Cmdr. Daniel Thompson, and Lieut. Cmdr. Yang Zhiqun. *Asian-Pacific Maritime Security: New Possibilities for Naval Cooperation?* 1994 (28 pages, $5.00).

Song, Jiuguang. *START and China's Policy on Nuclear Weapons and Disarmament in the 1990s*. 1991 (29 pages, $5.00).

Konstantin Sorokin. *Russia's Security in a Rapidly Changing World.* 1994 (95 pages, $10.00).

Roger D. Speed. *The International Control of Nuclear Weapons*. 1994 (59 pages, $11.00).

NEW István Szönyi. *The False Promise of an Institution: Can Cooperation between OSCE and NATO Be a Cure?* 1997 (34 pages, $6.00).

Terence Taylor. *Escaping the Prison of the Past: Rethinking Arms Control and Non-Proliferation Measures*. 1996 (65 pages, $10.00)

Terence Taylor and L. Celeste Johnson. *The Biotechnology Industry of the United States. A Census of Facilities*. 1995 (20 pages, $7.00).

Selected books available from other publishers

Herbert L. Abrams. *The President Has Been Shot: Confusion, Disability, and the 25th Amendment in the Aftermath of the Assassination Attempt on Ronald Reagan.* New York: W.W. Norton, 1992.

Coit D. Blacker. *Hostage to Revolution: Gorbachev and Soviet Security Policy.* New York: Council on Foreign Relations, 1993.

George Bunn. *Arms Control by Committee: Managing Negotiations with the Russians*. Studies in International Security and Arms Control. Stanford: Stanford University Press, 1992.

Gordon H. Chang. *Friends and Enemies: The United States, China, and the Soviet Union, 1948–1972.* Stanford: Stanford University Press, 1990.

Sergei Goncharov, John W. Lewis, and Xue Litai. *Uncertain Partners: Stalin, Mao, and the Korean War.* Stanford: Stanford University Press, 1993.

Seymour Goodman, Peter Wolcott, and Grey Burkhart. *An Examination of High-Performance Computing Export Cont rol Policy in the 1990s.* Los Alamitos, CA: IEEE Computer Society Press, 1996.

Robert Hamerton-Kelly. *The Gospel and the Sacred: The Poetics of Violence in the Gospel of Mark.* Fortress Press, 1994.

David Holloway and Norman Naimark, editors. *Reexamining the Soviet Experience: Essays in Honor of Alexander Dallin*. Boulder, CO: Westview Press, 1996.

David Holloway. *Stalin and the Bomb: The Soviet Union and Atomic Energy, 1939–1956.* New Haven: Yale University Press, 1994.

John Wilson Lewis and Xue Litai. *China's Strategic Seapower: The Politics of Force Modernization in the Nuclear Age.* Studies in International Security and Arms Control. Stanford: Stanford University Press, 1994.

Michael McFaul. *Post-Communist Politics: Democratic Prospects in Russia and Eastern Europe.* Washington, D.C.: CSIS, 1993.

Michael McFaul and Sergei Markov. *The Troubled Birth of Russian Democracy: Parties, Personalities, and Programs*. Stanford: Hoover Press, 1993.

Norman M. Naimark. *The Russians in Germany: A History of the Soviet Zone of Occupation*. Cambridge: Belknap/Harvard University Press, 1995.

Scott D. Sagan and Kenneth N. Waltz. *The Spread of Nuclear Weapons: A Debate*. New York: W W. Norton, 1995.

Scott D. Sagan. *The Limits of Safety: Organizations, Accidents, and Nuclear Weapons*. Princeton: Princeton University Press, 1993.

Judith Sedaitis, ed. *Commercializing High Technology: East and West*. Lanham, MD: Rowman & Littlefield Publishers, 1997.

Condoleezza Rice and Philip Zelikow. *Germany Unified and Europe Transformed: A Study in Statecraft.* Cambridge: Harvard University Press, 1995.