CRITICAL INFRASTRUCTURES PROTECTION: RESEARCH AGENDAS FOR INFORMATION SYSTEMS SECURITY

Stuart W. Katzke, Arthur E. Oldehoeft, and Shirley Radack
National Institute of Standards and Technology
Gaithersburg, Maryland 20899-8930

Abstract

Several national studies have examined the vulnerabilities and threats to the critical infrastructures upon which the U.S. depends for its national defense and economic growth, and have addressed measures needed to protect the critical infrastructures. These include systems for energy, banking and finance, transportation, human services, and telecommunications. Telecommunications and information systems interconnect the critical infrastructures, making them interdependent on one another and vulnerable to new threats, both domestic and international. When the interconnecting telecommunications and information systems fail to perform properly, the critical infrastructures are at risk. This paper summarizes the findings and recommendations related to computer and information security developed by the President's Commission on Critical Infrastructure Protection, the National Science and Technology Council, and the Computer Science and Technology Board of the National Research Council. Details of the research agendas developed by these three groups are presented for those aspects dealing with computer and information security. This paper does not attempt to synthesize a single agenda that bridges the substantial differences in philosophical approaches, emphases, and itemization of proposed research topics, but provides summaries that may help researchers and system developers who wish to focus their attention on critical infrastructure protection problems.

CRITICAL INFRASTRUCTURES PROTECTION: RESEARCH AGENDAS FOR INFORMATION SYSTEMS SECURITY

I. Introduction

The U.S. Government has launched new efforts to raise awareness about the vulnerabilities of the critical infrastructures upon which the U.S. depends for its national defense and economic growth, and to find new ways to protect the infrastructures and their supporting systems. The critical infrastructures include systems for energy, banking and finance, transportation, human services, and telecommunications. Telecommunications and information systems interconnect the critical infrastructures, making them interdependent on one another and vulnerable to new threats, both domestic and international. When the interconnecting telecommunications and information systems fail to perform properly, the critical infrastructures are at risk.

Several national studies of the vulnerabilities and threats to the critical infrastructures have been issued recently. This paper, which focuses on the information systems security issues that have an impact on critical infrastructure protection, summarizes the findings and recommendations of the President's Commission on Critical Infrastructure Protection, the National Science and Technology Council, and the Computer Science and Technology Board of the National Research Council.

Each of these groups has published proposed research agendas that address the need to protect the national critical infrastructures. In each study, a concept of protection or trust or high confidence encompasses broad issues of safety, reliability, and security as applied to computers, networks, and information. Each proposed research agenda calls for required advances in theory, design, and engineering. While recognizing that a holistic approach is required, this paper attempts to extract from each of these three agendas only those aspects that are particularly relevant to computer and information security. Because of substantial differences in philosophical approaches, emphases, and itemization of research topics, this paper does not attempt to coalesce the three agendas into a single agenda. The resulting three summaries may serve as useful guidelines for researchers and system developers who wish to focus their attention on problems in this somewhat more restricted subspace.

Other groups have completed related reports that concentrate on the policy issues affecting the critical infrastructures. The Center for Strategic and International Studies (CSIS) is a public policy research institution that conducts analyses and policy studies in areas such as international finance, U.S. domestic and economic policy, U.S. foreign policy, and national security issues. A CSIS study [1], entitled "Cybercrime, Cyberterrorism, Cyberwarfare," discusses the threats that strategic information warfare (SIW) poses to the critical infrastructures, and recommends that the U.S. adopt policies to protect against SIW attacks and make strategic information dominance a national security objective. Federal, state, and local governments are encouraged to take steps to assure that essential government services will be maintained and to work with the private sector in achieving secure systems.

The President's National Security Telecommunications Advisory Committee (NSTAC) has examined information assurance and information infrastructure protection issues. NSTAC advises the President on national security and emergency preparedness matters related to telecommunications and information systems. The NSTAC report [2] on critical infrastructures discussed recommendations for the development of policies, procedures, techniques, and tools to facilitate joint industry-government cooperation on cyber security, particularly for national security and law enforcement applications.

II. Critical Infrastructures Studies

A. President's Commission on Critical Infrastructure Protection (PCCIP)

The President's Commission on Critical Infrastructure Protection, established in July 1996 by Executive Order 13010, was asked to review the vulnerabilities and threats to U.S. critical infrastructures, and to develop a comprehensive national strategy for protecting the infrastructures. Composed of both private and public sector representatives, the PCCIP consulted with experts, identified physical and cyber threats, assessed the risks, developed policy recommendations, and issued its report, Critical Foundations, Protecting American's Infrastructures, in October 1997 [3].

Critical infrastructures were identified as systems providing telecommunications, electric power, transportation, oil and gas delivery and storage, banking and finance, water, emergency services and government services. The Commission recommended that the government increase its investment in research and development for infrastructure assurance to about $1 billion over a five year period, and indicated that a similar level of commitment would be needed from the private sector. This increase in private sector spending could occur as market demand for infrastructure assurance technology increases. The Commission called for close coordination and partnership among government, industry and academia to conduct a successful R&D effort, and proposed that a National Infrastructure Assurance Office be established to coordinate and oversee the R&D agenda. The Commission also recommended that the National Research Council define more fully a national infrastructure assurance research program based on the information in the Commission's report.

The Commission identified six areas for research and development:

- Information Assurance to protect the increasingly interconnected and complex communications infrastructure and the information created, stored, processed, and transmitted on it. New affordable means of protection are needed because of the increasing rate of incidents, new vulnerabilities, and the inadequacy of current solutions.

- Monitoring and Threat Detection to provide reliable, automated monitoring and detection systems, timely and effective information collection technologies, and efficient data reduction and analysis tools. These technologies are needed to identify and characterize attacks against systems and to support early warning systems.

- Vulnerability Assessment and Systems Analysis to provide advanced methods and tools to advance physical and cyber security in an integrated fashion. These methods are needed to identify critical nodes within infrastructures, to examine infrastructure interdependencies, and to help understand the behavior of complex systems. Modeling and simulation tools and test beds for studying infrastructure-related problems are needed for experimentation that cannot be performed on actual operational systems.

- Risk Management and Decision Support to help government and private sector decision-makers prioritize the use of finite resources to reduce risk. Methodologies and tools are needed to address risk from familiar threats, such as natural disasters and physical attacks, as well as emerging and future threats that may arise from the increasing interdependence and reliance on cyber systems.

- Protection and Mitigation Technologies for real-time system control, infrastructure hardening, and containment and isolation techniques to protect infrastructure systems against the entire threat spectrum. Advanced survivability, reliability, and assurance enhancement measures are needed.

- Incident Response and Recovery Technologies and Tools for planning for, responding to, and recovering from incidents, such as natural disasters and physical and cyber-based attacks that affect local or national infrastructures.

1. Presidential Decision Directive (PDD) 63

Presidential Decision Directive (PDD) 63, Protecting America's Critical Infrastructures [4], was issued in May 1998. The PDD builds on the recommendations of the President's Commission on Critical Infrastructure Protection. PDD-63 sets a goal of a reliable, interconnected, and secure information system infrastructure by the year 2003. It also calls for increased security for government systems by the year 2000, establishes a national center to warn of and respond to attacks, and advocates building the capability to protect critical infrastructures from intentional acts by 2003. Each federal government department and agency is expected to reduce its exposure to new threats and to serve as a model for national infrastructure protection. These goals are to be attained through the voluntary participation of private industry.

The PDD established new organizations to deal with the critical infrastructure, foreign terrorism and threats of domestic mass destruction (including biological weapons).

- The National Infrastructure Protection Center (NIPC) at the Federal Bureau of Investigation (FBI) will include representatives from the FBI, Departments of Defense, Energy, and Transportation, the U.S. Secret Service, the Intelligence Community, and the private sector, and will encourage information sharing among agencies in collaboration with the private sector. The NIPC will also provide the principal means of facilitating and coordinating the federal government's response to attacks on the infrastructures.

- The federal government will cooperate with the private sector in setting up an Information Sharing and Analysis Center (ISAC).

- A National Infrastructure Assurance Council drawn from private sector leaders and state/local officials will provide guidance to the formulation of a national plan to protect the critical infrastructures.

- The Critical Infrastructure Assurance Office will support the development of the national plan, will coordinate a national education and awareness program, and will conduct legislative and public affairs.

2. Critical Infrastructure Assurance Office

The Critical Infrastructure Assurance Office (CIAO), stipulated in PDD-63, has been established. The CIAO is responsible for integrating the sector plans that have been developed for the various critical infrastructures into a National Infrastructure Assurance Plan. The CIAO will coordinate analyses of the federal government's dependencies on critical infrastructures, conduct a national education and awareness program, and carry out legislative and public affairs activities. The Department of Commerce was designated the executive agent for the CIAO. Federal government organizations are encouraged to find new ways to work with industry in solving the critical infrastructures problems.

The CIAO conducted a four month study with the Transition Office of the President's Commission on Critical Infrastructure Protection to identify preliminary R&D topics and associated roadmapping information for the eight critical national infrastructures. The CIAO report, Preliminary Research and

Development Roadmap for Protecting and Assuring Critical National Infrastructures [5], focuses on perceived technology shortfalls, or "gaps," between infrastructure assurance technology needs and available technologies.

B. National Science and Technology Council

The National Science and Technology Council was established in 1993 to advise the President on science, space, and technology issues, and to coordinate the different parts of the federal government research and development activities. The Council establishes national goals for federal science and technology investments, and prepares research and development strategies

The National Science and Technology Council's High Confidence Systems Working Group of the Subcommittee on Computer, Information, and Communications Research and Development recently completed a report, A National Research Agenda for High Confidence Systems [6]. The Working Group describes a high confidence system as one in which the consequences of its behavior are well understood and predictable. It must withstand internal and external threats and must deal with naturally occurring hazards and well as malicious attacks from a sophisticated and well-funded adversary.

Working within the framework of the need for protection of critical infrastructures, the National Science and Technology Council published an agenda for research that focuses on critical information technologies necessary to achieve predictably high levels of system safety, security, reliability, and survivability. High confidence systems support transportation, health care, electric power generation, manufacturing, oil and gas production, chemical production, and financial services, as well as law enforcement, emergency services, and national defense.

Achieving high confidence is becoming more difficult as systems become more complex, according to the Council. Increased integration, continuous evolution, and larger scale systems are producing more complexity. New analysis techniques are needed for these more complex systems. The need is especially urgent since the U.S. is growing increasingly dependent on computing in many industries and in government. Information and communications systems are regularly subject to new and malicious attacks. As a result, new and advanced techniques are needed to protect against internal and external threats.

The Council envisions that its proposed research agenda will be conducted by the federal government agencies that are the participants in the high confidence systems initiative. The work will be coordinated by a High Confidence Systems Working Group, and will leverage the efforts of the participating organizations through collaborative research. Further, cooperative efforts with the private sector will be sought.

C. National Research Council

The National Research Council is an operating agency of the National Academy of Sciences and the National Academy of Engineering. It was organized in 1916 to associate the broad community of science and technology with the Academy's purposes of further knowledge and advising the federal government. The NRC provides services to the government, the public, and the scientific and engineering communities.

The PCCIP recommended that the National Research Council define a research program based on the PCCIP report on the vulnerability of critical systems to attack and their susceptibility to disaster. The Committee on Information Systems Trustworthiness was convened by the Computer Science and Telecommunications Board of the NRC to assess the nature of information systems trustworthiness and the prospects for technology to increase it. The Committee's study, Trust in Cyberspace [7], was developed at the request of the Defense Advanced Research Projects Agency (DARPA) and the National Security

Agency (NSA). The study addresses how the trustworthiness of networked information systems can be enhanced through improved computing and communications technology. A detailed research agenda was developed for improving information system trustworthiness.

III. National Research Agendas

This section presents a summaries of three proposed national research agendas in the broad area of computer and information security. The general philosophy of each of the three agendas is described. Specific itemization of some of the important recommended research programs and initiatives is included in Tables I, II, and III.

A. President's Commission on Critical Infrastructure Protection and Critical Infrastructure Assurance Office

In response to PDD-63, three major reports have been prepared. The first report [3], from the President's Commission on Critical Infrastructure Protection, consists of a study of the critical infrastructures that constitute the life support systems of the nation, a determination of their vulnerabilities, and proposals of strategies for future protection. The critical infrastructures addressed in this report are transportation, oil and gas production and storage, water supply, emergency services, government services, banking and finance, electrical power, and information and communications. The second report [5], from the Transition Office of the President's Commission on Critical Infrastructure Protection, builds on the first report and provides a foundation for steps to develop a national infrastructure assurance R&D Program. Preliminary R&D topics, along with roadmapping information, are provided for the national critical infrastructures. The study team identified more than 70 R&D topics for three timeframes (near-term [before 2002], before 2005, before 2010), and estimated that approximately $2 billion each would be needed for the near-term and mid-term requirements and that approximately $3 billion would be needed for the long-term requirements.

In a third and more recent report [8] to the President, the Critical Infrastructure Assurance Office is proposing the creation of a National Information Systems Defense Program. The details of this program are currently under review by the Critical Infrastructure Coordination Group and, as a consequence, its contents are subject to change.

The R&D topics correlate with the objectives of infrastructure assurance -- to reduce critical vulnerabilities by protecting infrastructures, detecting intrusions, mitigating the effects of disruptions, assisting in the management of incidents, and facilitating recovery. Other themes focus on developing analytical or supporting technologies to help meet those objectives. Vulnerability assessment provides supporting baseline information for all of the other themes. Elements within each theme apply to various infrastructures, making it possible to leverage expenditures to cover multiple topics.

1. National Information Systems Defense Program

The proposed Computer Applied Research Initiative (CARI) component of the overall National Information Systems Defense Program is summarized here in an effort to contrast it with other national research agendas. The proposed program consists of three component plans:

- the Critical Infrastructure Protection Plan for sectors of the economy,
- the Federal Information Assurance Plan for civilian agencies and departments of the Federal government, and
- the Defense National Information Security Plan for the Defense Departments and national security agencies.

Each plan has three goals, which are supported by eighteen programs:

- Goal 1: prepare and prevent – assess and eliminate significant vulnerabilities to information warfare,
- Goal 2: detect and respond – assess, warn, isolate, respond and reconstitute essential information dependent components, and
- Goal 3: build strong foundations for secure cyber-systems – public-private partnership, sound legal footing, widespread public understanding of importance and need, and international cooperation.

Five new concepts are introduced:

- creation of national intrusion detection networks,
- a Cybercorp program for college-level training in computer science,
- a computer applied research initiative (CARI) – Federally sponsored research in information assurance and defense practices,
- an information assurance/security foundation – Federally-funded national program of awareness and education on information and security practices, and
- a (ReconNet) Federal disaster recovery plan for reconstitution of computing systems and networks in event of a widespread outage.

Table I provides some details of recommended research projects in communications and computer security for Goals 1 and 2.

B. National Science and Technology Council

The National Science and Technology Council's research agenda focuses on activities that will protect the public by creating confidence in the safety, reliability, trustworthiness, security, timeliness, and survivability of critical systems. Other activities will protect the consumer by fostering higher reliability, safety and ease of use of commercial products, and will promote improved government services and national security.

A multi-agency research agenda is intended to provide the requisite planning to support a new Federal government R&D initiative in high confidence systems (HCS). Five technology goals are listed:

- provide a sound, theoretical and technological basis for assured construction of safe, secure systems,
- develop hardware, software and system engineering tools that incorporate ubiquitous, application-based, domain-based, and risk-based assurance,
- reduce the effort, time and cost of assurance and quality certification programs,
- provide a technological base of public domain, advanced-prototype implementations of high-confidence technologies to enable rapid adoption, and
- provide measures of results.

The specific recommendations for research to advance HCS are presented in Table II.

C. National Research Council

The NRC's agenda for research provides a science base and engineering expertise for building trustworthy information systems. The area of trustworthiness is identified as having numerous dimensions (e.g., correctness, security, reliability, safety, and survivability). In forming their recommendations, the NRC committee tracked the progress of the critical infrastructure protection efforts. Their report proposes, within the context of networked information systems (NISs), a detailed agenda for long-term research and the promotion of fundamental or revolutionary, rather than incremental advances. The agenda, targeted for federal research funding organizations such as DARPA and NSA, identifies specific science and technology advances that could potentially play a significant role.

For the single dimension of computer and communications security, the report recommends a new approach, especially for NISs where foreign and mobile code are commonplace. The report acknowledges that insecurity exists, that insecurity cannot be destroyed, and that insecurity can move around. Therefore, the recommendations for future research are not based on the concept of absolute security, but rather on techniques for identifying vulnerabilities and, in the light of anticipated threats, making design changes to reposition the vulnerabilities to counter the consequences of attacks. Needed research activities include:

- Better understanding of the design and engineering practices that foster trustworthiness. Needed are specific guidance to designers, implementers and managers, and protective measures for the public telephone network and the Internet.
- Research in techniques for composing subsystems in ways that contribute to trustworthiness. An understanding is needed of how subsystems interact with each other and with the other elements of a larger system, and of the relationship between commercial-off-the-shelf components and system trustworthiness. Research is needed for improving the integration of testing and formal methods.
- Research in techniques for identifying vulnerabilities and making design changes to reposition those vulnerabilities in light of anticipated threats. Better cryptographic protocols and faster encryption and authentication algorithms are needed to keep pace with increasing speeds of communications. Other security research needs are for application-layer firewalls, operating system support for access controls, and ways to defend against denial of service attacks.

- Ways to enhance the trustworthiness of untrustworthy components. Research is needed to determine where to place trustworthiness functionality with a system, and in the application of monitoring and detection practices.

Specific research recommendations for computer and information security are included in Table III.

IV. Conclusion

The critical infrastructures in the U.S. including systems for communications, finance, energy distribution and transportation are vulnerable to malicious attacks which could have devastating impacts on the U.S. national security and economy. Public awareness of the vulnerabilities and their potential consequences has been raised by the report of the President's Commission on Critical Infrastructures Protection. Under the PCCIP, a research agenda for the National Information Systems Defense Program was developed. The National Science and Technology Council developed an agenda to protect the U.S. government's high confidence systems, which are essential to protection of the public and the consumer and to improving national security and public services. The National Research Council identified specific R&D areas to protect the critical infrastructures, focusing specifically on their increasing dependence on networked information systems and on complex subsystems, which are often impaired by defects.

The studies agree that research and development investment is needed to promote safe, reliable, dependable, secure and survivable information systems for both public and private sectors. There is also agreement that government, industry, and research organizations must work together to achieve affordable and effective commercial, off-the-shelf products. Action over the next few years will be critical to accomplishing the needed research and improving the security of the critical infrastructures.

This paper summarizes the proposed research agendas for the areas of computer and information security as extracted from these studies in the broader areas of protection, trust, and high confidence. While there is some recognizable overlap in the three agendas, there are substantial differences in philosophical approaches, emphases, and itemization of specific topics. With varying degrees of applicability, each is singularly helpful to researchers and developers who focus their attention in security areas. A seemingly useful future study would be a synthesis of the three agendas into one common agenda.

References

1. Cybercrime...Cyberterrorism...Cyberwarfare..., Averting An Electronic Waterloo, A Report of the CSIS Global Organized Crime Project, Center for Strategic and International Studies, 1998.

2. Information Infrastructure Group Report, The President's National Security Telecommunications Advisory Committee, September 1998.

3. Critical Foundations – Protecting America's Infrastructures, The Report of the President's Commission on Critical Infrastructure Protection, October 1997.

4. Presidential Decision Directive 63, Protecting America's Critical Infrastructures, May 1998.

5. Preliminary Research and Development Roadmap for Protecting and Assuring Critical National Infrastructures, Transition Office of the President's Commission on Critical Infrastructure Protection, July 1998.

6. A National Research Agenda for High Confidence Systems, Subcommittee on Computing, Information, and Communications Research and Development, National Science and Technology Council, December 14, 1998.

7. Trust in Cyberspace; Committee on Information Systems Trustworthiness; Computer Science and Telecommunications Board; Commission on Physical Science, Mathematics, and Applications, National Research Council, Prepublication Copy, September 29, 1998.

8. The National Information Systems Defense Program, Critical Information Assurance Office, (Draft) November 22, 1998.

Authors

Stuart W. Katzke directed computer security related programs at the National Institute for Standards and Technology (NIST) from 1975-99, serving as Chief of the Computer Security Division for NIST's Information Technology Laboratory (ITL) from 1988-99. He is currently a Chief Scientist in the Evaluated Information Assurance Solutions Group at the National Security Agency. Arthur E. Oldehoeft is a Professor of Computer Science at Iowa State University, presently working in NIST's Computer Security Division under a U.S. Federal program established by the Intergovernmental Personnel Act. Shirley Radack is presently a technical consultant to the ITL.

Table I. Computer Applied Research Initiative (CARI) Program for Goals 1 and 2
National Information Systems Defense Program

Goal l, Prepare and Prevent. Research in the following areas will be needed to provide the U.S. information infrastructure with a capability to withstand hostile attacks, to characterize and locate attacks, and anticipate possible attacks through a vulnerability assessment of existing systems.

Fielding an Enhanced Vulnerability Detection, Assessment and Analysis Program. The objectives are to identify, collect, organize, and disseminate vulnerability information, and to develop technologies and methods to deal with vulnerabilities during product development and system integration. Anticipated results and products include: a threat/vulnerability lexicon and vulnerability/attack classification taxonomies; a database system to identify, collect, organize, and disseminate vulnerability information; enhanced abilities to predict and test new vulnerabilities; and development of technologies, methodologies, and automated tools to avoid, reduce, or eliminate vulnerabilities during product development and system integration phases

Development of Advanced Information Assurance Tools. The objectives are to develop tools and techniques for rigorous design, implementation, testing, and formal verification of components and their integration into systems. Products include new technologies, management procedures, security protocols, and advances in fundamental theory.

Development of Advanced Security Architectures. The objectives are to organize security components and services to provide confidentiality, integrity and availability for information and communication systems. This includes advances in public key infrastructures, directory and certificate management, security component interoperability, security policies for emerging technologies, advancements in firewall and packet-switching technologies for active dynamic networks, automated distribution of patches and security upgrade information, and scalability and optimization of security architectures. The resulting products would also be applicable for retrofitting legacy systems.

Development of Advanced Modeling and Simulation Tools. The objectives are to develop representative models and simulation tools necessary to create and evaluate the technologies required to protect the infrastructure. They will be used to assess risk, security, interoperability, and recovery issues, provide synthetic test beds for experimental studies, and would be applicable to both emergent technologies and retrofitting legacy systems.

Development of Advanced Tools for Risk Management, Performance Assessment, Security Testing, and Metrics. The objectives are to develop new metrics and measurement tools for real-time performance to assist in detection of degradations in advance of serious impact. There is also a focus on the development of measurement techniques to apply advanced security to manufacturing and building of supervisory systems. Anticipated products include technical specifications and test results for performance measurement algorithms for use in infrastructure protection applications, prototype test instruments, a national network performance monitoring center and supporting prototype instruments, and specification of testing and monitoring procedures, measurements, testing methodologies, and standards.

Development of Advanced Secure Supervisory Control and Data Acquisition (SCADA) Systems. The focus of research is on security issues and vulnerabilities of SCADA systems. The goals include an analysis of the vulnerabilities, determination of best practices and improved security features and protocols, and the development of new architectures to increase redundancy and reliability.

Development of Advanced Artificial Intelligence Software Tools for Trapdoor Analysis. The objectives of the program are to develop advanced tools and techniques for detection and elimination of trapdoors (surreptitiously installed) in software. Such tools can be used to increase the integrity of software products and reduce the problem of future penetrations and compromises of computers and networks.

Development of Tools for Automated Distribution, Installation and Tracking of Software Patches. The objectives are to develop a set of software tools for automated distribution for software patches in computers and networks, track use of patches, and detect systems in which patches are not in use or not properly installed.

Understanding Human Factors in Information Assurance. The objectives are to address human factors relevant to information assurance and to develop procedures and strategies to reduce the risks to the associated infrastructure. Anticipated products include mitigation strategies, recommended best practices, and standards for personnel.

Goal II, Detect and Respond. The following research would support a system that will minimize the duration and extent of reduced infrastructure performance, isolate any damage, repair damage or allow human intervention to discard damaged systems, and keep relevant personnel informed as to the status of systems.

Development of Advanced Intrusion and Incident Detection and Warning Techniques. This research will involve the development of tools and procedures to detect, respond to, and recover from intrusions, incidents, and loss of services with a long-term goal of developing automated indications and warning systems. Metrics are needed for evaluating false alarm rates, strategy-based intrusion detection technologies, scalable intrusion detection systems, and tools to trace intrusions back to sources.

Characterization and Notification of Threats. The objectives of the research are collect data that will assist in the characterization of threats (motivation and origin) and to develop tools and technology that will profile attackers (motivation, origin, capabilities) and pinpoint origins of attacks (physical or cyber). Expected products include sharable data and information, and procedures and tools for collecting, analyzing, and disseminating information.

Fielding an Enhanced Response and Recovery System. The objectives are to develop the necessary tools and procedures (both manual and automated) for rejection, containment, and ejection of intruders, the mitigation of damage, and timely recovery and reconstitution of services. Products include response/recovery tools and procedures, hardware/software technologies, management procedures, and advances in fundamental theory.

Developing Advanced Reliability, Survivability, and Robustness Tools. The objectives are to develop technologies that will increase network reliability, system survivability, and robustness of the infrastructure and its systems and components. Methods of analysis and tools for enhancing reliability and restoration in the event of outages will be developed. Also to be developed are studies of the vulnerabilities of the nationwide spectrum dependent systems to determine identifying characteristics of radio spectrum attacks, and developing general mitigation guidelines and methods. The R&D security technology will specifically focus on intrusion prevention for the entire infrastructure including: cryptography, public key infrastructure, and advanced network protocols; system assurance methods; fault tolerance technology; domain name services protection; anomaly and suspicious event detection; automated response and reconfiguration mechanisms; new assurance methods such as model-checking for formal specifications; hazard analysis; test generation techniques; and techniques for combining dependable components.

Table II. A National Research Agenda for High-Confidence Systems
National Science and Technology Council

The agenda is dominated by the need for R&D to develop strong engineering theories, tools, and practices rather than R&D that are specific to secure computer and communications systems. Since the latter must necessarily depend on the former, most of recommended research areas seem applicable. The summary agenda is organized by major components, following that of the committee report. In a few cases, the report does single out specific aspects of computer security that are required to achieve overall high confidence. The research efforts of the agenda are organized under four major components.

HCS Foundations to Develop Supporting Theory and a Scientific Base

Theory. New theoretical work is needed in cryptography, safety, network management, operating systems, and assurance. In the area of computer and communications security, physical protection has evolved into the need for more flexible information protection domains that may span machines, networks, and national boundaries. The result is a patchwork approach in areas of cryptography, network management, public-key infrastructure (PKI), intrusion detection and response, and others. Fundamental new approaches are needed to deal with problems of wide-scale key distribution, certificate management, and interoperability. There is a need for improved mathematical underpinnings to provide more efficient, flexible, and resilient to attack operation of cryptographic systems in diverse environments. Revolutionary approaches, such as quantum cryptography and quantum computing technology, have the potential of providing theoretically sound security. New paradigms are needed in networked information systems (NISs) to provide transparent protection, adaptability to changing environments, and availability of critical services. Research is needed to monitor the health of local and global networks along with techniques to cope with stressful conditions. High-confidence capabilities in mobile and wireless systems require strong identification and authorization techniques to identify legitimate users and agents.

Specification. Efficient and effective methods are needed to decompose a computational demanding global property into local properties whose verification is computationally simple. Specification tools would automate the correlation of different levels of abstraction between systems, subsystems, and components. Executable specification languages would support tools for system visualization, design, analysis, implementation and validation. A greater range of domain-specific analysis would accommodate hardware and software design for resolving high-confidence issues.

Interoperability. A variety of formal methods are used for modeling and reasoning about complex systems in order to verify and validate system properties. Interoperable approaches for reasoning are needed that would consider multiple properties of a system and their interactions and provide an explanation that spans different methods.

Composition. Scientific foundations are needed for the safe, secure composition of components with unknown behaviors and the use of distributed-object systems in high-risk environments. These foundations are requisite for languages and engineering tools used in building systems that robustly manage interaction/interference and limit the propagation of failures.

HCS Tools and Techniques to Build Capabilities that can be used in the Application of Engineering Science to Design and Build Large-Scale Systems

Programming Languages, Tools and Environments. Technology is needed that embeds assurance into the software and system construction process through languages, tools and environments that include support for reasoning about functional and nonfunctional concepts relevant to the application domain. Research is needed to create an integrated combination of logically precise, automated mechanisms that focuses on eliminating the sources of error and early detection and correction of errors. An integration of assurance measures would lower barriers to use and reduce the tendency to defer or omit them.

Modeling and Simulation. Model construction is prerequisite to both formal and informal reasoning about systems. Research is needed in tools for describing and analyzing system and software behavior, and in building and using domain theories for domain-specific tools and languages. Automated abstraction is needed for feasible analyses, and in the automated generation of software from specification and engineering design models. Research is needed to examine how current methods for verification and analysis can supplement modeling and simulations efforts, and in the semantics of visualization of system states.

Robust System Design. The HCS research is intended to yield systems that achieve greater resilience under failure or adverse operating systems and provide higher performance and function under normal conditions. Techniques are needed to enable systematic design of failure or degraded modes of operation that can limit some aspects of function and performance, to analyze and limit fault propagation between connected components, and to create robust understandable designs for human operation.

Monitoring and Detection. Generally, parameters must be defined to indicate when a system reaches a point where high-confidence can no longer be guaranteed. Specifically, in the area of computer security, research has focused on traditional mechanisms for confidentiality, integrity, and availability. Currently, there are no known sets of parameters for security that can adequately describe a highly secure, available, and high-integrity system. In general, research is needed concerning sensors and probes that will collect, collaborate, and share information in determining when adverse conditions occur along with analytic methods for determining and quantifying fault coverage and performance of the systems. Visualization techniques are needed to aid in the human analysis of data.

Validation. New validation and verification techniques and tools are needed to systematically combine formal methods and language-based assurance, analysis, testing, and simulation. Model-based mathematical methods are needed to evaluate and quantify relevant system properties (stability, robustness, and performance) over a range of uncertainties, including the most likely severe adverse conditions.

Evidence and Metrics. Research is needed to establish principles that provide the foundations of measurement (metrology) for information technologies. New approaches are needed to evaluate the level of confidence in a system that employs HCS technologies. Measurement values may require structured (rather than single numeric) units to capture evidence from different assurance activities.

Process. Research is required to determine/identify high-confidence processes (involving people, platforms, tools, environments, and management) used in producing a system in order achieve consistent, repeatable development of high-confidence systems over a broad spectrum of applications.

HCS Engineering and Experimentation to Provide Reference Implementations, Scalable Proofs-of-Concepts, and Reusable Tools, Libraries, and Techniques

HCS Building Blocks. Rapid engineering of HCS requires a technology base of components with understood high-confidence properties and a methodology for assembling those components so that the

resulting system can be understood, reasoned about, and validated. High assurance needs to be incorporated into commercial-off-the-shelf (COTS) technologies and technologies must be developed to permit construction of HCS from components that are themselves not necessarily of high confidence.

Software Control of Physical Systems. For digital control systems, assurance is needed for software logic that manages the transition of among modes, coordinates interacting controllers, and interacts with operators. The research is intended to yield a high-confidence implementation technology and provide reference implementations of system building blocks.

Hardware and Software Platforms. The overall objective is to produce high-assurance hardware at lower costs. The goal of this research is to develop an HCS reference implementation for a hardware verification environment based on a loosely-integrated collection of tools such as cycle simulation, model checking, theorem proving, and other techniques that accept VHDL specifications as their input. This is intended to unify ongoing work and provide a common environment for the development of a methodology for formally verified designs.

High-Mobility Systems. Research is needed to apply HCS foundations, tools and techniques to the construction of high-mobility systems. Issues include real-time control over potentially interruptible connections, rapid and secure reconfiguration, and dealing with unreliability due to power constraints and environmental disturbances.

Public Key Infrastructure. This research will be to employ HCS foundations and HCS tools and techniques to construct a high-confidence distributed large-scale reference PKI. Testing and evaluating PKI components in a high-confidence distributed large-scale environment will be necessary. Attention must be given to the problems of integrating a complete set of PKI management components including security policy specification and approval, and certificate archiving, issuance, and revocation. Research is required to understand, specify, and test assurance properties of components that make up the cryptographic modules in the NIS and the integration of the corresponding key management systems. A better understanding of trust implications of certificate management is needed to provide assurance of the bindings of user identities and user public keys through a chain of trusted (possible cross-certified) third parties. Such bindings are also needed for some surrogate programs or processes known as agents, e.g., those that update and modify critical information in switching nodes and routers, and those that map domain names to real addresses.

HCS Demonstration and Pilots to apply the technology to real-world problems in various user-agency domains

Public Key Infrastructure for the Next Generation Internet (NGI). The confidence and trust in using the NGI will be directly related to the confidence and trust in the accompanying security services. The objective is to demonstrate this through the design, integration, and testing of a PKI in the NGI.

High-Confidence Systems for Free Flight. New technologies for improved situational awareness will be needed to increase air system safety. The present aircraft traffic control (ATC) system will be replaced by an autonomous system where each aircraft will report its position, velocity, and intention to all other aircraft in the proximity and to the ATC. The fused information, including data from other sensors, will serve as the basis for decisions to maintain clearances. In the new global air traffic management system, high-integrity communications, navigation, and surveillance subsystems will be required.

Table III. Trust in Cyberspace Research Agenda
National Research Council

While the numerous dimensions of trustworthiness are not independent of each other, this summary of the Council recommendations are presented primarily along the single dimension of security. The research agenda is divided into the six areas.

Access Control Policies

Many security policies of practical interest cannot be formulated as discretionary and mandatory access control policies and have other shortcomings because they cannot model the effects of certain malicious and erroneous software. Further, they do not fully address availability of resources and services and assume that objects have uniform security levels. Formal policy models have limited expressive power in that they do not suitably capture application-dependent control mechanisms (modern programs/subsystems have their own control mechanisms and presuppose that organizational policies are static with precise and succinct characterizations). Demonstrating the correspondence between a system and a formal model is not a practical approach to NIS security.

Fundamental research is needed to remove the limitations of the types of security that are captured by many formal policy models. Rather than a philosophy of "absolute security," the feasibility of an alternative philosophy needs to be investigated, that is, to identify insecurities and make design changes to reposition them in the light of known threats (i.e., move them to less exposed and less vulnerable parts of the system). Practical (cost-effective, time-efficient) means for evaluating the security characteristics (security features and residual vulnerabilities) for COTS system components are required.

Identification and Authentication Mechanisms

Network-Based Authentication. Authentication relies on the underlying network (and possibly host computers) to identify the source of traffic. Reliability is tied to that of the network. When implemented with moderate assurance (relying on the network provider as a third party), the principle of least privilege is violated. In general, network authentication is not amenable to high-assurance implementations.

Cryptographic Authentication. While more amenable to high-assurance applications, a compromise of encryption keys is a compromise of the authentication process. The design of key exchange protocols is a subtle business and flaws have been found in numerous protocols. As the sophistication of attackers increases, the need for authentication at the packet level becomes more critical. For deployment in large contexts, trusted third parties are needed, but this is a potential vulnerability. If implemented without the use of auxiliary storage (e.g., smart card) and if the encryption key is derived from a conventional password, then password guessing is a threat.

Token-Based Mechanisms. Hardware tokens have gained in popularity and are evolving into full-fledged, personal cryptographic devices, capable of providing services beyond authentication. Typically, a personal identification number (PIN) is required to enable a token. The degree of tamper resistance varies widely so their resistance to attacks is uneven.

Biometric Techniques. Relying on personal characteristics of users, the vulnerabilities of PINs and passwords are eliminated. Cost and availability are currently problems, so deployment is usually

limited to high-threat environments. Experience indicates that there is some reluctance to interact with certain types of biometric devices. When used for digital signatures, an interface is needed so that user will know what is being signed. There is a need for security in capturing biometric information in order to deal with the threat of capturing and replaying the resultant bit stream. Encryption of authentication information is still required. Compromise of authentication information can be permanent.

Further Research into Cryptographic Techniques and Supporting Tools. New protocols are needed for multicast communication authentication, but the technology for verifying protocols is far from mature. More research is needed to address interface commonality issues for hardware tokens. For use in closed NISs, existing or envisioned biometric interfaces in personal computers (e.g., microphones, cameras) should be explored.

Cryptography and Public Key Infrastructure

Cryptographic mechanisms can provide strong physical, personnel, and procedural security for geographically distributed, heterogeneously-administered NISs. But there is much to be learned about the practical aspects of deployment and use. In private-key methods, key sharing is vulnerability. In public-key methods, speed (cost) is an inhibiting factor, leading to the use of hybrid methods.

For key management in NISs, key distribution centers (KDCs) are needed for private key systems and certification authorities (CAs) are commonly used for public-key systems. User-centric models can be used for the latter, but they do not tend to scale well and certificates do not have a common meaning. In a KDC, there are stringent availability requirements and compromise of the KDC is a vulnerability. On-line CAs also have availability requirements (e.g. for timely revocation of certificates in the event of a compromise). Although more difficult to exploit than in a KDC, a covert compromise of a CA is a vulnerability, especially when certificates are used for authentication since bogus certificates can be signed and issued.

In cryptography, research is needed in application programming interfaces in order to promote wider use in NISs. Faster encryption and authentication/integrity algorithms are required in order to keep pace with increase in computational speeds and for deployment in a wider range of applications (e.g., multicast groups). For PKI, research needs to focus on the client/consumer side (as well as the issuer side). In particular, most applications have poor certificate management interfaces for users and system administrators, thereby introducing unnecessary vulnerabilities. Toolkits for certificate processing are weak, and further attention must be given to the issues of timely revocation of certificates, recovery from compromise of CA private keys, and name space management. Some obstacles in PKI will not be known until wide-scale deployments are attempted.

Network Access Control Mechanisms

In closed user groups (CUGs), subscriber communication is controlled on the basis of identities represented by network addresses. But this technique is not relevant in open systems like the Internet.

A virtual private network (VPN) is an illusion created using a public network, usually through administrative controls in central switches. Cryptographic controls, needed to prevent wiretapping, typically employ proprietary protocols, which limit wide-scale deployments. Adoption of the Internet Protocol Security (IPsec) will allow widespread use. But these protocols do not defend against attacks on the resources used to build the VPN.

A firewall is typically deployed at the boundary of a trusted and an untrusted network, placing restrictions on inbound and outbound traffic so that only messages perceived as safe are allowed to transit the firewall. Widely deployed for numerous benefits derived from the additional layer of

security, limitations do exist. For example, insiders can set up a Web proxy server on some outside machine, protocols implemented at specific layers of a network do not prevent attacks at higher layers, and utility is limited when using end-to-end encryption. Since most firewalls are implemented at the application layer, they are vulnerable to attacks directed at the operating system.

Guards have been primarily used in military systems to control the flow of information governed by mandatory access control policies. They have limited utility in more open environments.

Substantial work is needed to facilitate wide-reaching and flexible VPNs (e.g. support for dynamic location of security gateways, accommodation of complex network topologies, negotiation of traffic security policies across administratively dependent domains, support for multicast communication, and development of better interfaces for network management). Firewalls with enhanced capabilities will be needed in support of VPNs. Application-layer firewalls will be needed along with the development of useful traffic policies at this level.

Foreign Code and Application-Level Security

The traditional solution of executing only shrink-wrapped software purchased from trusted commercial venders does not protect against threats posed by foreign code. Programs may be downloaded from the Internet; software may be associated with Web pages (e.g., Java applets and Active X modules); weak OS security facilities may allow virtually unconstrained access to resources on a PC; and active document technologies may hide untrustworthy code, such as postscript files that do more than print and macros that exist in Word documents.

The Java Approach executes code in a confining Java Virtual Machine Environment (JVM) with an interpretation of a platform independent, stack-based intermediate language, typed information carried with variables, a run-time stack, signatures of routines that are defined and executed, and load-time checking of conformity to safety rules. However, some compilers now generate platform-dependent code that bypasses the Java security system, the JKD 1.2 model is complex and must be mastered by programmers, and users/programmers must correctly assess and configure suitable access rights for foreign code.

The Active X Approach allows modules to be digitally signed, but the presumption that a user can decide to execute code based on knowing the identity or seeing the credentials of the vender has proven to be of questionable effectiveness against malicious attacks. Also, revocation is a problem if the signing key is compromised.

Use of Fine-Grained Access Control (FGAC) and Application Security is an effective mechanism in enforcing the principle of least privilege. However, FGAC is generally not supported by the operating system (OS) and users and system administrators must configure it for all resources and modules. This is an impractical requirement; there may be a mismatch between application-level security policies and the FGAC configuration of low-level objects and permissions; and FGAC may be rendered ineffective by an OS that does not adequately protect the execution environment.

Language-based Security approaches are in their infancy and may be a vanguard of a new approach to the enforcement of security policies. One known approach consists of modifying the program prior to its execution by adding checks to prevent security violations. Another approach consists of analyzing (and effectively proving), prior to execution, that the security policies of interest will not be violated.

OS implementations of FGAC would help support the principle of least privilege. This will requires new mechanisms with negligible impact on overall performance. Also, management of FGAC must be

feasible and attractive for individual users and administrators. Application-level security is likely a shared responsibility between the application program and the security mechanisms provided at lower system levels. Research is needed to determine how to partition the responsibilities and what mechanisms are required at various system levels. More work is required to better understand the assurance limitations associated with application-level security when employing a COTS operating system. Opportunities offered by language-based approaches to security need further investigation.

Denial of Service

Denial of service may come in various forms, e.g., monopolization of use of resources, inactivation of a critical subsystem, contamination of the Internet Name Service caches, inactivation of packet routing, and diversion of traffic from its intended destination. Ad hoc solutions employed in an OS, e.g., preemption and scheduling algorithms seem unsuitable for NISs – no single trusted entity controls the agents making requests; individual servers may ignore untrusted requests, but they cannot terminate the requestors; even if detection and termination of a requester is possible, the cost may be prohibitive; and an attacker may cause a large number of surrogate clients to make unreasonable requests. Defending against denial-of-service attacks is important for ensuring availability of an NIS, but no general mechanisms or systematic approaches exist.

Other Aspects of Trustworthiness

In other aspects of trustworthiness, a number of issues arise that have implications for strengthening security.

The presence of redundant functionality with diverse interfaces can be used as a deterrent against attacks against interfaces.

Systems that monitor, detect, and respond can be used to amplify the trustworthiness of system components, but more research in needed.

Best placement of functionality to amplify trustworthiness requires much research. Graceful degradation has implications for availability of services. New types of algorithms may have the potential for usefulness in defining trustworthy systems (e.g., self-stabilization, emergent behavior, and biological metaphors). Risk management may be useful in trustworthiness, but security risks are difficult to identify and quantify. While risk avoidance may maximize trustworthiness, the cost may be prohibitive and a risk mitigation strategy may be more pragmatic. Consequences may be unpredictable and may affect people with varying levels of security. Useful metrics are unlikely to be developed for security because the corresponding formal models are necessarily incomplete.

As an alternative to general metrics, standards and criteria constitute a response for appraising trustworthiness and mitigating problems arising from imperfect information. Standards serve to simplify decision making by producers/consumers by narrowing the choices. Standards attract scrutiny to reduce design flaws and promote trustworthiness. They provide wide availability of technical information that serves as a basis for assessing where the vulnerabilities lie. Unfortunately, this situation increases the likelihood that a successful attack in one system might prove effective on a broader scale. A precise and testable definition is required to assess whether a standard has been fulfilled. Security-based criteria is a broader notion that rates the extent of the security mechanisms (functionality) and the degree to which the mechanisms can be trusted to perform their functions correctly (assurance). Criteria may improve level of information available to consumers and producers of components, but may have difficulty keeping pace with evolving threats. A rigorous methodology is needed for assessing the security of NISs assembled from components.