Chapter Eleven

AN EXPLORATION OF CYBERSPACE SECURITY R&D INVESTMENT STRATEGIES FOR DARPA*

Robert H. Anderson and Anthony C. Hearn

INTRODUCTION

"The Day After . . ." exercise methodology, developed over the past several years under the leadership of Roger Molander, has proven useful in eliciting thinking about complex strategic issues from groups of up to about 60 individuals. The exercises are also useful in "awareness building"—exposing participants to the possible ramifications of current trends, and options for altering those trends. For examples of previous uses of this methodology to explore the national security policy implications of the continued diffusion of nuclear weapons capabilities, see Millot, Molander and Wilson (1993); Mesic, Molander and Wilson (1995); Molander, Wilson, Mesic and Gardiner (1994); and Molander, Riddile and Wilson (1995). A recent application of the methodology to issues of strategic information warfare is presented in Molander, Riddile and Wilson (1996).

The U.S. Defense Advanced Research Projects Agency (DARPA) is interested in understanding strategies for the investment of research and development funds for securing the U.S. information infrastructure against "information warfare" (IW) attacks. (As Roger Molander

*Robert H. Anderson and Anthony C. Hearn, *An Exploration of Cyberspace Security R&D Investment Strategies for DARPA: "The Day After—in Cyberspace II,"* Santa Monica, Calif.: RAND, MR-797-DARPA, 1996. Some figures, tables, and text were omitted for this version.

put it, tongue in cheek, during his opening remarks at the exercise described in this report: "OK, you guys built the ARPAnet, which has become the Internet; now fix it!") A variety of recent studies (e.g., Hundley and Anderson, 1995) have documented the web of interrelated information systems comprising the national information infrastructure and its heavy dependence on the public switched telephone network. These systems are attacked every day by hackers worldwide and, less commonly but more insidiously, by trusted insiders, organized groups, commercial organizations, intelligence agencies, and other agencies of foreign governments. As our society becomes more dependent on this information infrastructure, concern rises about what strategies and technology might best be employed to substantially strengthen the infrastructure against deliberate attacks.

The Purpose of This Exercise

The purpose of this particular exercise was "to conduct an exercise informing ARPA staff and selected representatives of the user community of the principal features of (defensive) information warfare (IW) and identifying for participants the future demands that IW may place on ARPA information technology programs."[1] Dr. Howard Frank of DARPA's Information Technology Office acted as the project monitor.

In subsequent discussions with Dr. Frank and among RAND staff, we referred to the exercise purpose as helping inform DARPA's investment strategy for research and development on the integrity and reliability of information systems on which the security and safety of the nation depend.

The Scenario and Methodology Used for This Exercise

The original "The Day After . . ." exercise methodology used a three-step process: (1) preparing a memo to a senior government executive regarding problems occurring about five years in the future, in the early stages of a crisis; (2) addressing additional problems several days to a week later, as the crisis worsens; and (3) preparation of a memo "today" (i.e., 1996) discussing measures that should be taken now to avoid problems such as those described in steps 1 and 2.[2]

In several dry runs of the DARPA exercise, conducted using RAND staff both in Santa Monica and in Washington D.C., we determined that participants became frustrated in steps 1 and 2 because there was little that could be done in the short term to ameliorate or halt the series of cyberspace-based attacks on the U.S. infrastructure. Participants also felt that there was too little time left in the exercise to discuss possible R&D programs that could be instituted today to prevent or greatly reduce such attacks in the future. For these reasons, we decided to modify the exercise so that it contained just two steps: (1) IW attacks occurring five years in the future; and (2) a discussion of what could be done beginning today to cope better with those future attacks.

A second dry run using this new methodology proved successful. Participants developed heightened awareness of the problems that could be encountered in the future in Step 1, but then had ample time left to discuss R&D measures in the new Step 2. Because the purpose of this exercise was to develop R&D strategies, this new two-step approach was clearly superior for our purposes.

We began with an existing scenario of cyberspace attacks on U.S. infrastructure used in a previous exercise[3] and tuned and expanded the cyberspace attacks for our particular purposes. We wanted to illustrate the diversity of infrastructure systems dependent on "cyberspace" that might be subject to attack, from transportation control systems to power control to key financial systems. Since the participants for this exercise were to be technologically sophisticated, we added some indications of *how* these attacks might be performed, to increase their believability and counter any possible reactions that "that couldn't possibly happen!"

The set of cyberspace incidents we evolved for the scenario used in this exercise is shown in Table 11.1.

The Conduct of the Exercise

The exercise was held on Saturday morning, March 23, 1996, in RAND's Washington, D.C. offices. After a plenary introductory

Table 11.1
Cyberspace Incidents Used in Scenario

Year 2000 background	
general	software agents roaming net and Web
1999	MEII discussed but not yet established
1998	electronic "looting" of Saudi Arabian bank ($1.2 billion)
1999	attempted placement of Trojan horse in AB-330 flight control software
1999	sniffers and logic bombs in Israeli C2 systems
general	electronic "looting" of U.S. and European banks by Russians
1998	computer virus in software causes Yen crisis in Japan
1998–99	Infonet Threat Center established in U.S.
1999	flight control software alert regarding U.S. commercial aircraft
The Crisis—Step 1	
2000 May 11	power in Cairo (90%) out for several hours — perpetrator uncertain
2000 May 11	public switched telephone network (PSTN), massive failure in Riyadh, Saudi Arabia
2000 May 11	PSTN, Ft. Lewis, WA, mass dialing attack
2000 May 11	Saudi PSTN, apparent "trap door" in switching code
2000 May 13	control malfunction, Aramco refinery, Saudi Arabia — perpetrator uncertain
2000 May 14	control malfunction, Bundesbahn train crash, Germany — perpetrator uncertain
2000 May 16	sniffers, Bank of England funds transfer system
2000 May 16	power grid for Rhein Main airbase, Germany, fails
2000 May 17	non-governmental organization "Consortium for Planetary Peace" mobilization via Internet and other media
2000 May 18	PSTN in Delaware and Maryland fails — affects air traffic control at Dover AFB
Continuing Crisis—Step 1	
2000 May 20	Automated Teller Machine networks malfunction in Georgia
2000 May 20	CNN off air for 12 minutes; issues special report
2000 May 20	worm, corrupting data in Time Phased Force Deployment List (TPFDL)
2000 May 22	flight control software malfunction; AB-340; plane crash at O'Hare
2000 May 22	recommendation that all late-model AB-340 and -330s be grounded
2000 May 22	TV signal in Saudi Arabia replaced by other broadcast
2000 May 23	PSTN, Saudi, fails; trap doors similar to earlier Saudi PSTN failure
2000 May 23	full-scale IW attack at CONUS military bases involved in deployment
2000 May 23	Chicago Commodity Exchange subjected to electronic manipulation
2000 May 23	PSTN failed, Wash./Baltimore area, similar to Saudi PSTN failure

session to review the scenario and some recent developments, approximately 60 participants were placed into five groups of about 12 persons each to discuss the Step 1 scenario.

In Step 1, participants were told to act as members of "a technical tiger team advising the Secretary of Defense and the Director of ARPA, in a time-urgent process. The group's task is to revise a draft memo to the SECDEF in preparation for the ARPA Director's meeting with the SECDEF scheduled for a few hours hence."[4]

In Step 2, participants were brought back to the "very near future—say the late spring of 1996." They were told that they were "again in the role of a top advisor to the Director of ARPA, preparing him for a meeting with the Secretary of Defense on a national R&D investment strategy for information systems security and related issues."[5]

The following section contains findings and research suggestions resulting from the groups' deliberations.

FINDINGS AND RESEARCH SUGGESTIONS

The format of the exercise, described in the previous section, lends itself naturally to two types of observations and findings: those from Step 1, involving short-term actions that can be taken to reduce or ameliorate a set of cyberspace incidents in progress; and those from Step 2 regarding longer-term research and development initiatives that might prevent or greatly reduce the likelihood of such incidents occurring in the future. We present below the key findings and recommendations from group deliberations of steps 1 and 2, concentrating on new observations arising from the discussions, rather than ideas presented in the draft memos given to the participants to stimulate their discussion. The materials presented in this section result both from the group presentations at the plenary sessions and from notes taken by RAND observers who monitored the deliberations of each individual group.

Step 1. Observations and Findings

At the conclusion of their deliberations regarding the Step 1 incidents occurring in the year 2000, the five groups presented the following observations and findings. In what follows, we have edited

their remarks to omit obvious and redundant observations, concentrating on items that might affect DARPA research and development investment initiatives.

In the following discussion, we do not rigidly follow the structure of the "Memo to the SECDEF" in Step 1 of the scenario, because the issues raised there are primarily oriented toward "consciousness-raising" among the participants. Since the scenario in the year 2000 is hypothetical, so are the explicit recommendations made in response to it. We concentrate instead on broader observations about the state of U.S. information vulnerability in the year 2000 and on the tradeoffs and compromises that might be required to deal with attacks on that vulnerability.

"Safe Havens" Should Be Developed As a Fallback Means for Systems When Under Attack. The information systems supporting our nation's infrastructure have become increasingly interconnected during the past several decades. Regional power grids now exchange information and signals more substantially than before; the more than 1500 telecommunication companies providing public-switched telephone service share a common signaling system; and financial trading and exchange systems are linked worldwide with real-time networks. Because of these interdependencies, a vulnerability in one portion of a system can be used to exploit, disrupt, or deny service in other portions—at times geographically remote from the original source of entry.

A possible solution strategy to this problem is to configure these infrastructure systems so that they can quickly be isolated into self-sufficient regional systems. If, in a matter of seconds or minutes, the energy grids or telecommunication systems could be isolated into smaller units, the resulting smaller units might become safe havens protected from remote attack. At a later safe time, the units might be reassembled into an interconnected system. (See the suggestion on the use of "human firewalls" to oversee this reconnection process, under the subhead "Operational aspects of security . . ." below.)

It was also mentioned that key portions of the infrastructure should have backup repositories of software code (e.g., for telecommunication switches) positioned locally, stored in a manner in which such code can be verified as authentic and accurate. This code could be

used for "rebaselining" systems that may have been corrupted. Its local storage is important in case the system in question has been disconnected from other systems, which might prevent downloading the code from a central repository.

Tactical Warning/Attack Assessment (TW/AA) Is an Important Concept for Cyberspace Security. There was considerable discussion (prompted by the draft memo to the SECDEF that was part of the Step 1 materials) regarding the concepts of tactical warning and attack assessment.[6] It was agreed that TW/AA is important, and that there is currently little infrastructure in place to perform these activities.

The main reaction was "Who's in charge?" For TW/AA to be successful, there must be a clearinghouse (a "National IW Center"?) to collect, collate, and uncover patterns in cyberspace attacks that span systems in all key infrastructures: transportation, power, finance, communication, defense, and so forth. At present, there is no agency or entity that is mandated/empowered to collect this information, much less process it.

It was noted that, if such a center existed, it would need software tools to distinguish coordinated attacks from uncoordinated ones.

One possible activity of such a coordinating center would be to design and implement "trigger levels" of activity that would cause alerts to be broadcast to key parts of the U.S. information infrastructure. These alerts might be analogous to the DoD "DEFCON" levels used to represent the state of alert for Defense organizations.

Operational Aspects of Security (Dealing with People, Procedures, Regulations) Are Vitally Important to Any Solution. Although this exercise was focused on R&D initiatives of the type DARPA typically supports, there was considerable discussion of "operational" aspects of security that may be less amenable to R&D, but are deemed vitally important to any overall security posture. It was clear that issues related to people, procedures, regulations, training, education, and so on were a critical adjunct to any successful security technology initiative.

The following operational aspects were specifically mentioned:

The concept of "cyberspace hot pursuit" needs attention. We need software tools to aid in the backtracing of incidents, to discover the perpetrator. As such backtracing begins within the U.S. but then crosses country borders, we need clear laws and regulations stating which U.S. or international agencies are authorized to conduct such "cyberspace pursuits," what cooperation should be expected from foreign governments and organizations, and what might be done (in real time, if possible) to disable the means by which the perpetrator is instigating the incidents.

We need procedures for the prepositioning of backup systems and software. As mentioned above, the concept of "safe havens" in information systems was discussed, along with the related idea of prepositioning verifiably accurate software (and possibly hardware) for rebaselining corrupted systems. Are there standard procedures that can be developed and used for such baselining? Is each portion of the infrastructure responsible for prepositioning needed systems components, or is some more central organization and coordination desirable?

"Red teams" are needed to test system defenses. The groups tended to concur that active testing of system defenses is an important means for assessing system security. The pioneering tests by the Defense Information Systems Agency (DISA) and the Air Force Information Warfare Center (AFIWC) at Kelly Air Force Base are examples of such testing. The testing concept should be expanded to cover all key national information infrastructure systems. Among the questions needing attention are: What agencies should do the testing? Under what auspices? Would such testing be voluntary or mandatory? What safeguards are needed to protect against unintentional damage or denial of service in these infrastructures as the result of tests? What are the possible legal liabilities as a result of such tests?

Map the networks. Cyberspace is a loose concept describing interconnected information systems, with the Internet and the telephone system (PSTN) on which it depends as key—but certainly not the only—components. We need maps of the interconnections among the networks of cyberspace to resolve a number of questions, such as: How do energy grid control systems depend on the PSTN? If a perpetrator appears to be linking into the networks from Iran, or North Korea, or wherever, what are the routes that he or she may

take, and can they be blocked? Some agency(ies) should be tasked with maintaining an updated map of the tens of thousands of links and interrelationships and interdependencies among key networks. A subsequent question then arises: Would that map then be widely available to inform discussions of cyberspace security, or classified so that only a select few could access it?

Personal ID verification systems should be employed. Participants felt it was important to employ such systems on all links into the infrastructure, including access through dial-in maintenance ports. In this way perpetrators may have an additional hurdle to cross, and an audit trail can be maintained to assign responsibility or blame for incidents.

The concept of "human firewalls" should be considered in an emergency. As systems are decomposed into "safe havens" (see above) when an attack is imminent, or during an attack, it might be possible to insert a human as an intelligent verification device to pass judgment before various people and systems are allowed to obtain access to critical nodes and links in the infrastructure.

A "two-person rule" might be used for critical decisions or system changes. Just as firing a nuclear missile requires the cooperation of (at least) two individuals, we should consider the advantages (weighed against additional costs and impediments) of requiring two persons to authorize and allow any key change to critical system software, or to implement a decision regarding critical links or nodes. This idea would require considerable analysis to see if it could be practical. See also the discussion of the need for research on the design of secure information systems, below. The "two-person rule" might be a part of the procedures for secure system design and implementation.

Consider better pay and status for critical system operators. Personnel might then be less vulnerable to bribes, and less likely to become disgruntled or disaffected. It is widely understood that the trusted insider poses the greatest threat to critical information systems.

Some Notable Quotations Recorded During Step 1 Deliberations. We thought the following comments added information and insight to the proceedings, and were worthy of retention.

"If the power system is at risk, everything is at risk."

Many felt that the power system was critical to literally every other component of the infrastructure.

"Corrupting compilers is a very powerful, invidious attack."

Control of compilers is a key component of an overall secure process for software development.

"There are several examples already where perpetrators have spent 18 months inserting trapdoors, etc., into financial software before beginning to steal money."

Carefully orchestrated and planned attacks are being seen, not just hackers doing their thing.

"The U.S. has two main tasks (when under cyberspace attack): (1) recover from what has occurred; and (2) prevent what has not yet occurred."

"Consider putting encryption on all critical control links (e.g., in the power system, the FAA, . . .)."

Step 2. Observations and Findings

Step 2 of the scenario involved the editing and development of a memorandum to the Secretary of Defense regarding steps that could be initiated "today" to reduce U.S. vulnerability to cyberspace-based attacks in the future. Some of the observations of Step 1, above, were reiterated. Perhaps the most interesting new observation dealt with analogies the U.S. government might consider in considering its posture and relationship with industry in working toward better cyberspace security. Three specific analogies were mentioned:

Automobile Safety Regulations. The U.S. government, in cooperation with the auto industry, created regulations that raised the safety level of automobiles. These regulations also raised awareness of safety issues within the U.S. populace in general. The safety and se-

curity of cyberspace is now in a situation analogous to that of the automobile industry many years ago. With appropriate regulations, the market could be influenced in a substantial way. This is important because market forces will ultimately have the major influence on the safety and security of U.S. information systems.

The U.S. Centers for Disease Control (CDC). The CDC acts as a worldwide clearinghouse for health and disease information; it is a central source for information when needed, from routine queries to tracking the spread of epidemics. This same clearinghouse function is needed to collect and assess information on disparate cyberspace security incidents.

Underwriters' Laboratory. It may be possible to create an institution for the testing and evaluation of the security provisions of telecommunications and other infrastructure software and systems. Perhaps, eventually, systems that don't have this "seal of approval" would not be allowed to interconnect to the infrastructure. It is an open question, however, if the safety and security of complex operating systems and application programs comprising millions of lines of source code could in fact be so tested. The evolution of software systems (multiple versions and releases, new system components, etc.) may be too rapid for this task to be accomplished in reasonable time or at reasonable expense.

R&D Investment Suggestions

We believe the following are the most important specific research and development suggestions made during the course of Step 2 deliberations.

Study "Distributable Secure Adaptable Architectures." The group that coined the phrase "distributable secure adaptable architectures" believed each word in the phrase was important. Although much research has been done on secure operating systems for individual computers or workstations, new advances are needed for systems that are inherently distributable (over telecommunication links and networks, over geographic distances, among disparate groups). These systems should be secure and adaptable, because rigid system solutions are bypassed or trashed as the environment in which they must work evolves. They must be architectural, dealing with all sys-

tem levels, rather than "silver bullets" meant to solve narrow specific problems. This topic was meant as a theme for a research program, not just an individual project.

Study "Rapid Recovery" Strategies and Systems. Participants despaired of the design and implementation of verifiably secure information systems throughout the nation's infrastructure—at least in their lifetimes. But perhaps even near-absolute security would be much less necessary if systems were designed for rapid recovery. If any link or node might be disabled by a perpetrator, but could be restored in milliseconds, or at most seconds or minutes, and if the system in addition had considerable redundancy—then perhaps that would suffice for most systems and applications. What portions of the infrastructure might be amenable to such a solution? How might systems be designed with rapid recovery from malevolent (or inadvertent) acts as a design criterion?

Study "Understanding and Managing Complex Systems." The information systems controlling our national infrastructure are some of the most complex systems ever designed. They have millions of interacting components. Often, each node is controlled by millions of lines of code. We need a better science of complex systems, or at least tools for helping to understand their dynamic operation and vagaries. Among the tools that were suggested at the exercise were:

- *Data probes and selective sampling* as a means of ascertaining the health and vitality of a system during its operation;
- *Intelligent modeling tools* for representing such complexity at various levels of abstraction;
- *Tools for the visualization of information flows.* With proper visualization could abnormal patterns of activity be detected before they became destructive?
- *Interactive and multiple-scale global analysis.* How can analysis be conducted at various levels of the system, interactively during system operation?

Study the Design of Processes for Developing Secure Software Systems. Through the efforts of the Software Engineering Institute, among others, a "science" of software engineering is slowly emerging. They are developing standards for assessing the level of maturity

of software development groups. We need comparable processes and an engineering discipline devoted to the design and implementation of secure information systems. Such processes must include a variety of procedures to ensure the validity of the compiler being used and protect access to it, which may require a "two-man rule" for making critical system changes (see "Operational aspects...", above), and numerous other procedural and technical safeguards. An entire science and discipline of secure system development is needed.

Study the Concept of a Minimal Essential Information Infrastructure (MEII). The scenario materials given to the participants presented for their consideration the concept of a Minimal Essential Information Infrastructure. Groups generally supported exploration of the idea, and encouraged study of

- *the essential services it must protect and carry.* How many are there? What are their information demands?
- *the functionality that must be guaranteed.* Participants stressed attention to functionality, rather than becoming absorbed in the "nuts and bolts" of specific hardware and system components.
- *the appropriate telecommunications architecture.* Do existing telecommunication systems provide the appropriate redundancy and architecture, or are alternative designs needed?
- *a global management structure.* We come back to the question: Who's in charge? Is an MEII managed in a decentralized manner, or centrally? What regulations and guidelines govern its use?
- *prototyping and exercising the system*. It was widely understood that an MEII could not be created and "put on the shelf" for use in emergencies only. The information environment is much too dynamic for such a warehoused system to remain viable. It must be used regularly to remain relevant.

Some felt that encouraging diversity in infrastructure systems (of both paths and system architecture) was more important than attempting to design or develop an MEII. Others stated that "DoD, for cost reasons, will have to fall back on a reduced functionality system like MilStar, rather than attempting to secure, or duplicate, portions of the nation's existing telecommunications system." It was unclear,

however, whether such satellite links could be extended to cover the communications required by non-Defense portions of critical national infrastructures.

Study the Minimum Essential Functionality for Various Segments of Our Society. This question is related to the previous topic. Research should be undertaken to ascertain the minimum amount of information infrastructure that would sustain our society for limited periods of time. If the energy system could only provide half the normal power, would that suffice for a week? Would 2/3 of banking systems suffice; if so, for how long? If 1/4 the air traffic control systems were inoperable for 48 hours, could air transportation continue, and if so with what throughput compared to normal? Such a study would allow estimates to be generated of the minimum essential communication capacity that would be needed in an emergency, as a function of time. These estimates would in turn inform the studies of an MEII (see above).

Study the Analogy of "Biological Diversity" for Complex Information Systems. Considerable concern was expressed at the exercise about the limited diversity in our key infrastructure systems. Most telephone switches are made by one of only a few companies (e.g., Nortel, Siemens, AT&T), and these switches are almost exclusively based on the Unix or VMS operating systems. Most Internet nodes run common versions of the Unix operating system. The telephone signaling system uses the Internet's SMTP message transfer protocol. And so on. Once perpetrators discover a flaw in such systems, that flaw can be quickly exploited in thousands of copies of that system component. Biologists have long extolled the virtues of biological diversity, so that crops such as corn, wheat, etc. are not genetically identical and subject to the same diseases or infestations. In the same way, government may be called upon to mandate that sufficient dissimilarity be engineered into critical systems. Without such intervention, the market is tending toward uniformity in system components to achieve savings from mass production, replication, training, and documentation.

Consider the Biological Immune System Metaphor for Software. The Step 2 draft memo handed to group discussants mentioned as a possible research idea the concept of modeling system defenses on

the tactics used by the human immune system to discover and immobilize "intruders." As described in Hundley and Anderson (1995):

> The biological agents providing the active defense portion of the immune system employ certain critical capabilities: the ability to distinguish "self" from "nonself"; the ability to create and transmit recognition templates and killer mechanisms throughout the organism; and the ability to evolve defenses as the "threat" changes.
>
> Software agents providing a cyberspace active defense analogue to these biological antibodies would need the same capabilities.
>
> The message of this metaphor is clear: Cyberspace security would be enhanced by active defenses capable of evolving over time.

Some existing research is under way based on this metaphor, for example, see Forrest et al. (1994) and Kephart (1994). Discussants at the exercise were intrigued by the concept and recommended further exploration of its possibilities.

Study "Dynamic Diversity" in Infrastructure Information Systems. A security problem with existing infrastructure systems is their stability and consistency. Once a flaw is discovered, it can be exploited for months and on multiple instances of that system throughout the country. Groups talked about the possibility of dynamic diversity, wherein software at all levels of these systems modified itself frequently in a way that didn't affect functionality, but that could foil attempts to exploit known security flaws. Perhaps if file names changed, the location of software modules moved, alternate protocols were used, and so on, it would preclude broad attacks on multiple identical system components. Is such dynamic diversity possible, while retaining the ability to perform maintenance, upgrades, training, and other activities that depend on stability in systems? The related topic of a system performing dynamic self-configuring around corrupted elements was also mentioned; this is another biologically related metaphor that recurred in group discussions.

Replace Software with Firmware? Software is modifiable. Firmware (instructions burned into read-only memory (ROM) or related memory devices) is much less so. Can software in critical systems be replaced by firmware so that it cannot be "hacked" by intruders? If so,

which systems are amenable to this approach? How would the security improvements of this approach weigh against the greater difficulty of upgrading and maintaining—e.g., by the changing of ROM chips rather than remotely downloading software—the instructions controlling system behavior?

Is It Possible to "Sterilize" Data Passing Through Our Telecommunications Systems? Billions of bits of data pass through our national information infrastructure each second. Some of those bits represent information about individual citizens' login and password combinations, social security and credit card numbers, account information, health status, and innumerable other sensitive information items. Our nation has superb communications monitoring tools, housed primarily in the National Security Agency. However, the NSA is precluded by law from collecting information about U.S. citizens. When incidents of "information warfare" are being waged against U.S. systems, could key data flows be "sterilized" or "sanitized" by computer hardware and/or software in such a manner that the NSA could help monitor and track perpetrators in cyberspace without violating these laws? This topic was raised during exercise discussions. We have not studied all the relevant laws and regulations to assess whether such sterilization measures would allow the power of NSA's analyses to be brought to bear on telecommunications involving U.S. citizens, but perhaps the topic merits further investigation. If so, what kinds of pattern detection and replacement algorithms would suffice to accomplish this goal?

Study the Ability to Reengineer or Retrofit Legacy Information Systems to Enhance Their Security. There are thousands of existing information systems and components supporting the national information infrastructure, including individual PSTN switches, pipeline control systems, the air traffic control system, Internet routers, and so on. It is clearly not possible, in the next decade or two, to redesign and reprogram all these systems to enhance their security significantly. Is it possible, however, to retrofit these systems with special hardware/software devices for greater security? An analogy might be the "TCP Wrapper" technology pioneered by Wietse Venema[7] and others that is used as a software retrofit on a key Internet protocol. Are other security-enhancing "wrappers" possible in other circumstances? The entire topic of retrofitting existing sys-

tems could use substantial R&D if significant progress on infrastructure security is to be made on any reasonable time scale.

Sponsor Development of an Aircraft-Like "Black Box" Recording Device. When a cyberspace security incident happens, it is often not detected in real time, and the trail back to the perpetrator becomes lost. Could a "black box" recording device be developed, to be attached to key nodes or links of cyberspace systems, that would record every transaction passing through that node or link during the last n minutes (where n = 5 or 10, for example)? If so, that record would be invaluable in tracing the source of incidents, whether they are accidental or deliberately perpetrated. Thousands of such systems would be required to cover key links or nodes; could they be made robust, inexpensive, and ultra-reliable?

Sponsor Development of Devices That Would Record Tamper-Proof Audit Trails for Information Systems. This concept is related to the previous one. A variety of critical infrastructure systems retain some level of audit trail of system activity, to help in diagnosing problems. Many such audit trails are merely data recorded into a file for later analysis. If a perpetrator gains root access to a system, he or she can tamper with the audit trail to remove any indication of the perpetrator's presence and activities. How should systems create tamper-proof audit trails that can become accurate records of system activity? Since it is impossible for many systems to retain a record of *all* activity over lengthy periods of time, such tamper-proof audit trails may well need to be "FIFO queues" (first-in first-out), where the newest information recorded pushes out the oldest information because of limited recording space.

Develop Software That Can Perform Real-Time Pattern Detection As an Aid to Attack Assessment. Systems are currently under development, and being fielded, that monitor for suspicious or abnormal activity in real time during a system's operation. Examples include SRI's Next Generation Intrusion Detection Expert System (NIDES)[8] and work at the Air Force Information Warfare Center. Research should be conducted to evolve the capabilities of such real-time pattern detection systems, since they form a vital component of any information security program. Participants mentioned that neural nets are one appropriate technology to be considered, since they can be self-adapting as patterns of system activity change. We are aware

that some existing systems already incorporate both neural-net and rule-based components. These use biological metaphors analogous to those we discussed earlier.

REFERENCES

Anderson, D., T. Fribold, and A. Valdes (1995). *Next Generation Intrusion Detection Expert System (NIDES): A Summary,* SRI-CSL-95-07. Menlo Park, CA: SRI International.

Anderson, D., T. F. Lunt, H. Javitz, A. Tamaru, and A. Valdes (1995). *Detecting Unusual Program Behavior Using the Statistical Component of the Next Generation Intrusion Detection Expert System (NIDES),* SRI-CSL-95-06. Menlo Park, CA: SRI International.

Forrest, S., A. S. Perelson, L. Allen, and R. Cherukuri (1994). "Self-nonself discrimination in a computer," in *Proc. 1994 IEEE Symposium on Research in Security and Privacy.*

Hundley, R., and R. Anderson (1995). "Emerging Challenge: Security and Safety in Cyberspace," *IEEE Technology and Society Magazine,* Vol. 14, No. 4, Winter 1995-1996, pp. 19-28. Reprinted in RAND RP-484.

Kephart, J. O. (1994). "A Biologically Inspired Immune System for Computers," in R. A. Brooks and P. Maes (eds.), *Artificial Life IV, Proceedings of the Fourth International Workshop on Synthesis and Simulation of Living Systems.* Cambridge, MA: MIT Press, pp. 130-139.

Mesic, R., R. Molander, and P. Wilson (1995). *Strategic Futures: Evolving Missions for Traditional Strategic Delivery Vehicles,* RAND, MR-375-DAG.

Millot, D., R. Molander, and P. Wilson (1993). *The Day After... Study: Nuclear Proliferation in the Post-Cold War World,* Vols. I–III. RAND, MR-266-AF, MR-253-AF, MR-267-AF.

Molander, R., A. Riddile, and P. Wilson (1995). "Nuclear Command, Control, Communications and Intelligence Review Adjunct," RAND, internal paper.

Molander, R., A. Riddile, and P. Wilson (1996). *Strategic Information Warfare: A New Face of War,* RAND, MR-661-OSD.

Molander, R., P. Wilson, R. Mesic, and S. Gardiner (1994). *Under the Nuclear Shadow: Power Projection in the Post-Cold War World,* RAND, MR-513-AF.

Venema, W. (1992). "TCP Wrapper: Network Monitoring, Access Control, and Booby Traps," in *Proceedings of the 3rd Unix Security Symposium,* Baltimore, MD, September 1992. Also available via Web site ftp://ftp.win.tue.nl/pub/security/index.html.

NOTES

[1]From the Project Description, August 25, 1995. At the time of its writing, DARPA was referred to as ARPA. In this report, when quoting original materials we use the terminology of those materials.

[2]See the research reports cited in the first paragraph of this section for descriptions of previous exercises using this three-step exercise methodology.

[3]See Molander, Riddile and Wilson (1996).

[4]From the Step 1 scenario instructions.

[5]From the Step 2 scenario instructions.

[6]Tactical warning provides information about an attack in progress; attack assessment determines the extent and characteristics of an attack, including information on targets, consequences, and perpetrators.

[7]See Venema (1992).

[8]Anderson, Fribold and Valdes (1995); Anderson, Lunt, Javitz, Tamaru and Valdes (1995).