Cyber Security Research and Development Agenda

January 2003

Institute for
Information
Infrastructure
Protection

The Institute for Information Infrastructure Protection

The Institute for Information Infrastructure Protection (I3P) is a consortium of twenty-three academic and not-for-profit research organizations focused on cyber security and information infrastructure protection research and development (R&D). Its mission is to help protect the information infrastructure of the United States by developing a comprehensive, prioritized research and development agenda for cyber security, and promoting collaboration and information sharing among academia, industry, and government. The I3P embodies a concept developed and validated in studies between 1998 and 2000 by the President's Committee of Advisors on Science and Technology (PCAST), the Institute for Defense Analyses (IDA), and the Office of Science and Technology Policy (OSTP).

I3P Consortium Members

The I3P is currently composed of the following member organizations:

- Center for Information Security at the University of Tulsa
- Computer Security Laboratory at the University of California, Davis
- Critical Infrastructure Protection Project at George Mason University School of Law
- Georgia Tech Information Security Center
- H. John Heinz III School of Public Policy and Management at Carnegie Mellon University
- Information Security Laboratory at Oregon State University
- Institute for Civil Infrastructure Systems at New York University
- Institute for Security Technology Studies at Dartmouth College
- Johns Hopkins University Information Security Institute
- Los Alamos National Laboratory
- MIT Laboratory for Computer Science
- MIT Lincoln Laboratory
- The MITRE Corporation
- Mitretek Systems
- Pacific Northwest National Laboratory
- The RAND Corporation
- Sandia National Laboratories
- Software Engineering Institute at Carnegie Mellon University
- SRI International
- Stanford University Computer Science Department
- University of California, Berkeley
- University of Illinois at Urbana-Champaign
- University of Virginia

Abstract

This initial Cyber Security Research and Development (R&D) Agenda identifies R&D topics of significant value to the security of the information infrastructure that are either not funded or under-funded by the collection of private sector and government-sponsored research activities in the United States. The Agenda is based on information gathered and analyzed during the 2002 calendar year and reflects the input of experts in industry, government, and academia. The Agenda, together with that supporting information, is intended to aid researchers in identifying problems and R&D program managers in defining program directions. Areas in which new or additional research is needed include:

- Enterprise Security Management
- Trust Among Distributed Autonomous Parties
- Discovery and Analysis of Security Properties and Vulnerabilities
- Secure System and Network Response and Recovery
- Traceback, Identification, and Forensics
- Wireless Security
- Metrics and Models
- Law, Policy, and Economic Issues

Executive Summary

The Institute for Information Infrastructure Protection (I3P) is a consortium of twenty-three academic and not-for-profit research organizations focused on cyber security and information infrastructure protection research and development (R&D). Its mission is to help protect the information infrastructure of the United States by developing a comprehensive, prioritized R&D Agenda for cyber security and promoting collaboration and information sharing among academia, industry, and government. This document constitutes the initial Cyber Security R&D Agenda. The I3P embodies a concept developed and validated in studies from 1998 and 2000 by the President's Committee of Advisors on Science and Technology (PCAST), the Institute for Defense Analyses (IDA), and the Office of Science and Technology Policy (OSTP).

The information infrastructure consists of technologies and capabilities for gathering, handling, and sharing information that are accessible to, or commonly depended upon by, multiple organizations, whether within a single enterprise, a critical infrastructure sector such as banking and finance, the U.S. Government, the nation as a whole, or transnationally. The information infrastructure includes well-engineered systems as well as poorly configured systems in businesses and homes.

The Internet is perhaps the most obvious element of the information infrastructure; other easily recognized components include such widely used products as desktop operating systems and routers, the devices that handle message transfers between computers. The development of this infrastructure over the past two decades has been swift and has permanently changed the way the nation conducts business, operates its governmental structures and armed forces, keeps its people healthy and safe, and spends its leisure time. The changes have been fundamental and are all but irrevocable.

The information infrastructure, taken as a whole, is not an engineered system. It is the result of the entrepreneurial efforts and the collective genius of the nation, working to improve efficiency and provide new opportunities for people and businesses. Security was not a significant consideration at its inception, and security concerns today do not override market pressures for new uses of technology or innovation, in spite of frequent stories of hackers, criminals, and, increasingly, terrorists and nations using or planning to use the information infrastructure as a weapon to harm the United States.

In the United States, the private, academic, and public sectors invest significant resources in cyber security. The commercial sector primarily performs cyber security research as an investment in future products and services. While the public sector also funds R&D in cyber security, the majority of this activity focuses on the specific missions of the government agency funding the work.

Thus, broad areas of cyber security remain neglected or underdeveloped. This Cyber Security R&D Agenda is intended to identify the highest-priority gaps: R&D problems of significant value to the security of the information infrastructure that are either not funded or under-funded within the collection of private sector and government-sponsored research in the United States. The Agenda is based on information gathered and analyzed during the 2002 calendar year, and it reflects the input of experts in industry, government, and academia. The Agenda builds on the work of other research, industry, and government organizations that have focused on cyber security issues.

Supporting documents include a gap analysis that compares user needs to existing products and research; a survey of products, tools, and services; a survey of research and development; a survey of related roadmaps and R&D agendas; and a capstone document that provides context for and identifies crosscutting and pervasive issues arising from the three surveys. These substantial documents can be found on the I3P web portal, at http://www.thei3p.org/ecommunities/abouti3p.jsp.

Areas in which new or additional R&D is needed to improve the security posture of the information infrastructure include:

- **Enterprise Security Management**
  Each piece of the information infrastructure may be owned by individuals or enterprises, but we are all interconnected. Therefore, the enterprise security management (ESM) challenge is to integrate diverse security mechanisms into a coherent capability for managing access to and use of enterprise resources, monitoring behavior on enterprise systems, and detecting and responding to suspicious or unacceptable behavior. While the marketplace offers product suites under the rubric of ESM, the problem area is broader than the fragmented capabilities provided by existing products. Research needs remain in the areas of enterprise policy definition and management, definition and maintenance of a targeted risk posture, and definition of, and protection at, security boundaries. IT-based collaboration with partner organizations, and increased services to home users make these boundaries more complex and extend the definition of "insiders." Further research is needed to address the insider threat.

- **Trust Among Distributed Autonomous Parties**
  In cyberspace, entities—individuals, organizations, software, and devices—need to establish relationships dynamically and without recourse to a central authority or previously determined trusted third party. Existing research, particularly in terms of the techniques entities use to establish trust in the security of other entities, is expected to address many of the needs articulated by enterprise users. However, solutions are needed that address the autonomy, scale, complexity, and dynamism of critical infrastructures. Research needs exist for trust models for autonomous entities that are geographically or organizationally distributed, definition and management of dynamic security relationships in peer-to-peer settings, techniques for developing trust relationships between systems and end-user devices such as cell phones or laptops, and approaches to establish trust in data.

- **Discovery and Analysis of Security Properties and Vulnerabilities**
  The information infrastructure has a large number and variety of components, in different forms: hardware, firmware, software, communications media, storage media, and information. Frequently, the properties of these components are poorly understood, due to undocumented functionality, flaws in their design or implementation, or unanticipated uses. Products and systems commonly include vulnerabilities and inadequately understood security properties. Moreover, the security properties of a system or subsystem cannot be derived or deduced from those of its components, and emergent properties of large-scale systems are difficult to describe, much less predict. Considerable effort has been applied to the problem of ensuring the presence of desired security properties and preventing (or determining the presence of) vulnerabilities. The need is acute for ways to determine, throughout a product or system's life cycle (development, integration, update and maintenance, decommissioning, or replacement of components), whether exploitable defects have been introduced or unanticipated security properties have emerged or escalated. Research is needed into techniques, embodied in tools to ensure their utility, to analyze code, devices, and systems in dynamic and large-scale environments.

- **Secure System and Network Response and Recovery**
  The proliferation of numbers and types of computing devices has resulted in the increasing size and complexity of the information infrastructure. Response to and recovery from attacks against such multifaceted systems are hindered by this inherent complexity. As a result, response across a set of organizations is often uneven and difficult to coordinate, and reconstitution to a secure state can be difficult. The potential for survivability from attacks and in making intrusion detection systems more proactive has driven research into secure response and recovery. Current research, however, does not adequately address the issues of scale, coordination across different administrative and policy domains, or coordination across the highly diverse systems that are the hallmarks of information infrastructure protection. Research needs remain in the areas of prediction or pre-incident detection, as well as recovery and reconstitution for systems of systems.

- **Traceback, Identification, and Forensics**
  During and after an attack, responding organizations must have prompt and reliable information to determine and implement an appropriate response. Current capabilities are oriented toward enabling the enterprise to detect and respond internally to suspected attacks. Research is needed into capabilities that enable responders to trace back, or identify the source location of the attack; to identify the individual, group, or organization originating the attack; and to determine the actual nature of the attack. Companion research is needed to address the legal and policy implications of such capabilities.

- **Wireless Security**
  Wireless technologies are increasingly crucial to enterprise systems and across critical infrastructure sectors. Wireless networks include not only wireless telecommunications *per se*, but an increasingly diverse set of end devices, such as sensors, process controllers, and information appliances for home and business users; in some cases, end devices may also provide wireless telecommunications services. In principle, many of the security concerns for wireless networks mirror those for the wired world; in practice, solutions developed for wired networks may not be viable in wireless environments. Private sector concern, and thus investment, focuses on proprietary or enterprise solutions. Research is needed to make security a fundamental component of wireless networks, develop the basic science of wireless security, develop security solutions that can be integrated into the wireless device itself, investigate the security implications of existing wireless protocols, integrate security mechanisms across all protocol layers, and integrate wireless security into larger systems and networks. In particular, research is needed into security situation awareness techniques for wireless networks and strategies to address distributed denial-of-service attacks.

- **Metrics and Models**
  Individuals, organizations, and critical infrastructure sectors bear the risks of relying on the information infrastructure. For organizations to manage cyber security risks—to accept a given level of risk, transfer or externalize risk, or apply resources to decrease the level of risk to an appropriate balance—decision makers need a clear and defensible basis for making investment decisions that can be related to organizational missions and strategies. That basis should be founded on rigorous and generally accepted models and metrics for cyber security. Research is needed to provide a foundation of data about the current investment and risk levels. Research is also needed to define metrics that express the costs, benefits, and impacts of security controls from multiple perspectives—economic, organizational, technical, and risk—so that the dynamics at work in making security decisions can be better understood. Finally, research is needed into techniques for modeling the security-related behavior of the information infrastructure and predicting consequences of risk management choices.

- **Law, Policy, and Economic Issues**
  Decisions that affect the security posture of the information infrastructure are made in a poorly understood context of economic factors, laws, regulations, and government policy. Research is needed to determine the actual magnitude of the cyber security problem and enable a better understanding of the relationships between the forces that shape information infrastructure protection (i.e., research into the structure of the market, and to determine how changes in laws, policy, and economic conditions, as well as technology, affect one another). For any emerging technology, companion research is needed into the legal, policy, and economic implications as well as the cyber security implications of the technology and its possible uses. Research is needed to describe the structure and

dynamics of the cyber security marketplace, as well as the impacts of various interventions—changes in enterprise purchasing patterns, cyber security laws, regulations, government acquisition practices, policies, auditing practices, insurance and other factors—on cyber security in general, and on the development, deployment, and use of cyber security technology in particular. Research is also needed into the implications of implementing alternative strategies for allocating responsibility for security in cyberspace, and into tradeoffs among stakeholder concerns. In particular, there is a need for research into the role of standards and best practices in improving the security posture of the information infrastructure, the policy and legal considerations associated with collecting and retaining data about the information infrastructure and its uses, and the implications of potential changes to laws or policies that would be intended to enable direct responses to attacks.

Participants in the R&D Agenda development process frequently noted the importance of education, training, and awareness; quality assurance methodologies; information sharing and coordination; practicable procedures; and physical security. While these areas are not identified as research gaps in this Agenda, they are important considerations for researchers who seek ultimately to affect the practice of cyber security. Similarly, technology transfer was frequently identified as problematic, highlighting the need for cyber security R&D programs to explore innovative strategies for improving the flow of ideas and technologies between researchers, product developers, system integrators, and end user organizations.

Table of Contents

Section 1: Introduction

This document presents an initial Cyber Security Research and Development (R&D) Agenda for information infrastructure protection that focuses on high-leverage, under-served areas of cyber security R&D.[1] This Agenda was developed by the Institute for Information Infrastructure Protection (I3P) based on information gathered and analyzed during the 2002 calendar year, and it reflects the input of experts in industry, government, and academia.[2] The purpose of the Cyber Security R&D Agenda is to identify priority R&D areas for information infrastructure protection that are either not funded or under-funded by the collection of private sector and government-sponsored research activities.[3] By identifying and prioritizing high-leverage cyber security research "gaps" of national importance, this Agenda also serves as the basis for the I3P's program planning and its research funding and evaluation processes.

This introductory section describes the history of the I3P, the information infrastructure protection problem, and the nature and scope of this R&D Agenda. Appendix A describes on-line resources provided by the I3P. Section 2 briefly summarizes the process the I3P used to construct this Agenda; Appendix B describes that process in more detail. Section 3 presents the outcome of that process: eight research areas are described, representing the gaps between existing or anticipated capabilities and the highest-priority needs identified by experts and stakeholders. Section 4 provides brief conclusions.

1.1 The Origin of the I3P

In 1998, the President's Committee of Advisors on Science and Technology (PCAST) recognized that investments in information security R&D were made primarily on a tactical basis—to fulfill an immediate perceived need—or for private sector commercial reasons. There was no institution or collection of institutions that (1) looked at the landscape defined by the state of the art in information security and the existing body of ongoing public and private R&D, and (2) identified the gaps in the national information security R&D portfolio. In late 1998, the PCAST recommended that the government fund an independent, non-governmental, non-commercial laboratory that would accomplish this important task by articulating the nation's information security requirements, cataloging ongoing R&D efforts, identifying gaps in the nation's R&D portfolio, and conducting research in these critical gap areas. The PCAST also recommended that this institution should have $100 million available per year to fund these activities after a start-up period.[4]

In early 1999, the President agreed with the importance of protecting the nation's information infrastructure and directed the National Security Council (NSC) and the

[1] The preparation of the I3P 2003 Cyber Security R&D Agenda was performed under the sponsorship of the U.S. Department of Commerce, National Institute of Standards and Technology.
[2] See Section 2, "Methodology," for a more detailed description of the R&D Agenda development process.
[3] This collection of existing publicly and privately funded cyber security research will be referred to herein as the "national R&D portfolio." The areas that were identified as important but were not sufficiently funded in the existing national R&D portfolio are called "gaps."
[4] PCAST Letter to the President, December 10, 1998.

Office of Science and Technology Policy (OSTP) to perform an immediate review of the PCAST proposal. The NSC and OSTP staff asked the Institute for Defense Analyses (IDA) to study the matter.

While IDA was working on its study, in late 1999, the OSTP and PCAST met with 15 Chief Technology Officers from leading information technology corporations, who endorsed the concept. In early 2000, the President's fiscal year 2001 budget proposal requested $50 million to fund the I3P.

In the spring of 2000, IDA completed its study, which supported the establishment of the I3P, recommended that it perform the functions described by the PCAST, and supported the PCAST proposal that the mature institute receive funding of $100 million per year. The Institute, it suggested, would disburse most of this money to outside cyber security researchers. Shortly thereafter, the OSTP issued a white paper finalizing the concept and endorsing the basic structure of the I3P.

In both FY 2001 and FY 2002, Congress appropriated seed funding to the Institute for Security Technology Studies (ISTS) at Dartmouth College to establish the I3P.

1.2 The Information Infrastructure Protection Problem and the Growing Threat

The information infrastructure consists of technologies and capabilities for gathering, handling, and sharing information that are shared, or commonly depended upon, by multiple organizations, whether within a single enterprise, a critical infrastructure sector such as banking and finance, the U.S. Government, the nation as a whole, or transnationally. Almost every aspect of contemporary life, for individuals, businesses, and governments, depends in some way on the information infrastructure. Business, the delivery of essential services, national security, leisure, and the conduct of our personal affairs increasingly rely on our ability to connect and communicate with people and our environment using information technologies. We are far down this path of information reliance; pre-information age ways of doing things are rapidly being replaced, and in many cases they are no longer available to us. There is no going back.

The significant gains in productivity seen in the last ten years, made possible in large part by the information infrastructure, came with significantly increased dependence on these technologies and the environment in which they are developed, fielded, and used.[5] Yet, the information infrastructure is vulnerable to breaches, cyber attack, and cascading or pervasive failures.

Vulnerabilities in the information infrastructure arise from many sources, including the lack of inherent security in new technologies, flaws in commonly used products, and

[5] Chairman Alan Greenspan of the Federal Reserve has often testified before Congress about rapid increases in productivity. See, for example, his testimony before the Banking, Housing and Urban Affairs Committee on July 20, 2000, at http://www.federalreserve.gov/BoardDocs/hh/2000/July/testimony.htm. Also, see increases in Gross Domestic Product, www.bea.gov/bea/dn/nipaweb/TableViewFixed.asp#Mid.

organizational failures to address security concerns in system design and use.[6] While criticism has on occasion been directed at specific organizations with varying degrees of justification, many information infrastructure vulnerabilities arise from the facts that information technology is evolving rapidly, it is being used in unanticipated ways to enable new business practices, and systems and products are interconnected—and becoming interdependent—in ways that lead to unintended consequences. Complexity is now a defining characteristic of the information infrastructure.

There is a wide array of actors, ranging from teenage hackers seeking bragging rights, to criminal organizations and terrorists who would do us harm, to foreign nations conducting espionage or military operations, that have the motives, capabilities, and opportunities to exploit these vulnerabilities. Although attack methods can be physical (e.g., bombs), social (e.g., rumors, disinformation, deception), cyber (i.e., relying on information technology), or a combination thereof, the I3P focus is on R&D for protecting the information infrastructure solely from cyber threats. Vulnerabilities to cyber threats can be created deliberately or erroneously by, or accidentally triggered by the actions or decisions of, end users, administrators, developers or integrators, or strategic planners defining information technology (IT) uses and architectures. These vulnerabilities can also be the unintended consequences of decisions made by policy-makers.

The threat to the information infrastructure continues to grow. The nation's dependence on information and computer networks for communications, data management, and the operation of critical infrastructures renders it increasingly vulnerable to computer-based, or cyber, attacks against our information infrastructure, including the Internet, telecommunications networks/backbones, and interconnected computer systems.[7] Cyber attacks now threaten not only our information infrastructure but also other critical infrastructures—such as banking and finance, transportation, and energy—that rely on information technology. Moreover, because these infrastructures are highly interdependent, attacks on one infrastructure can damage other infrastructures as well. Concentrated infrastructure attacks could thus have a significant effect on our national security and economy.

Significantly, the problem is not confined to the stereotypical teenage hacker who attacks systems for the mere challenge of it. Rather, recent years have seen a growth in attacks coming from much more sophisticated actors, such as organized crime, terrorist groups, and, most significantly, foreign nations that are developing cyber attack techniques for activities ranging from covert espionage against U.S. Government agencies or U.S. industry to information warfare against the United States. As a recent Defense Science Board report stated, "At some future time, the United States will be attacked, not by

[6] These are but a few of the significant areas highlighted repeatedly in I3P workshops, surveys, and consultations with experts.

[7] The last few years have witnessed an ever-growing number of damaging attacks, including viruses and worms, distributed denial-of-service (DDoS) attacks, and unauthorized intrusions. The Computer Emergency Response Team (CERT) Coordination Center indicates that the number of reported computer security incidents has more than doubled during each of the past three years (http://www.cert.org/stats/cert_stats.html).

hackers, but by a sophisticated adversary using an effective array of information warfare tools and techniques. Two choices are available: adapt before the attack or afterward."[8]

Cyber security is also vital to protecting personal privacy, an issue about which the American public is becoming increasingly concerned. The Federal Trade Commission recently observed that Americans are especially worried about "the specter of identity theft."[9] Thieves may use stolen personal information to access a victim's existing accounts, create new accounts in the victim's name, or commit other types of fraud. Since terrorists use identity theft to facilitate their movements and operations, it is one of several areas in which privacy and national security concerns overlap.

But while the cyber security problem continues to grow and public concern increases commensurately, the state of our technical defenses is not keeping pace. Indeed, the widening knowledge gap between cyber attackers and defenders led the Chairman of the Defense Science Board to conclude that the "[Department of Defense] cannot today defend itself against an Information Operations attack by a sophisticated nation state adversary."[10] Yet, Office of Management and Budget (OMB) reports on computer security at federal agencies have consistently identified the Department of Defense as a relative bright spot. The OMB recently reported that "many [other] agencies have virtually no meaningful systems to test or monitor system activity and therefore are unable to detect intrusions, suspected intrusions, or virus infections."[11] The private sector is in much the same shape, with some corporations taking proactive measures, while others are using much more rudimentary security, if any at all. And just as in the government, no company is impervious to sophisticated attack. Even private companies that employ intrusion detection systems (IDSs) and other computer security measures find that their vulnerabilities are increasing, and this is borne out by statistics on vulnerabilities and attacks.[12] On the one hand, the core technologies underlying the Internet were not built with security in mind. On the other hand, the growing complexity of information technologies multiplies attack routes and makes it harder to anticipate how problems will cascade through information networks.

In the words of Richard Clarke, Chairman of the President's Critical Infrastructure Protection Board, "Our infrastructure is fragile." The United States urgently needs new technologies that will identify and fix vulnerabilities, and "harden" and protect our information infrastructure, making it more robust and resilient in the face of attacks.

[8] *Protecting the Homeland*, Report of the Defense Science Board Task Force on Information Operations, 2000 Summer Study, Volume 2 (http://www.acq.osd.mil/dsb/reports.htm).

[9] Prepared statement of the Federal Trade Commission on Identity Theft: The FTC's Response before the subcommittee on Technology, Terrorism and Government Information of the Senate Judiciary Committee March 20, 2002 (http://www.ftc.gov/os/2002/03/idthefttest.htm).

[10] *Protecting the Homeland*, Report of the Defense Science Board Task Force on Information Operations, 2000 Summer Study, Volume 2 (http://www.acq.osd.mil/dsb/reports.htm).

[11] *Office of Management and Budget FY 2001 Report to Congress on Federal Government Information Security Reform* (http://www.whitehouse.gov/omb/inforeg/fy01securityactreport.pdf).

[12] See, for example, the statistics regularly reported by the CERT/CC on incidents and vulnerabilities, at http://www.cert.org/stats/cert_stats.html.

1.3 Nature and Scope of the I3P R&D Agenda

This R&D Agenda is intended to serve as a national agenda, that is, to identify R&D topics of national importance that can be of use to all researchers and U.S. R&D program managers, public or private. The I3P's focus is on high-leverage cyber security R&D to address *information infrastructure protection* problems of national importance that are under-funded or not funded in the national R&D portfolio. The national R&D portfolio has two major components, private sector and publicly funded R&D.[13] The commercial sector generally performs cyber security research as an investment in future products and services. This investment is, by definition, motivated by profits, and the marketplace is not structured to address the full range of information infrastructure vulnerabilities for several reasons. First, responsibility for vulnerabilities that arise from interconnections and interdependencies is often unassigned, particularly when the security postures of organizations or infrastructures that rely upon each other are not comparable.[14] In these cases, there is often insufficient economic motive to address such vulnerabilities. Second, security capabilities also compete with—and frequently impede—functional capabilities that consumers want. Since the demand for security is quite small compared to the demand for functionality, security features are relatively unattractive to profit-oriented organizations. Finally, security risks are often transferred to consumers who frequently do not understand them and lack recourse when they become victims of a cyber attack.

The public sector also funds R&D in cyber security, but the majority of this activity, quite reasonably, focuses on the specific missions of the government agency funding the work. Once again, responsibility for common security is not firmly attributed to any agency.[15]

Thus, the need for a national R&D agenda to tackle neglected areas is clear. This R&D Agenda does not try to address the entire scope of national R&D needs for cyber security. Instead, it addresses the gaps in the national R&D portfolio, and complements existing R&D efforts in the public and private sectors. This document could also serve as one input into the R&D component of the National Strategy to Secure Cyberspace being

[13] In surveying the national R&D portfolio in 2002, the I3P restricted its attention to the open literature. While proprietary R&D in the private sector and classified R&D in the public sector may eventually result in improvements to the security posture of the information infrastructure, the transition process is frequently problematic.

[14] Vulnerabilities caused by the interdependencies of infrastructures are numerous. If company A relies on company B for a service but views confidentiality or integrity as crucial, while company B views availability as its top priority, the security postures are not comparable and the security needs of company A are not met. For example, a bank that provides online financial services depends on the telecommunications infrastructure to move the information needed to complete this transaction. It therefore inherits some of the vulnerabilities of the telecommunications providers over which this information passes. For a rigorous discussion of interdependencies, see Rinaldi, Peerenboom, and Kelly, *Complexities in Identifying, Understanding, and Analyzing Critical Infrastructure Interdependencies,* IEEE Control, December 2001.

[15] The one exception to this mission orientation is the National Science Foundation, whose charter is to fund research that advances science.

developed by the PCIPB in coordination with the OSTP, the Homeland Security Strategy, and the National Security Strategy.

It is worth noting that the I3P's efforts to develop a cyber security R&D Agenda differ from other efforts in important ways. First, the process (described briefly in Section 2 and in greater depth in Appendix B) is not a one-time effort. It is cyclical, based on recognition of the dynamic nature of the information infrastructure and the fact that any cyber security R&D Agenda that is not periodically updated will quickly lose its value. It is also procedurally rigorous; great pains were taken to identify cyber security requirements, promote the evolving state of cyber security and research, identify gaps in the national R&D portfolio, identify priority R&D areas, craft an R&D Agenda that addresses them, and validate that Agenda with key stakeholders. To do this, the I3P elicited input from a spectrum of stakeholders, including representatives from key critical infrastructures, nationally recognized experts, and government officials. Importantly, this is not a pro bono effort; the I3P is a full-time operation, with dedicated staff and expert assistance.

Furthermore, the I3P is attempting to forge a dynamic relationship with and among researchers in this field through the I3P Web portal, which was specifically constructed with this goal in mind. The I3P also seeks to provide resources to the cyber security research community, both through the I3P Web portal and through its Digital Archive (both described in Appendix A). Finally, given sufficient funding, the I3P will fund R&D in the important areas identified in Section 3.

Section 2: Methodology

This section summarizes the methodology or process that the I3P used to develop the Cyber Security R&D Agenda. Appendix B contains a detailed description of the process.

First, the I3P determined cyber security needs, unconstrained by the limitations of stakeholder resources, using four complementary and concurrent approaches. The first was a series of stakeholder workshops and supporting Web-based surveys, to elicit input about needs and perceptions of threats, vulnerabilities, and consequences from the Banking and Finance, Energy, Chemical, Water, IT and Telecommunications, Emergency Response, Manufacturers and Vendors, Researchers, Government, and Transportation sectors. Second, the I3P supported three cluster groups—researchers, technologists, and policy experts focused on high-importance security topics—in the areas of Enterprise Security Management, Wireless Network Security, and Legal, Policy, and Economic Issues. Third, the I3P identified the requirements and priorities of several public and private sector organizations by reviewing their R&D agendas and roadmaps. This R&D Agenda builds on the extensive efforts of many others.[16] Finally, the I3P sought the insights of the nation's leading experts in information security, many of whom are members of the I3P Consortium.

Simultaneously, the I3P established a baseline for the current and anticipated state of information infrastructure protection. That baseline includes information about and analysis of the cyber security marketplace and existing R&D that address current and future information infrastructure needs. To do this, the I3P surveyed and assessed cyber security products and services, and identified ongoing and planned research in the public and private sectors that could be applied to information infrastructure protection.

Next, the I3P performed a gap analysis to identify the high-payoff, under-funded gaps in the national R&D portfolio.[17] The I3P identified gaps by comparing the results of the baseline and requirements determination efforts described above to establish the sufficiency of products and current or planned R&D. These results were grouped into research areas. Then, the I3P priority research areas were identified using data from the surveys and workshops, the input of the cluster groups, and the perspective of experts. The determination of priority areas also considered factors such as the relative importance of different needs, the likelihood of R&D success, and the timeframe in which R&D occurs.

As stated earlier, an important aspect of the I3P effort is that this is not a one-time-only Cyber Security R&D Agenda. It will be reviewed and updated periodically, as the state of the information infrastructure changes and as I3P research is funded.

[16] See the list of Related Roadmaps, R&D Agendas, and Studies in the References, including the studies and agendas developed by the National Research Council. See also specific references in the gap analysis and baseline surveys.

[17] The gap analysis, baseline surveys, cluster group reports, and summaries of the stakeholder workshops are available at the I3P Web portal. See Appendix B for more details.

Section 3: Research Areas for Information Infrastructure Protection

This section describes the results of the Cyber Security R&D Agenda development process described in Section 2. Workshop participants, surveyed experts, cluster group members, I3P Consortium members, and other experts in critical fields repeatedly indicated that the research areas identified here are crucial to information infrastructure protection. These are also gap areas in the national cyber security R&D portfolio. Some of these areas are relatively broad, and in some there is substantial ongoing research. They are identified here, however, because important research topics within them require additional attention and resources.

It should be noted that the complex nature of cyber security makes it impossible to define mutually exclusive research areas. For example, overlap exists between the areas of wireless security and enterprise security management. Some important topics, such as privacy, are reflected in multiple areas. In addition, not all research areas are of comparable scope. The research areas presented below were specified to highlight most effectively the nation's cyber security needs and to focus research on the issues of greatest importance. The resulting variations in levels of abstraction, and possible overlaps in definitions, should not materially detract from the quality or merit of the identified research areas.

The research areas identified by the I3P are:

- Enterprise Security Management
- Trust Among Distributed Autonomous Parties
- Discovery and Analysis of Security Properties and Vulnerabilities
- Secure System and Network Response And Recovery
- Traceback, Identification, and Forensics
- Wireless Security
- Metrics and Models
- Law, Policy, and Economic Issues

A few general observations are in order before discussing the research areas. First, this R&D Agenda focuses on cyber security research, not general IT research. However, the two fields cannot be cleanly partitioned. For example, one of the recurring stakeholder observations was that, to be maximally effective, security controls must be easy to use, adjust, and update; thus, cyber security research needs to incorporate the results of research into human-computer interfaces and visualization. Consequently, elements of research that are not exclusively cyber-security-oriented have been included as necessary to motivate a viable and complete security solution.

Second, significant challenges must be addressed when extending existing security models to emerging technologies, architectures, and applications of technologies. The assumptions behind existing models (e.g., amount of local storage and computing power, centralization of administration) do not hold consistently across the information infrastructure. Identification and analysis of vulnerabilities of emerging technologies, architectures, and uses is crucial to the secure evolution of the information infrastructure. This issue is addressed in several research areas, notably wireless security.

Third, while the goal of providing cyber security measures for information infrastructure protection poses many complex technical challenges, a purely technical approach will be inadequate. Workshop participants repeatedly emphasized the importance of the human and organizational elements, and the need to address the larger context of legal, policy, social, and economic factors. Technical solutions must be usable, compatible with organizational objectives and processes, and viable in the larger context. This theme is visible in several research areas.

Finally, the development of security solutions is complicated by such characteristics of the information infrastructure as dynamism, complexity, diffusion of control, and increasingly sophisticated threats. Thus, cyber security research needs to consider issues arising from uneven technology transition (e.g., the transition from IPv4 to IPv6, the transition from circuit switched to IP-based voice networks), differences in policies and risk management strategies as embodied in different technologies, and the expectation that the effectiveness of any given solution will degrade over time.

3.1 Enterprise Security Management

Brief Problem Description

While the information infrastructure includes widely used products and standards, its most visible components are systems and networks owned and managed by individual enterprises. Each enterprise must define policies for secure and appropriate use of its information resources and translate those policies into practice. Enterprise systems are increasingly complex, due to the pace of technological change, the exploration of new uses of IT, and changes to the enterprise itself such as mergers, acquisitions, and reorganizations. Changes to enterprise systems, as well as IT-based collaboration with partner organizations, extend the definition of "insiders," even as the insider threat is increasingly recognized as a significant problem.[18] As a result, the Chief Information Officer of an organization may be unable to state with confidence exactly which policies are actually enforced by its own enterprise systems, or

The enterprise security management (ESM) challenge is to integrate diverse security mechanisms into a coherent capability for managing access to and use of enterprise resources, monitoring behavior on enterprise systems, and detecting and responding to suspicious or unacceptable behavior.

[18] An insider may be defined as anyone who has or had an one time authorized access to a system as an employee, but also may include contracted users, partnering vendors, temporary employees, etc. See, for example, *Research and Development Initiatives Focused on Preventing, Detecting, and Responding to Insider Misuse of Critical Defense Information Systems*, CF-151-OSD, 1999, http://www.rand.org/publications/CF/CF151/CF151.pdf and *Research on Mitigating the Insider Threat to Information Systems - #2: Proceedings of a Workshop Held August, 2000*, http://www.rand.org/publications/CF/CF163, and *Trends in Proprietary Information Loss*, September 2002, sponsored by PricewaterhouseCoopers, U.S. Chamber of Commerce, and the ASIS Foundation, http://www.asisonline.org/pdf/spi2.pdf.

even what components and software are included in those systems. As attacks increasingly exploit connectivity and functional dependencies across critical infrastructure sectors, the impacts of decisions about policies, processes, procedural and technical safeguards, and risks propagate beyond an enterprise's own systems.

Existing Research and Capabilities

While the marketplace offers product suites under the rubric of ESM, the problem area is broader than the fragmented capabilities provided by existing products. Stakeholders repeatedly identified ESM as a critical area due to the difficulties involved with linking security concerns within and across organizations to disparate, non-standard, and oftentimes separately controlled security mechanisms.

R&D in the area of ESM has generally focused on specific problem areas, including configuration management, monitoring of activities by users (particularly privileged users), patch management, and integration. The area of security configuration management has been and continues to be addressed by both the research community and commercial vendors. The ability effectively to manage and track configuration changes within enterprise systems is considered good engineering practice, and applying security principles to this area has been a focus of these efforts for some time. Efforts increasingly focus on automatic generation of configurations consistent with a stated security policy, as well as automatic checking of configurations.

Monitoring of user activity, including privileged users such as administrators, is an active R&D area. Increasing concern about the insider threat has led to development of initial capabilities for insider monitoring. Monitoring of security management activities has been primarily addressed through the capabilities of general-purpose systems that can log privileged activities, including those that affect the security posture of a system. While work is progressing in this area, research is needed to link monitoring activities with security policy definition and enforcement in order to improve the overall effectiveness of monitoring.

Security patch management is a critical ESM problem that is being addressed within the vendor community. Security administrators constantly struggle to keep abreast of, test, and install patches for critical security issues and to ensure there are no unintended consequences when the patches are deployed in their complex enterprises. Work has been ongoing to address how to make these patches widely available, yet still establish trust in the validity and security of these patches (e.g., to ensure that a downloaded patch has not been modified to include malicious or unwanted code). No significant work, however, is being done on providing an environment that allows for repeatable testing of applications before or in conjunction with identifying and installing a patch. Additional research is needed into ways to determine the effects of security patches on as-used systems (e.g., one patch can undo the effects of another, make an existing application unusable, or conflict with the policy of the system's owner).

Finally, existing R&D has also focused on discovering ways to manage security across heterogeneous environments. This includes research into role-based access control (RBAC) and modeling and mapping of access control attributes. Efforts to provide a uniform interface to diverse products have mostly been driven by third-party vendors (i.e., those that are creating businesses based on providing these capabilities). Current research into meta-languages and taxonomies to correlate security management and monitoring data is expected to enable better integration of capabilities across architectural layers (e.g., integration of authentication across operating systems, middleware, and applications; integration of intrusion detection across hosts and networks).

Gaps between stakeholder needs and current or anticipated capabilities arise in part from the difficulty of translating between organizational and system behavior. On the one hand, an enterprise needs to define policies for the appropriate use of its information resources in a way that is consistent with its goals and that can be translated into practice by implementing, configuring, managing, and administering security capabilities. On the other hand, it may not be clear which policies are actually enforced by enterprise systems. This difficulty is increased by the diffusion of control across such information infrastructure components as systems, distributed applications, and communications and storage networks. Due to the inherent complexity of the information infrastructure, gaps also arise in the areas of defining, expressing, and maintaining a level of cyber security risk. Finally, gaps emerge as a result of the increasingly dynamic nature of the information infrastructure; defining and providing protection at security perimeters is increasingly problematic as system boundaries are less well-defined.

I3P Research Areas

The I3P gap analysis indicates that new or additional R&D in the following areas would have high leverage in improving the state of information infrastructure protection:

- Enterprise policy definition and management

  As stated above, the problem of defining and managing enterprise policies for the use of information resources is two-fold: first, how can policies be defined in a way that is both consistent with the goals of the enterprise and translatable into practice? Second, how can security managers more easily and effectively determine which policies are actually enforced by enterprise systems? This area is of concern within an enterprise (particularly since a given enterprise can include multiple security policy domains). It is also of concern in the context of inter-enterprise resource sharing. Implementing controls across heterogeneous technologies to meet a stated policy is an important problem for the information infrastructure.

  From security configuration management to user authentication and authorization, each technology brings with it unique capabilities, interfaces, and controls. Issues such as the scope and granularity of control, the strength of protection, and the ease of configuration all impact the consistent enforcement of policy across technologies. Efforts to provide administrative interfaces across similar products (e.g., user account setup) have been limited and have not addressed the larger need to apply security controls consistently across various security components. Comprehensive tools for creating an accurate, real-time IT inventory of all components (e.g., hardware, software, applications, configurations) would help security managers define and enforce security policies. Better, more accessible interfaces are needed both to assist security managers and system administrators in expressing the effects of policies and to help users understand them. Visualization techniques, such as user-friendly graphic displays, that extend to wireless networks and dynamic systems are especially needed. More broadly, research is needed into questions of which enterprise security management decisions could be completely automated, and what the practical, organizational, and policy implications of such automation would be.

Further, there is a particular need for definition of authorization and access controls, and management of user access (e.g., setup, role-based profiling, changes in roles and responsibilities, terminations) across infrastructure components. This includes negotiating authorization rights among infrastructure components that may employ vastly different methods and degrees of control. Research is needed into issues of ensuring the security of ESM technologies, data, and data exchange in heterogeneous environments. Research must also include an examination of the concept of trusted insiders, including consideration of different types of insiders (e.g., users, administrators, outsourcing providers, policy makers, and strategic planners). Research must consider how issues of cross-domain access, auditing, monitoring, evidence collection, and investigation should be handled, and the privacy implications of different strategies for addressing these issues. Finally, research must address social and organizational, as well as technical, aspects of these problems.

- Definition and maintenance of a target risk posture

  Making decisions on how to spend limited resources on security solutions has become increasingly difficult, due to the complexity and heterogeneity of existing environments and the continuous introduction of new technologies and configuration changes into these environments. The questions of how much to spend on security and where to spend it are very difficult to answer in the absence of a way to define a target risk posture and assess the current risk posture of enterprise systems.[19]

  The target risk posture may not only be defined within an individual organization, but also may be influenced or led by overarching guidance or requirements based

[19] The risk posture of a set of information resources (e.g., a system, a network, the data and applications needed to perform a given business function) is defined in terms of the threats and undesirable consequences that concern the enterprise. A risk posture can consist of a single value (level of risk), or of a set of values (e.g., level of vulnerability to hackers, level of vulnerability to insider threats, likelihood of extended unavailability, likelihood of loss of confidential data). Problems of how to define a risk posture (i.e., of how to model risk) and how to translate technical risks into business risks, are addressed in Section 3.7, "Metrics and Models."

on industry-specific best practices or regulations. Research is therefore needed into a common language, definitions, and assessment methods that an enterprise can use to demonstrate adherence to applicable guidelines and regulations.

The trend toward increased interdependencies among systems, networks, and other capabilities owned by different organizations is expected to continue. It is extremely difficult to determine whether connectivity to an external organization poses an acceptable risk and which security controls are needed to provide adequate protection. The critical aspects of this problem are not just the challenges of defining technical risks, but also the challenges associated with translating technical risks into "business" or "organizational" risks, since it is at the organizational level that the risk management decisions will be made. [20]

Therefore, research is needed to develop appropriate, useful, and effective methodologies to assess an enterprise's risk posture that take into account both technical and non-technical aspects of the environment. New capabilities must be able to translate risk postures into specific procedural, architectural, and technical requirements for widely diverse environments, and to be able to do this even as these environments change over time. Furthermore, research is required into capabilities to assess "as-built" environments in real time such that the current enterprise risk posture can be determined or the adherence to defined risk management strategies can be assessed.

- Definition of and protection at security perimeters

  Historically, a security perimeter—the boundary within which one security policy can be enforced, and beyond which the security policies and enforcement mechanisms can only be assumed—was defined by one or more communications or network interfaces, specifically the interfaces between the organization's network and some external (untrusted) network. This perimeter model is of decreasing validity, as organizations have become more geographically dispersed, permitted remote user access into their networks and applications, and become more interconnected and integrated with customers and partner organizations. Essentially, the perimeter must be defined not only at the IP layer, but at all layers. The security perimeter, however, may be different at different layers; for example, the operating system of a desktop computer controls a different set of resources than, and defines its security perimeter differently from, a distributed application. Therefore, research is needed into new models for enterprise protection. Specific challenges include how to protect against intrusions and limit damage when internal systems are externally accessible, and how to define architectural alternatives and defense-in-depth strategies.

[20] Enterprise risk management involves identifying, assessing, and managing the risks that the enterprise will be unable to perform its stated mission. Cyber security risks are but one class of enterprise risks. Security risk management strategies include risk acceptance, risk transfer (e.g., via insurance), risk mitigation (e.g., by applying more stringent access controls), and risk avoidance (e.g., by refusing to allow connectivity to another system or network).

Progress in these areas will require expertise in information management and security technologies, as well as an understanding of policy requirements, business models, and organizational processes. Information infrastructure protection-oriented research on these issues will emphasize defining and managing security relationships among the individual enterprise, its partner organizations, and information infrastructure providers (e.g., trust relationships, the types of enterprise security management information that must be shared, and the externalization or assumption of risk).

3.2 Trust Among Distributed Autonomous Parties

Brief Problem Description

The information infrastructure enables and relies upon interactions among diverse parties, including organizations, systems, individuals, and devices, from mobile phones to desktop computers. The decisions such parties make about what interactions to allow are predicated upon trust relationships.[21] As the nature and use of the information infrastructure evolve, so do needs for defining, establishing, and enforcing trust relationships. In the absence of models and technologies for trust among entities that are organizationally or geographically distributed and largely autonomous, interactions may be based on wishful thinking disguised as poorly articulated assumptions.

In a trust relationship, each party, explicitly or (more typically) implicitly, asserts its identity and describes both the behaviors it can be expected to exhibit (e.g., a service provider claims that it will maintain the confidentiality of the information it receives) and the behaviors it can be expected not to exhibit (e.g., the provider promises it will not launch a denial-of-service attack). Each party also describes the behaviors which it requires or expects of other parties and which behaviors it precludes, and may state conditions on the identity of parties with which it will interact; in effect, each party asserts its policy for participating in an interaction.[22] The IT marketplace has been evolving trust models—which resources the owner seeks to control, which entities can seek to use those resources, on what basis decisions are made—ever since interactive computing made its way into commercial arenas. With the increased interconnectivity and interdependence that mark the information infrastructure, trust models need to include conditions on transitivity and delegation of trust.[23]

[21] Trust relationships in cyberspace reflect those among individuals and organizations. In a trust relationship, each party states which uses it will allow the other party (or parties) to make of its information resources. For example, an organization can specify which data a business partner may access; a network service provider can specify, for a customer, which communications channels it will make available. In a trust relationship, each party also states what information it will rely on the other party or parties to provide. For example, one system can rely on another to identify and authenticate a user.

[22] Possible behaviors depend on the nature of the entity; for example, an individual or a process acting on an individual's behalf can seek access to resources controlled by a system; a system can control a device; and a device can provide data to another device. For example, a system may grant a user remote access to the user's e-mail only if the user makes the request over a secure channel, and it may require a mobile code to be digitally signed.

[23] As an example of transitivity, if one system accepts the identification and authentication (I&A) of users provided by a second, and the second system accepts user I&A from a third, then the first system in effect

> "When everything happens at once, wide and fast moving problems simply route around any central authority. Therefore overall governance must arise from the most humble interdependent acts done locally in parallel, and not from a central command."
> Kevin Kelly, *Out of Control*

Established techniques largely rely on a central authority and relatively static relationships. Models and technologies that do not rely on a previously determined trusted third party are needed for trust in dynamic environments, in which autonomous parties previously unknown to one another need to decide whether and how to interact. Such parties are distributed across organizations, physical locations, and (in the case of software acting on behalf of an individual or organization) layers or components in system architectures.

Existing Research and Capabilities

Because trust relationships involve identification and expectations, considerable R&D effort has focused on identification and authentication of distributed parties, in both homogeneous and heterogeneous system architectures.

For homogeneous environments, expectations are defined and managed by internal policy definition and security administration capabilities. Current solutions rely upon either a central authority or a previously determined trusted third party—an organization or information infrastructure component on which others can rely to provide security-related information or services. Enterprise products provide methods of authentication and determination of allowable behavior in a homogeneous distributed environment, through a recognized central domain authority that serves as a common internal organizational reference point or base level service for determining those principles. A Public Key Infrastructure (PKI) extends identification and authentication to heterogeneous components within an administrative boundary: a third party, the certificate authority (CA), authenticates each participant in a transaction. (That third party is trusted to have properly validated the identity of the individual or organization to which it issued a certificate, which is used for authentication.) Static trust relationships between CAs enable recognition of entities across administrative boundaries.

R&D investment continues to address problems associated with existing solutions, including scalability (how well a solution works as the problem size increases), consistent implementation and management across the enterprise, integration, and ease of use. Existing research also addresses some of the problems associated with defining and communicating expectations. The focus is on specific authentication mechanisms and on authorization to access data, use resources, or participate in a transaction. Specifically, existing research has developed ways to determine how authorizations can be encoded in

trusts the I&A of the third. Transitivity is often implicit: if one organization relies on a second to keep shared information confidential, and the second shares some information resources with a third, then the first organization trusts the third not to exploit its relationship with the second to obtain that confidential information. In turn, one process in a transaction can provide its authorization to access data to a second process; it thereby delegates its authority to act.

digital certificates or other data structures, how they can be revoked in a timely and predictable way, and what types of behavior can or should involve authorization.[24]

Moving beyond existing solutions, current research includes peer-to-peer (P2P) trust models, the dynamic establishment of trust, the revocation of trust between autonomous parties, and data validation and trust in the absence of trusted peers. The relationships being investigated include those between individuals (e.g., classic P2P information sharing), between individuals and enterprises (e.g., e-commerce), and between enterprises (e.g., business-to-business). The emphasis of ongoing research remains on authorization, rather than on defining and communicating expectations regarding other security-related activities, particularly monitoring, auditing, and correlation of behavior traceable to a single individual or organization.

I3P Research Areas

Existing research, particularly in the areas of dynamic trust management, security for grid computing, and P2P trust models, is expected to address many of the needs articulated by enterprise users. However, solutions are needed that address the autonomy, scale, complexity, and dynamism of critical infrastructure sectors. A sector can have many co-equal and independent organizational components, such as a federation of regional banks participating in a transaction-clearing network. A public utility infrastructure can have multiple regional domains, dynamic constituent elements or member utility companies, and tens of thousands of Supervisory Control and Data Acquisition (SCADA) devices. Relationships can be highly dynamic and established under conditions that stress both systems and the organizations that use them, as in the case of emergency response.

This research area touches all aspects of trust among diverse and autonomous parties, including identification; authentication (which could be bound to attributes other than identity); authorization of actions that span security management domains; dynamic negotiation, enforcement, and demonstration of enforcement of security-related agreements; trade-offs among security and functional objectives important to different parties (e.g., privacy vs. convenience); definition and management of roles and responsibilities across participating elements in an infrastructure; and definition and revocation of trust relationships in an increasingly dynamic information infrastructure.

The I3P gap analysis indicates that new or additional R&D in the following areas would have high leverage in improving the state of information infrastructure protection:

- Trust models for distributed autonomous parties

  The problem of establishing and maintaining trust in the absence of a previously determined trusted third party presents conceptual, technical, and social challenges. Conceptual challenges include determining an adequate basis for trust between parties in a given environment (e.g., via reputation, via negotiation to

[24] The Security Assertion Markup Language (SAML) provides a framework, based on the eXtensible Markup Language (XML), for exchanging authentication and authorization information.

agree on a trusted third party); identifying behaviors of concern (e.g., access to system resources, modification of information, issuance of control instructions to a device, monitoring, sharing of personally identifiable information); and specifying attributes of a party that are important for establishing trust (e.g., identity, location, history). Other questions include: To what extent should mistrust be assumed? What are the limits of a given trust model?

Technical challenges include defining new protocols that embody policies; identifying methods and mechanisms by which agreements on allowable behavior can be automatically negotiated, enforced, and their enforcement demonstrated; identifying techniques for managing dynamic changes in agreements; managing trust inheritance or delegation; and managing trust revocation.

Social issues for which research is needed to increase confidence in the information infrastructure include exploration of the implications of automated negotiation of trust relationships: under what circumstances might different approaches be acceptable to individual and enterprise users? Research is needed into the application, adaptation, and extension of models of trust in the social sciences—how trust is established and maintained, who trusts whom and why, and levels or nuances of trust.

- Dynamic security relationships in P2P settings

  The need to address the general trust model problems described above is particularly acute for P2P settings: computation, networking, and information sharing. P2P relationships are becoming more popular, but they increase the vulnerability of the computing platforms or mobile devices that participate in them. Specific challenges include defining a lowest common denominator for trust negotiation and defining stable, converging protocols that allow a number of parties to arrive at the right level of trust (or rejection of trust). Exploration of alternatives for an initial trust negotiation—how much information must each peer provide, and what information (if any) must be obtained about or from the environment (e.g., geolocation)—is also needed. Examples of other questions include: how can established security agreements or overall security policies be used to guide or restrict the trust relationships negotiated between peers? What is required to add 1-to-n peers to an existing trust relationship? To what extent does the lack of a predetermined trusted third party reduce or restrict the possible trust outcomes in a P2P approach, and how can or cannot a P2P approach be mixed with a traditional third-party authority? How can trust relationships be modified in response to changes in the environment (e.g., decreased bandwidth)?

- Devices as parties in dynamic trust relationships

  Considerable attention has been paid to dynamic trust relationships involving software (e.g., mobile code, mobile agents). Increasingly, devices dynamically interface with, interact with, and disconnect from the information infrastructure.

The types of devices are increasingly diverse, including consumer items such as mobile phones and personal digital assistants, enterprise resources such as laptops, and special-purpose devices such as sensors and process control devices. The general trust model problems described above must be addressed in the more specific context of devices, both those that are infrastructure components and those that transiently interact with such components. Questions such as how much intelligence is needed within the device to negotiate and enforce trust relationships, how much mistrust is needed, and how other infrastructure components can protect devices which (for reasons of cost or power) cannot protect themselves, apply to process control devices such as SCADA.

- Establishing trust in data

  In an increasingly distributed, peer-oriented environment, the need to migrate trust from individual parties (whether components, users, or software) to the data itself is increasing. Can software-oriented techniques such as proof-carrying code be extended to data?[25] A key question is under what conditions, and how, can the accuracy, correctness, and absence of malice in a set of presented data be assessed without having established a security context for the source? Research is needed into techniques that provide communications path independence, so that data maintains its integrity and trust level even in the absence of a context of how it arrived. (For example, in public utility infrastructures, techniques must be able to determine the trustworthiness of readings from a remote sensor either when suspecting compromise in an intervening component or, because of system failures, when retrieving the data through non-typical means from the remote environment.)

3.3 Discovery and Analysis of Security Properties and Vulnerabilities

Brief Problem Description

Systems being developed and deployed today are rife with vulnerabilities and poorly understood security properties.[26] Users are frustrated by an unending stream of security alerts and patches, the effects of which are not always understood or, in the context of an as-used system, benign. The information infrastructure is subject to a constant barrage of attacks that result in operational and financial costs. New approaches and technologies are needed to determine, throughout all stages of the life cycle (development, integration, update and maintenance, decommissioning, or replacement of components), whether exploitable defects have been introduced or unanticipated security properties (e.g., restrictions on what a long-used application can do) have been introduced or have

[25] Proof-carrying code enables an application to demonstrate that it will enforce its stated security policy; recipients can use that application without regard to the level of trust they place in the distributor.

[26] The security properties of a system or component include how, and how well, it achieves the traditional security objectives of confidentiality, integrity, availability, accountability, and nonrepudiation. The security properties of a component or system also include how, and the extent to which, it is vulnerable to a given threat or class of threats, as well as what security policies it is capable or incapable of enforcing.

escalated. Most urgently needed are methodologies and tools to analyze source code, object code, and as-integrated systems to identify exploitable vulnerabilities in software.

The information infrastructure has a large number and variety of components, in different forms: hardware, firmware, software, communications media, storage media, and information. Frequently, the properties of those components are poorly understood, due to undocumented functionality, flaws in their design or implementation, or unanticipated uses. Moreover, the properties of a system or subsystem cannot be derived or deduced from those of its components, and emergent properties of large-scale systems and systems-of-systems are difficult to describe, much less predict.[27]

The scope of the problem should not be underestimated. First, computer products being marketed today have numerous software vulnerabilities.[28] Often, no sooner has an operating system or application been released than a patch or update is announced. The magnitude of this issue is so large, and has become so culturally ingrained, that an entire industry has formed around patch management. Yet, software vendors, both large and small, have recognized that there are both financial and competitive advantages to be reaped from marketing more secure software.

> "Let's acknowledge a sad truth about software: any code of significant scope and power will have bugs in it."
> Steve Ballmer, *Connecting with Customers*,
> http://microsoft.com/mscorp/execmail/2002/10-02customers.asp, 2 October 2002

Second, the existence, and exploitation, of undocumented functionality is a problem. Such functionality has been found in processor chips, networking and operating system software, and applications that are central to the information infrastructure.[29] Undocumented functionality frequently constitutes a security vulnerability: it can include backdoors, exploitable code, time bombs, or other malicious capabilities.[30]

Third, this problem is exacerbated by the inclusion of hardware and software components from multiple sources into products: vendors routinely integrate third-party products but cannot easily determine whether a component includes vulnerabilities, trapdoors, or malicious code.[31] In large systems of systems, the sheer complexity of the components, interdependencies among functions, and failure to understand the desired behavior result in flaws. Indeed, due to interdependencies and interactions, the security properties of a

[27] See, for example, *Emergent Algorithms: A New Method for Enhancing Survivability in Unbounded Systems* (http://www.cert.org/archive/html/emergent-algor.html).

[28] Incident statistics are available through the CERT Coordination Center (http://www.cert.org/stats/cert_stats.html); documented vulnerabilities, searchable by platform or software, can be found at http://www.cve.mitre.org.

[29] "Easter eggs"—unadvertised functionality that can be triggered by undocumented user actions—are a common, and often benign, example found in software. See http://www.eggheaven2000.com for a database of Easter eggs.

[30] See Ken Thompson's *Reflections on Trusting Trust* (http://www.acm.org/classics/sep95).

[31] In a classic case in 1991, one company inadvertently sent diskettes infected with the Stoned-3 virus to their own customers (http://www.ciac.org/ciac/bulletins/c-11.shtml).

system or subsystem *cannot* be derived or deduced from those of its components.[32] In addition, the interactions between a system (which might include sensors or control mechanisms) and its environment (which can change dynamically) cannot easily be predicted.

It is evident that today, as the size, complexity, distributed nature of software/hardware development, and the degree of integration of systems increases, new or significantly enhanced capabilities are required to analyze hardware and software systems to identify vulnerabilities before the systems are used operationally. This need will become even greater in the upcoming years.

Existing Research and Capabilities

Not surprisingly, considerable research effort has been applied to the problem of ensuring the presence of desired security properties and preventing (or determining the presence of) vulnerabilities. Most rigorously, the study of formal methods for ensuring that hardware and software implementations conform to stated expectations has long been an active field of research.[33] Yet, while formal methods have had successes with focused capabilities, their broad-based applicability is limited, and, thus far, they have not proved practical for large systems or systems of systems. Current research includes techniques for automatically generating test suites from specification of required or unacceptable behavior, and for analyzing the security properties of cryptographic protocol specifications.

Many vulnerabilities result from coding errors. Research intended to prevent, or decrease the likelihood of, vulnerabilities in code includes the development of "safe" programming languages and automatic generation of code from specifications.

Current research seeks to develop methodologies and tools to identify and analyze malicious code. While the virus detection market seeks particular signatures within code, research is still in its infancy for the more general problems of detection of polymorphic malicious code (code that changes to avoid fitting a known profile) and of vulnerability identification in source or object code. Current research on source code analysis includes exploration of pattern matching, feature extraction, and code slicing analysis techniques; for object code, comparative analyses, and disassembly-based techniques are being explored. Though some progress is being made, no capability with the completeness and quality (e.g., robustness, scalability) required to examine large code bases exists.

To detect vulnerabilities in as-built and as-used systems, a wide range of scanning tools examines externally presented interfaces. The Security Administrator Tool for Analyzing Networks (SATAN), written in 1995, was one of the first comprehensive vulnerability

[32] This is true even when the security properties of those components are well understood. See *Turning Multiple Evaluated Products Into Trusted Systems*, National Computer Security Center (NCSC) Technical Report-003, Library No. S-241,353, July 1994 (http://www.radium.ncsc.mil/tpep/library/rainbow/NCSC-TR-003.pdf).

[33] A compendium of recent activity may be found at http://www.afm.sbu.ac.uk/meetings.

scanners; today, there are many host-based and network-based scanners that test for a wide range of vulnerabilities.[34] In addition, commercial and freeware software and system configuration checkers are available. In both cases, however, the tools' effectiveness is limited in that they are not based on a comprehensive analysis strategy, they look for known vulnerabilities but have no way to discover other defects, and they cannot scale to large, complex systems. More promising is the use of a "Red Team," an independent, simulated "enemy force" contracted to identify vulnerabilities, including previously unknown ones, in systems and their operational environments.

I3P Research Areas

Needs remain across the board for the analysis of hardware and software components and systems to identify vulnerabilities and/or malicious code.[35] To be useful, identification and analysis methodologies must also be embodied in readily usable tools; meaningful research should demonstrate the feasibility of a tool and lead to determination of tool requirements. The I3P gap analysis indicates that new or additional R&D in the following areas would have high leverage in improving the state of information infrastructure protection:

- Code scanning tools and techniques

  Research in capabilities to identify vulnerabilities in code is an immediate and critical need. This need can be refined into two sub-categories:

  – Source code-scanning tools

    Tools are required that can scan source code and identify potential vulnerabilities. These tools would be used by code developers as a quality assurance mechanism, by system integrators, and by users who have access to source code (either through purchase agreements or because the code is open source). Given that enterprises develop in-house applications and extensions to purchased software, freely available or low-cost tools could have a broad impact. Such tools could aid in identifying vulnerabilities in legacy and evolving enterprise systems, enable small companies developing innovative software to address security better, and support the open source community in providing more secure software, thereby raising the bar for software quality.

    The tools must at a minimum be able to identify, and if possible remove, common sources of vulnerabilities (e.g., susceptibility to buffer overflow attacks) and must work on large-scale software systems (e.g., the source code base for a commercial operating system or database management system).

[34] Although somewhat dated (January 2001), *Network Computing*'s article on "Vulnerability Assessment Scanners" (http://www.networkcomputing.com/1201/1201f1b1.html) provides a good overview.

[35] More precisely, the need is to identify the security properties of components and systems. A threat model or an enterprise security policy is needed to determine whether an attribute of code, a component, or a system constitutes a vulnerability.

The tools must be automated and must provide sufficient inherent capability to be useful to non-experts.

- Object code-scanning tools

  The research needs (e.g., vulnerability identification, scalability, ease of use) are similar to those noted above, except that they apply to object code. Such tools would be applied later in the product's life cycle, either by system integrators or by end users.

  This research must produce the analysis science and methodology, as well as proof-of-concept prototypes to validate the concepts, and explore scalability, performance, usability, and other parameters of practical importance.

- Device-scanning tools

  Software is not the only source of vulnerabilities or unexpected functionality. Research is needed into automated analyses of other elements of the information infrastructure computing components, such as hardware, firmware, communications media, and storage media.

- Discovery and analysis methodologies

  In addition to specific code and device-scanning techniques and implementations, broader methodologies for discovery and analysis of security properties are needed.[36] Examples of specific areas in which more general methodologies need to be applied include discovery and analysis of control and configuration settings, analysis of protocols, discovery of system security properties, and behavioral scanning of systems.[37]

  In addition, a broader methodology must provide a framework for addressing composability issues. (For example, while two software or other components, when evaluated independently, may be free of vulnerabilities, together they may open easily exploited vulnerabilities.) The methodology, then, must (1) address composition of like components (such as security functionality overlap or duplication, conflicts between components which degrade or obviate the security functionality offered by individual components), (2) address cross-element analyses (such as hardware dependencies embedded in software), (3) enable development of a minimal security configuration, and (4) be inclusive of higher

[36] In part, a discovery and analysis methodology could be used to formulate metrics (to address questions such as how do differing vulnerability scanning implementations compare, and what is the relative "susceptibility to vulnerabilities" of two comparable code bases). See Section 3.7, "Models and Metrics."

[37] The need is particularly acute to analyze the security properties of wireless protocols; see Section 3.6. Techniques for discovery of system security properties support Enterprise Security Management; see Section 3.1. Behavioral scanning—observation of how the system is being used—can serve as an indicator that a previously unknown vulnerability is being exploited. It also supports Secure System and Network Response and Recovery; see Section 3.4.

order constructs (such as policy interaction, administration, and architecture/design flaws) to identify vulnerabilities at a higher level of abstraction.

3.4 Secure System and Network Response and Recovery

Brief Problem Description

The proliferation of number and types of computing and communications devices has resulted in increasing size and complexity of the information infrastructure. Critical infrastructure sectors rely on these systems-of-systems, including some systems belonging to sector organizations and others (e.g., networks) belonging to private service providers or government.[38] Even enterprise systems are frequently loose federations of systems managed by different operating centers. Response to and recovery from[39] attacks against such complex systems are hindered by this inherent complexity, making reconstitution to a secure state extremely difficult.[40] Simply understanding the effects of a proposed responsive action is increasingly difficult for enterprise security managers. Response across a set of organizations is often uneven and difficult to coordinate.

Both the operational and research communities have long recognized the need for response and recovery at the system or enterprise level. Due to the enterprise focus, disaster recovery or continuity-of-operations planning was long distinct from cyber security as a discipline. The cyber security community's R&D emphasis to date has primarily been on information protection, with research focused on detection. There has been a marked lack of progress on secure response and recovery for complex and heterogeneous systems, but from an infrastructure perspective, these capabilities carry a high level of importance.

During response and recovery, system behavior is unpredictable and difficult to manage. In current architectures, response and recovery activities largely rely on the (compromised) networks, enabling an attacker to monitor those activities. Thus, periods of degraded or disrupted operations constitute a significant window of vulnerability to further attacks and particularly to insertion of malicious code or data.

[38] A system-of-systems is a collection of individual systems, owned or administered by different organizations, or acquired separately or at different times, that support a common mission, and thus are treated as a single entity for purposes of systems engineering or risk management. The component systems can include networks; thus, the phrase "system-of-systems" includes "systems of networks" and "networks of networks."

[39] Response is an action or set of actions triggered by an indication of a possible intrusion or abuse of system resources (e.g., insider malfeasance). Recovery is the process of taking a system from an unacceptable level of performance to some minimum acceptable level. Reconstitution is the process of taking a system from an unacceptable or minimally acceptable level of performance to full performance.

[40] While diversity, and resultant complexity, introduce many difficulties, uniformity or homogeneity introduces others.

Existing Research and Capabilities

Increasing interest in survivability and in making IDSs less reactive has driven further research into secure response and recovery. Major thrusts are to make enterprise systems and networks more proactive and capable in the face of potential attack; to identify suspicious activities before they have an impact; to perform situational assessments and to provide a picture of the status of enterprise systems that can be understood by decision makers; to tolerate intrusions and continue to operate in the presence of ongoing attacks; and to respond in a more timely and effective manner.

Current detection capabilities provide data that often indicate an intrusion has occurred or provide indicators that cannot be acted upon early enough to thwart an intrusion. Current research is investigating predictive techniques that can provide early warnings of impending attacks and that may allow attacks to be detected and stopped before they cause actual harm. Data-mining techniques are increasingly used. Current research also includes specification-based intrusion or anomaly detection, in which the behavior of an application or component is specified and monitored, with variances from specified behavior indicating potential abuse or intrusion.

Situation awareness and security event data management are also crucial concerns. With the prevalence of proprietary security mechanisms and the lack of standards associated with vulnerability analysis and monitoring/detection capabilities, determining the overall security status within an organization has proven to be virtually impossible. Establishing a "Common Operating Picture" for security would allow organizations to prioritize resource assignments, respond to critical attacks more effectively, and limit damage from malicious activities. Associated with creating this Common Operating Picture are the following issues:

- Correlation between detectable (and detected) events and activities that are deemed to be threatening;
- Scalability of monitoring and detection capabilities when the number and types of sensors continue to rise;
- De-duplication, event correlation, and analysis to reduce false positives such that trust is established in generated alerts; and
- Large-scale data collection, mining, and analysis to identify complex threat actions (e.g., those taking place over relatively long timeframes) that cannot be detected via real-time monitoring.

The application of ongoing Department of Defense research into a Common Operating Picture to critical infrastructure sectors is problematic, given the diversity and lack of hierarchical control over sector systems.

The major research focus is on recovery from damaging impacts, and specifically on the reconstitution of operational capabilities, for the system, network, and enterprise. Such research addresses not only functionality but also architectures in which systems can automatically degrade to operational levels that allow critical capabilities to remain intact while recovery operations take place. This includes research into survivable architectures, as well as into intrusion-tolerant systems that allow for the degradation or destruction of certain capabilities, while ensuring that critical functionality remains available. The current focus is on recovering operational functions, rather than on using distributed security capabilities to ensure that degraded operations still maintain an acceptable degree of risk.

Current projects in the private, government, and academic sectors include research into autonomic systems—systems that can sense and reason about their internal components and state—and "recovery-oriented computing."[41] Component capabilities have begun to be developed, including self-evolving systems that can monitor themselves and adapt to some change. This research is still in the early stages, with a focus on resource allocation and hardware and firmware failure.

I3P Research Areas

While response and recovery is an active research area, current research does not incorporate attack scenarios or security consequences and is focused at the system/network or enterprise levels. Thus, current research does not address the issues of scale, coordination across different administrative and policy domains, or coordination across highly diverse systems that are the hallmarks of information infrastructure protection.

The I3P gap analysis indicates that new or additional R&D in the following areas would have high leverage in improving the state of information infrastructure protection:

- Prediction/pre-incident detection

  Research is needed into techniques that can complement or supercede current pattern recognition (e.g., signature-based) approaches for intrusion and anomaly detection, particularly for complex systems. For a comprehensive and understandable picture of the status of enterprise systems, further research is needed not only into the correlation, analysis, and presentation of monitoring or status data from diverse components (network, host, application, device), but also into a single intrusion/anomaly model, how an intrusion detection system can monitor the status of its sensors (e.g., has a sensor been blinded), which observable behaviors should alter the monitoring strategy, and the impacts of intrusion detection mechanisms on the system (e.g., bandwidth demands).

  A particular challenge to existing detection strategies is the increased use of encryption throughout the protocol stack. Research is needed into techniques (e.g., intelligent traffic analysis) for detecting suspicious activities.

[41] See, for example, *Autonomic Computing: IBM's Perspective on the State of Information Technology*, IBM Corporation (2001), available at http://www.research.ibm.com/autonomic.

In addition, research is particularly needed into detection of actions that indicate preparation for a large-scale attack. Key components of such research include Internet-scale attack models, tools to look at shared data about anomalous behavior from different perspectives (e.g., by critical infrastructure sector, by geography) and to determine the footprint of an attack, and modeling and simulation to understand attack goals and intent. Research is also needed into techniques for sharing information about suspicious behavior outside the enterprise without increasing the exposure of enterprise systems; identification of information about suspicious behavior that can be garnered by information infrastructure providers (e.g., telecommunications providers); and techniques for aggregating and analyzing such information that respect privacy concerns.

- Recovery and reconstitution for systems-of-systems

  Current and planned research into intrusion-tolerant, self-healing, context-aware, or self-stabilizing systems needs to be extended and applied to systems-of-systems. Research challenges include definition and comparative analysis of different response strategies (e.g., notification within and beyond the enterprise, damage limitation, containment, and observation of attacker activities). In particular, models and decision support tools are needed to help manage tradeoffs between recovery of system capabilities as quickly as possible and reconstitution to a secure state. Autonomic systems require incorporation of a wide range of capabilities, including system, network, and environmental awareness; detection of potential threats; error detection and anticipation of failures; and sophisticated, complex logic and decision tools.

  Finally, with operating environments as complex as they are, and the fact that critical operations may be shared across organizations, it has become extremely difficult for procedural (human) processes to be relied upon as reconstitution and recovery mechanisms. Research is needed into ways to automate these functions in a manner that re-establishes functionality, ensures security is maintained, and can be trusted by those responsible for administering and managing these environments. Research is also needed into better ways to present information about system status and the implications of different courses of action to decision-makers.

3.5 Traceback, Identification, and Forensics

Brief Problem Description

During and following an attack, organizations (e.g., infrastructure users or owners, law enforcement, military services) must have prompt and reliable information regarding what the attack is and who is launching it to determine and pursue an appropriate organizational response (e.g., notification of other organizations, containment, mitigation, system reconstitution, investigation and prosecution, military action). This information must be consistently accurate in the face of perishable data, falsified data, and other obfuscation techniques.[42] A particularly urgent need exists, then, for capabilities to identify the source location of the attack; to identify the individual, group, or organization originating the attack; to determine the actual nature of the attack; and to maintain evidence that can be used to justify the chosen response.

Effective response to incidents must be predicated on the most complete analysis and understanding of the incident possible within existing time constraints. The analytic goal is to provide sufficiently complete reconstruction of the event (including answering questions about who, what, where, when, why, and how) to motivate an appropriate response consistent with the source, intent, and consequences of the incident. The data foundation for such analysis includes both locally available data (e.g., local logs and other incident artifacts) as well as data actively collected during and after the incident from external—and possibly foreign—sources.

The nature of the technical challenges mandates an infrastructure-supported solution. That is, the data available at end systems is generally insufficient to perform adequate analysis; data from the infrastructure(s) involved is often also needed. For example, an incident on an end system reached via dialup over the local telephone system will require support from the telecommunications service provider to identify the point of origin. Moreover, viable solutions will often require coordination and collaboration to collect and correlate data since significant incidents will likely transcend individual infrastructures and technologies. For example, an attack on a SCADA system will impact the end system itself, but may also involve Internet Service Provider (ISP)-owned networks, intermediate publicly or privately owned computing devices, long-haul telephone communications networks, wireless access points, and so forth.

Existing Research and Capabilities

Current capabilities are oriented toward enabling the enterprise to detect and respond internally to suspected attacks. For example, incident identification is relatively mature and is the beneficiary of significant ongoing investment. The IDS industry is well established and is developing and evolving technologies to identify attacks on systems as they happen. Today's IDS technology solutions span the range from host- to network-based techniques, from use of attack signatures to statistical-analysis-based anomaly detection, to hybrid systems intended for enterprise-wide protection.[43] Numerous IDS articles and conference papers are published each year, helping to push the state of the art.[44] The virus detection industry is equally robust and active in technology development, with numerous products and relevant conferences.[45]

[42] A spoofed source computer network address is an example of falsified data.

[43] Numerous products exist, with several Web-based resources identifying capabilities. The list maintained by the Center for Education and Research in Information Assurance and Security (CERIAS) provides a representative snapshot of the intrusion detection software landscape today (http://www.cerias.purdue.edu/coast/intrusion-detection/ids.html).

[44] See, for example, *Recent Advances in Intrusion Detection* (http://www.raid-symposium.org).

[45] Numerous products exist, with several Web-based resources identifying different capabilities. The ICSA-maintained list provides a representative snapshot of the virus detection software landscape today (http://www.icsalabs.com/html/communities/antivirus/certifiedproducts.shtml). The Virus Bulletin

Some research has been performed specifically for incident origination and incident characterization. For example, in the area of traceback (to identify the source of an incident), some research has been done, but it has been exploratory in nature and has not yet yielded concrete viable solutions.[46] The Internet Engineering Task Force (IETF) started a working group to develop a standard for backwards tracing, but to date the group has produced neither an Internet Draft nor a Request for Comments (RFC) and is presently dormant.[47] Computer forensics (i.e., the use of forensic techniques to uncover evidence from computer systems or networks) is a well-established discipline with ongoing technology development across the community as well as a range of existing commercial tools.[48] But these tools tend to analyze individual systems, with an emphasis on data collection and analysis. In terms of providing an enterprise-wide holistic approach to data collection, synthesis, analysis, and results production, current tools are lacking.

The focus herein is on the especially difficult technical hurdles which industry and the research community have not addressed with vigor but that offer large returns with respect to satisfying the nation's infrastructure protection needs. These specific areas were alluded to repeatedly, though obliquely, throughout the input received via the I3P data collection mechanism—frequently, the input assumed the form of concern over the cost of defensive mechanisms, incident response options, liability and insurance issues, and so forth. As expected, law enforcement personnel most succinctly identified the specific gaps.[49] Companion research into the legal and policy implications of new capabilities in this area is required.

I3P Research Areas

The I3P gap analysis indicates that new or additional R&D in the following areas would have high leverage in improving the state of information infrastructure protection:

- Incident origin

  The fundamental questions are: when and where—from what device, system, and geographic location—did the incident originate? Most simplistically and most naïvely, these questions refer to identifying the origin of the IP packets received

Conference (http://www.virusbtn.com/conference/overview/index.xml) is an example of a forum for presenting research results.

[46] See, for example: http://www.cs.washington.edu/homes/savage/traceback.html or http://www.perrig.net/~dawnsong/papers/iptrace.pdf.

[47] See http://www.ietf.org/html.charters/itrace-charter.html.

[48] The community includes the Department of Defense Computer Forensics Laboratory (http://www.dcfl.gov/) and the Computer Forensics Research Development Center at Utica College (www.ecii.edu/cfrdc.html). Examples of products include the Coroner's Toolkit (http://www.porcupine.org/forensics/tct.html), ForensiX (http://www.all.net/ForensiX/), and EnCase (http://www.guidancesoftware.com/).

[49] See *Law Enforcement Tools and Technologies for Investigating Cyber Attacks: A National Needs Assessment* (http://www.ists.dartmouth.edu/lep/lena.htm).

by an end system, given that the packet addresses can be spoofed.[50] A complete solution, however, must address such additional complications as use of intermediate systems that serve to mask the true originating system (such as previously compromised hosts, anonymizers, and the like); use of non-IP-based intermediate protocols; and cross-infrastructure paths (for example, paths that transition from wireless—802.11a or 802.11b—networks, to cellular, to traditional land-line telephone networks, to Ethernet networks, and so forth). Research is also needed into techniques for assigning a trustworthy timestamp to events in the chain(s) of events in the incident.

- Originator identity

  Identifying the physical incident origin does not always definitively identify the responsible individual, group, or organization. For example, the source may be a system accessible by multiple individuals, a mobile system with no link to an individual user, or a device of unknown provenance. Thus, analyses to identify incident origin must be augmented by techniques to determine attacker identities.

  The use of criminal profiling by law enforcement agencies is well established and effective. The application to computer security has long been pursued in the intrusion detection arena with varying success.[51] However, the macro-level application of such techniques to traffic patterns (e.g., volume, distribution, timing), technology applications and tool signatures, incident artifacts, and other observable manifestations is a significant gap that has yet to be addressed.

- Originator actions

  Computer forensics is a developing field. Nevertheless, significant gaps have been identified in capabilities that can handle encrypted data, steganographically encoded data (data hidden within larger data, such as a message hidden within an image), and multi-source log data; tools to retrieve, store, and analyze large media devices; attack-specific analysis tools; and timeline reconstruction.[52] In particular, the capability to collaborate across organizations and to correlate data from different sources, in order to reconstruct and analyze attacks, is repeatedly highlighted as a need.[53]

Ideal solutions would be (1) fully integrated across functional boundaries—for example, linking enterprise security functions into incident detection, response, and forensics functions, to produce data suitable for legal or military prosecution; (2) fully integrated

[50] For details on the IP, see RFC 791 at http://www.ietf.org/rfc/rfc0791.txt?number=791.
[51] For example, see the seminal paper "The SRI IDES Statistical Anomaly Detector" published in the *Proceedings of the IEEE Symposium on Research in Security and Privacy*, pages 316-326, 1991.
[52] See *Law Enforcement Tools and Technologies for Investigating Cyber Attacks: A National Needs Assessment* (http://www.ists.dartmouth.edu/lep/lena.htm). Capabilities for establishing trust in data, as discussed in Section 3.2, are central to more effective computer forensics as well as to establishing incident origin and originator identity.
[53] Attack reconstruction can be used to validate or improve attack models, as discussed in Section 3.4.

across infrastructures—for example, computer and network solutions integrated with SCADA solutions, telecommunications solutions, wireless solutions, satellite communications solutions, and military signals capabilities; and (3) suitable for use by novices as well as experts. However, the community would be well served by intermediate stand-alone capabilities that address the gaps outlined above.

3.6 Wireless Security

Brief Problem Description

Wireless technologies are increasingly crucial to enterprise systems and across critical infrastructure sectors. Wireless networks include not only wireless telecommunications *per se*, but an increasingly diverse set of end devices, including sensors, process controllers, and information appliances for home and business users. In some cases, end devices may also provide wireless telecommunications services.[54] End devices in wireless networks often have limited resources (e.g., power, processing, storage), thus limiting the security capabilities they can provide. Components in a wireless network may be out of any individual's control or in the hands of a user for whom security is not a concern, limiting the effectiveness of physical or procedural controls. Because of the unique aspects of wireless networks, new vulnerabilities and new security concerns regarding the integrity and confidentiality of wireless networks are emerging. Solutions developed for wired networks may not translate to or be implementable in wireless systems, requiring a dedicated research focus into cyber security for wireless networks.

The need for improved security in wireless networks was stated by many of the individuals and organizations involved throughout the R&D Agenda development process, with an emphasis on wireless local area networks (LANs), wide area networks (WANs), and P2P dynamic networking.[55] Wireless networking security concerns were also identified for communications among sensors, process controllers, and information retrieval and analysis platforms. Wireless networking assumes and enables mobility, whether of laptops, personal digital assistants, Web-enabled cell phones, water meter readers, or any of the numerous and inventive ways in which wireless technology is being deployed.

Clearly, wireless technologies are becoming increasingly crucial to enterprise systems and across critical infrastructure sectors. Wired and wireless networks both have many of the same vulnerabilities. However, because of inherent differences between wired point-to-point communications and the dynamic, broadcast nature of wireless networks, solutions developed for wired networks may not be viable in wireless environments. Similarly, aspects of existing vulnerabilities and problems in wired networks may become more difficult or complex in wireless networks, where mobile wireless nodes constantly enter, traverse, and leave wireless networks. A notable example is the

[54] See *Embedded, Everywhere: A Research Agenda for Networked Systems of Embedded Computers* (http://www7.nationalacademies.org/cstb/pub_embedded.html) for a discussion of issues arising for networks that take advantage of wireless telecommunications.
[55] Wireless technologies include radio frequency, infrared, and ultrasonic communications.

emerging challenge associated with the management of policies and services in wireless networks. New advances in configuration management, including intrusion detection and response, security policy management, and policy definition, are needed to address the growing complexity within large-scale, dynamic, ad hoc wireless networks. Current protocols used for managing authentication, such as key management and managing trust relationships, are also insufficient for emerging wireless technologies.

Because wireless nodes cannot be physically isolated for protection, wireless communications are more susceptible than wired communications to a host of security attacks, including disruption (e.g., jamming), observation (e.g., eavesdropping, traffic analysis), and misuse (e.g., theft of service).

Attacks unique to wireless networks include the capture and abuse of control channels, spoofing at or near the boundaries of network cells to capture traffic with mobile units, and direct attacks at the wireless power source. In addition, dynamic wireless networks may also be susceptible to attacks directed at the databases and/or services needed for maintaining configuration and/or security policy management. Established intrusion detection techniques become problematic in the wireless network setting.

Attacks against wireless communications can degrade the integrity of any critical infrastructure that uses them, perhaps most notably emergency response and transportation. Attacks against wireless communications can also threaten individual privacy and enable identity theft. The growing size and complexity of wireless networks complicate issues such as monitoring or understanding intrusions and network health, authentication, and network resiliency, further compounding the need for new research in wireless cyber security.

Existing Research and Capabilities

Current government-sponsored research in wireless security includes authentication, key management, analytic tools for understanding the health of the network, as well as intrusion detection and recovery. Other current research includes the development of scalable, dynamic, self-configuring networks for wireless networking. Research is also currently being done in the areas of encryption, robustness, and new protocols for wireless networks.

Examples of current industry work in this area include diagnostic tools for monitoring 802.11 networks. Work is also being done in industry to embed security into mobile components and hardware to help improve system and information integrity.

Despite the ongoing research, the basic and underlying science of wireless security has not been emphasized. Products and current research have provided incomplete “add-on” security capabilities as opposed to providing security as a fundamental component of wireless networks. Research is needed to extend security across all protocol layers and to embed security into wireless devices.

I3P Research Areas

The I3P gap analysis indicates that new or additional R&D in the following areas would have high leverage in improving the state of information infrastructure protection:

- Security as a fundamental component of wireless networks

  Security concepts and solutions developed for wired systems (e.g., strong encryption) may never be completely applicable to wireless, mobile units because of the disparity between the capabilities of wired and wireless systems due to constraints in areas such as range, processing power, bandwidth, and energy consumption. The asynchrony between wired and wireless networks is expected to be a continuous source of security vulnerabilities.

- Basic science of wireless security

  Development of a scientific and mathematical approach to understanding the security properties of wireless networks is in its infancy at present. Research is needed to identify and understand fundamental properties of wireless network security to enable modeling, measurement, and design of secure, large-scale wireless networks.

- Development of security at the device level for wireless systems

  In particular, research is needed into embedded security for wireless systems and into making wireless security more transparent to the user.

- Research into wireless security at the protocol level

  In the near term, more research needs to be done into understanding the security implications of existing protocols in wireless networks. Mid-term research needs to focus on the development of secure, resilient protocols tailored for wireless networks.

- Integration (and management in a coordinated way) of security mechanisms across all protocol layers

  This area is intended to include all of the layers, from application through presentation, session, transport, network, and data link to the physical layer. Current security protocols reside principally in one layer. Research is needed to provide an integrated, multi-layer defense that would enable improved protection against wireless cyber attacks. This integration would also need to take into consideration the fall-back capabilities (e.g., wired communications) and dependencies.

- Integration of wireless security into larger systems, networks, and systems-of-systems

  Wireless security solutions need to account for mobility, develop new methods for secure distributed authentication, and in the long term provide transparent security solutions.[56] Wireless communications enable more diffuse control of network resources (e.g., Web mesh networks), so that devices whose security is crucial to the network are currently in the hands of individuals who lack expertise or interest in security. Research is needed to improve the embedded security of these systems such that the security of wireless nodes does not depend on the level of security expertise of its users.

- Security situation awareness to permit understanding or visualizing the health of the wireless network at any point in time

  In a wireless environment, the network topology is in constant flux as nodes are intentionally added, moved, or removed. Intermittent connectivity, node and link failures, and compromises must also be detected to characterize the network adequately. In the near term, further investigation is needed into better ways to monitor and represent the status of wireless networks simply to understand their security posture. In the mid term, the focus shifts to dealing with attacks, including intelligent survivability and adaptive connectivity to respond to attacks.

- Addressing DDoS attacks

  Further research is also needed into identification and containment of potentially compromised wireless nodes and countering DDoS attacks. While this is a known and difficult problem in wired networks, the broadcast, dynamic, and mobile nature of wireless networks makes identifying and isolating compromised nodes particularly difficult. Similarly, wireless networks are susceptible to forms of DDoS, including jamming and the ability of a single source to affect multiple wireless nodes simultaneously that are not threats to typical wired networks.

3.7 Metrics and Models

Brief Problem Description

Individuals, organizations, and critical infrastructure sectors bear the risks of relying on the information infrastructure.[57] To accept a given level of risk, to transfer or externalize

[56] Security management for wireless networks (e.g., consistent and timely deployment of security patches) presents particular challenges. See Section 3.1.

[57] Risk can be defined as a combination of several factors: consequence, threat, vulnerability, and safeguards or other mitigation techniques. In the context of information infrastructure protection, consequences include loss of privacy or confidentiality; corruption of data; introduction of unanticipated or undesirable behaviors into information systems; degradation or denial of service; repudiation of, or loss of accountability for, actions taken by individuals, organizations, or system components; and usurpation or misuse of information resources.

risk, or to apply resources to decrease the level of risk to an appropriate balance, a clear and defensible basis founded on rigorous and generally accepted models and metrics for cyber security is necessary.[58] Decision makers lack a foundation of data about the current investment and risk levels; metrics that express the costs, benefits, and impacts of security controls from an economic perspective, technical perspective, and risk perspective; and ways to predict consequences of risk management choices.

Stakeholders repeatedly express needs for metrics and for the conceptual models underlying those metrics to support technology-related decisions. These needs occur at multiple levels, for example:

- Enterprises: Models and metrics are needed to support strategic planning and to inform investment decisions: Which security architecture best supports the enterprise architecture? What are the tradeoffs between security risks and other business risks? Has the level of risk increased or decreased? To what can the change be attributed—changes in usage patterns, new threats, or degraded effectiveness of security mechanisms, for example? What are the tradeoffs associated with different courses of action? How much should be spent on security? How much security is enough to protect the company's interests, and how much is too much?

- Home users: For the home and organizational user, product metrics are needed to inform investment decisions: Which security products are most effective? What are their associated costs?

Stakeholders also identify needs for measures and assessment methods to determine compliance with business policies, regulatory and statutory requirements, and "standards of good practice." Finally, the need is frequently expressed for public policy to be informed by measurements of such risk factors as threat intensity, vulnerability severity, and consequences (in particular, the costs associated with incidents or of vulnerability remediation); investments in security; and the costs and benefits associated with those investments.

Existing Research and Capabilities

Government programs provide approaches to product assessment, compliance assessment, aspects of risk assessment, and assessment of relative investment in cyber security. For example, product assessments are performed against the Common Criteria under the Common Criteria Evaluation and Validation Scheme (CCEVS).[59] The National Institute of Standards and Technology (NIST) provides guidance on capability

[58] The term "metric" is problematic, since it implies a system of measurement and for many connotes scientific rigor (e.g., a metric function in mathematics) or a single value. While other terms are often proposed, "metric" is widely used. In this R&D Agenda, the term "metric" is intended to include partially ordered sets (e.g., red, yellow, green) and vectors, the components of which measure different security-related properties.

[59] See the NIST page at http://csrc.nist.gov/cc/index.html or http://www.commoncriteria.org.

evaluations.[60] In addition, commercial companies routinely perform and publish comparative assessments.

However, the effectiveness of these methods is hindered by insufficient information, inadequate linkage to operational environments, measurement of properties in isolation, and lack of linkage to business and risk models. In addition, measurement processes are frequently ill-defined or cumbersome and lack timelines. Research tends to focus on measurement of specific properties rather than on assessment of system properties from properties of components.[61] The context of information infrastructure protection (e.g., a broad and diverse span of control, limited information sharing, and heterogeneous technology) adds complexity to the development of meaningful metrics and models.

Risk assessment and dependency modeling for cyber security remain in an immature state with only little momentum in the security marketplace. While targeted risk assessment capabilities (e.g., network- and host-based vulnerability scanners) have emerged to address risks associated with component vulnerabilities, products capable of identifying intra- and inter-enterprise risks are generally not available. In recent years, the few tools that were developed and marketed have mostly not succeeded commercially and are no longer available.

I3P Research Areas

The I3P gap analysis indicates that new or additional R&D in the following areas would have high leverage in improving the state of information infrastructure protection:

- Development of a foundation of data to support analyses

  Better information is needed to clarify the relationship between cyber security risks and other types of risks (e.g., physical). Current data about levels of investment and perceived risk are self-reported at best and frequently are anecdotal. Research is needed into feasible information-gathering techniques, methods to produce meaningful measures, and ways to communicate the results of the measurement process effectively. In particular, research from the public health sector into population risk assessment and into effective risk communication may provide an appropriate model for the information infrastructure protection domain.

[60] See NIST Special Publication 800-36 (draft), *Guide to Selecting Information Technology Security Products;* NIST Special Publication 800-35 (draft), *Guide to Information Technology Security Services*; NIST Special Publication 800-55 (draft), *Security Metrics Guide for Information Technology Systems*, October 2002; and NIST Special Publication 800-26, *Security Self-Assessment Guide for Information Technology Systems*, at http://csrc.nist.gov/publications/drafts.html.

[61] The Information Processing Technology Office (IPTO) of the Defense Advanced Research Projects Agency (DARPA) is beginning research in Information Assurance Measurement (IAM). It is expected that research gaps will remain, particularly regarding metrics and models that relate to business models, and models of infrastructure dependencies.

- Metrics and models to support decision making

  Crucial needs remain for metrics, supported by conceptually sound models, to assess the costs and benefits of cyber security alternatives. Metrics are needed to support decisions ranging from those of managers of information systems (e.g., product selection, configuration restrictions) to enterprise executives (e.g., how much to spend, or whether to transition to a wireless communications infrastructure) to policy-makers (e.g., determine expected consequences of proposed regulations or legal theories). Metrics, and supporting conceptual models, are needed to support several forms of analysis that inform decision-makers:

  - Economic analysis

    Research is needed to define cost-benefit models for cyber security that include life-cycle costs, address the existence of externalities, enable determination of productivity and opportunity benefits as well as of opportunity costs, and represent costs and benefits in ways consistent with other business cost models. Of particular interest are cost and benefit models and metrics that can be used to make comparisons over time; research is needed into how to represent assumptions about technology and the operational environment that support or disallow the use of a model or measurement system.

    Organizational models and metrics are also needed. Current efforts focus on defining levels of maturity for organizational security programs and assessing compliance with those levels or with regulatory requirements. Research is needed into the security implications of different organizational structures and processes (e.g., centralized vs. decentralized management, hierarchical vs. P2P reporting structures, incident response procedures).

  - Risk analysis

    Research is needed into models and metrics for risk and its component factors: threat, vulnerability, consequence, and safeguards. Of particular interest is development of insider threat models and associated measurable indicators of insider threat activity. While the insider threat is of increasing interest in the research community, major needs remain to address aspects specific to the information infrastructure protection context, including the transient and provisionary nature of some insider relationships.[62] These relationships raise issues of privacy and confidentiality of enterprise information in sharing records about insider activity.

[62] See Section 3.1.

Investigation is needed into the potential definition and use of risk-sharing models, where risks are not contained within a single organization but are shared between organizations sharing infrastructure responsibilities. This also includes the possibility that not all organizations will (or are willing to) equally share risks.

  - Technical analysis

    While product assessment techniques and processes exist, the need for metrics that can inform product selection remains. Product performance benchmarks could provide a framework for more useful security metrics. Research is needed to define models of, and metrics for, security properties and dependencies of different types of information infrastructure components.[63]

    A critical need is for techniques to determine the properties of a system from the properties of its components. The research community has made repeated attempts to address this composability problem, but no general solutions have been found, and some problems are specific to the information infrastructure protection context. These include composition of systems from components that can be assessed or modeled to different degrees of fidelity, the need to make large-scale assessments (e.g., enterprise- or sector-wide), and the need to combine data gathered at different times.

- Simulation

  Research is needed to construct simulations (based on sufficiently complete models of infrastructure components and their interdependencies) of the security-related behavior of the information infrastructure. Simulations are needed to explore how vulnerabilities are exploited, how networks are affected by attacks, how attacks can metastasize based on infrastructure interdependencies, the impact of heterogeneous systems and enforcement of different policies in different parts of the network, and how attacks could be coordinated across multiple systems or infrastructure sectors. The simulations must also be capable of representing dynamic composition of the information infrastructure, such as the addition or removal of network nodes. Models of interdependency need to incorporate and support models of risk, so the consequences of an attack—to the enterprise and its ability to perform its mission, as well as to a given system—can be better understood. Finally, the simulations must include time as a dimension to support the examination of evolving scenarios over time (e.g., sequences of actions, such as initial attack, protective response, follow-up attack).

[63] Such models can guide the development of scanning and analysis tools for software, devices, and systems; conversely, the output of such tools can be used to define metrics. See Section 3.3.

3.8 Law, Policy, and Economic Issues

Brief Problem Description

Throughout the course of the I3P's workshops, surveys, and interviews, it became increasingly clear that the framework of economic factors, laws, regulations, and government policy in which the information infrastructure exists and develops is of fundamental importance. In this equation, market factors, laws, regulations, and policies are the independent variables, and technology—what is developed, what is deployed, and how it is used—as well as cyber security practices, are the dependent variables. In a market-based economic system, it is not surprising that the market for IT and cyber security products defines the state of cyber security.

Two closely related questions appear to drive decisions on how security products and services are acquired and used: (1) what are the cyber security risks to the enterprise and how do they fit into the overall risk equation of a company, and (2) what is the value of cyber security—how much financial benefit it provides. There are no clear answers to these questions.

To achieve a cyber security posture that most benefits the nation, we must also achieve a clearer understanding of the threats to our information systems, our vulnerabilities, their consequences, and the parameters of externalities and any associated market failures. We must supply decision makers in the public and private sectors with tools that help them understand cyber security risks, cyber security market forces and structure, the likelihood of success of different approaches to the problem, which levers they control for affecting cyber security, and the consequences of their use. These levers include such diverse interventions as tax policy, regulation, changes to liability legislation, insurance requirements, education and training programs, and standards and best practices.

Existing Research and Capabilities

While many believe that there are real cyber security threats, vulnerabilities, and consequences, there has been no comprehensive analysis of this problem to demonstrate it more largely to policymakers, legislators and other decision-makers, and indeed there are skeptics who question the importance and scope of the problem.[64] This is the one research area in which the gaps in the national R&D portfolio are practically all-inclusive. Other than a few well-defined issues, all areas are in need of significant research.[65] Currently, decisions, legislation, and policy priorities appear to be based on

[64] See, for example, Joshua Green, *The Myth of Cyber Terrorism,* The Washington Monthly Online (http://www.washingtonmonthly.com/features/2001/0211.green.html).

[65] This is not to say that no research is being done, but rather that all areas are, at a minimum, in need of further research. See, for example, the papers presented at the 2002 Workshop on Economics and Information Security (http://www.sims.berkeley.edu/resources/affiliates/workshops/econsecurity), the RBAC economic analysis (http://www.nist.gov/director/prog-ofc/report02-1.pdf), and the 1997 reports to the President's Commission on Critical Infrastructure Protection (http://www.ciao.gov/resource/pccip/pccip_documents.htm).

statistics that only describe part of the problem, anecdotal evidence, and impressions. From both public and private sector perspectives, stakeholders emphasized that the fundamental questions underlying the information infrastructure protection problem—what could and should be done about it, the efficacy of prospective actions, and who should be responsible and why—should motivate this R&D Agenda.

I3P Research Areas

Specific areas for research include:

- Problem definition

  Research is needed to enable a better understanding of the dynamics that shape information infrastructure protection: how do changes in legal, policy, or economic factors, as well as technology, affect the others? Such an understanding is fundamental to the development of cyber security solutions that will positively address economic competitiveness, national security, homeland security, and public health and safety. In particular, research is needed to establish the scope and magnitude of the information infrastructure protection problem and to provide solid analysis in terms meaningful to business and government that address its potential and likely impact on the economy, security, and public health and safety.

  For any emerging technology, companion research is needed into the legal, policy, and economic implications as well as the cyber security implications of the technology and its possible uses. Without this companion research, decisions will continue to be made on the local, regional, and national levels that affect the information infrastructure and cyber security without a full or nuanced understanding of their complexities and risks. This research would address, among other things, issues of liability and indemnification; whether or not to regulate and the impact of various regulatory regimes, including voluntary or mandatory compliance; and the nature and scope of privacy issues.

- Market issues

  As stated above, what technology is developed, whether it is deployed, and how it is used are fundamentally functions of the market for IT and cyber security products and services. Some argue that only changes in market structure, whether through actions that affect private sector forces or government policies or regulation, can bring about real cyber security change in our largely private sector owned and operated information infrastructure.[66] An analytically based understanding of the market components and how various interventions would likely affect the market is of fundamental importance. Research is needed to describe the structure and dynamics of this market. To do this, the usefulness and timing of various interventions—changes in enterprise purchasing patterns, cyber

[66] It should be noted that this R&D Agenda in general and this assertion in particular do not relate to classified government systems.

security laws, regulations, government acquisition practices, policies, auditing practices, insurance, and other factors—on cyber security in general, and on the development, deployment, and use of cyber security technology in particular, need to be understood. For example, several I3P workshops highlighted as particularly important the development of an understanding of the likely effect of changes in liability laws on all aspects of cyber security.[67]

Currently, the "market" value of cyber security to an individual organization is not well understood and may be underestimated. Indeed, the value of investments designed to keep bad things from happening, if successful, is difficult to quantify in the absence of those bad things. Research is needed to describe the value of cyber security in terms meaningful to businesses and the market.

- Tradeoffs

  Several approaches are possible for addressing cyber security. Currently, the federal government's approach relies on public-private partnerships and the influence of persuasion; more rigorous analysis needs to be done on the prospects for success of this approach. Furthermore, some are beginning to refer to much of cyberspace as a "commons." While each piece of the information infrastructure belongs to some company or government, responsibility for the collective information infrastructure is unassigned. Looking out from the boundaries of a given organization, the information infrastructure does indeed resemble a commons—anyone can use much of it to send, receive, or look for information, or perform malicious acts. In particular, the owner of any given portion of the information infrastructure is only responsive to the security needs of his or her own organization, regardless of the needs of or impact on other users of that infrastructure. For example, individual computer users with broadband connections do not typically choose robust security, believing it is not worth the cost or trouble. This lapse creates the potential that their computer may be used as a "zombie" machine in a DDoS attack.[68] ISPs, which could supply security software and require all customers to use it, do not see it as their responsibility to supply or demand firewalls or other security precautions of their customers and believe that if they did so they would lose customers to competitors who do not require this extra expense. Enterprises may or may not purchase sufficient security for their enterprises, and they may be even less likely to purchase security to protect other enterprises, individuals, or the infrastructure itself. Finally, the government has not decided how or whether to require such security and investment throughout the infrastructure.

[67] Specific areas of liability included software liability and the liability of corporations and individuals whose cyber security lapses adversely affected others.

[68] In a DDoS attack, an attacker finds unsecured computers with broadband connections and installs software that permits him or her to use that machine—a zombie in the common usage—to bombard a victim with unwanted information, tying up the target of the attack and preventing the target from using its system for any traffic on the information infrastructure.

Aggressive approaches that more fully use the powers of the federal and state governments are also possible, but the costs and benefits are not well understood and the reasons for a general reluctance to regulate are well known. This statement raises the question of who is responsible for security in this information infrastructure "commons" and who should pay for it.

Another current belief, frequently voiced in I3P meetings with key stakeholders, is that there is a fundamental tradeoff between technological capability, and privacy and civil liberties. These concerns can hinder the development and fielding of technology, as seen in news stories surrounding the Department of Defense's research into information mining and organization systems.[69] Yet, technology can also be used to protect civil liberties if its development and use is governed by thoughtful policies.

To address these concerns, research is needed to develop a fundamental understanding of the limits and likelihood of success of cyber security public-private partnerships, and to develop an analytically based understanding of the implications of other approaches that use governmental intervention such as regulation and changes in liability law, funding policy, and federal acquisition practices. Research is also needed to analyze the issue of cyber security burden sharing between the public and private sectors—which should pay for security, how much security is needed—and between the federal government and state and local governments. Finally, research into understanding and articulating the tradeoffs between technology, and privacy and civil liberties is needed to develop and articulate our understanding.

- Standards and generally accepted security principles and practices

  Standards and generally accepted security principles and practices could be critical tools in addressing cyber security and could apply directly to several key questions highlighted above. Standards can be industry-created or regulatory, while generally accepted security principles and practices are guidelines that may be voluntarily adopted but that can also affect important factors such as liability. Additionally, both of these have potentially important economic and technical implications. On the one hand, they might affect U.S. industry competitiveness in the global information infrastructure market (e.g., by bounding the development of innovative products to only those that fit within a given regulatory scheme), while on the other hand they might make possible security solutions that could not exist without some structured framework. Research is needed to articulate the implications of cyber security standards, analyze the effectiveness of best practices in terms of likely economic impact and liability implications, and give an analytic basis for decisions on issues of ownership (e.g., who should articulate

[69] See, for example, John Markoff's article in the *New York Times,* "Pentagon Plans a Computer System That Would Peek at Personal Data of Americans," November 9, 2002, and the subsequent *Times* Opinion Column by William Safire, "You are a Suspect," on November 14, 2002.

standards or define best practices; who, if anyone, should enforce them; and what the implications for security and economic competitiveness are).

- Modeling, metrics, and data

  For corporate and national leaders to make informed decisions on cyber security, they must have ways to measure the magnitude of the problem and the likely effect of proposed solutions, the data that adequately describes the state of the information infrastructure, and the tools that will help them understand the dynamics of the problem. This is a particularly important component of any solution, and metrics and modeling research topics are described in depth in Section 3.7. However, data collection and maintenance have clear legal and policy implications. There is a need for research to determine what data is needed to accomplish modeling and forensic tasks, the implications of its collection for other issues such as privacy and civil liberties, who should collect it, how long it should be held, and other issues affecting our ability to assess and understand the problem.

- Direct response

  A final area of concern for government leaders and lawmakers is direct response to attacks, which is sometimes referred to as "hack back" or "active defense." The Law, Policy and Economics cluster group highlighted this as an important area needing research, yet it is clearly a controversial one as well. Some workshop participants also held strong opinions on the scope of and appropriate entities for conducting direct response, ranging from those who believe that only properly authorized government bodies should be permitted to take these types of actions to those who believe that more permissive policies are needed. These discussions made clear that physical metaphors—protection of property being the most prevalent—need to be examined carefully before they are used as analogies in understanding the complications and implications of action in this field. For example, accosting an intruder in one's home would almost never have foreign policy implications, yet many malicious acts in cyberspace originate or pass through foreign countries. Policies or laws permitting action against these foreign entities could be of international significance. Furthermore, the technical difficulties of properly identifying an actual attacker, rather than just the last server in a chain that has been used to conduct the attack, cloud the matter.[70]

  Yet arguments that government and private entities should be able to protect themselves by rendering an attacker harmless merit careful consideration. Research needs to be conducted in this field to develop an understanding of the potential value of changes to laws or policies that would permit direct responses in given circumstances and their implications.

[70] See Section 3.5, "Traceback, Identification and Forensics," for a discussion of this issue.

Section 4: Conclusion

This initial I3P Cyber Security R&D Agenda identifies topics of significant value to the security of the information infrastructure that are either not funded or under-funded by the collection of private sector, academic, and government-sponsored research in the United States. The R&D Agenda is based on information gathered and analyzed during the 2002 calendar year, and reflects the input of experts in industry, government, and academia. Areas in which new or additional research is needed include:

- Enterprise Security Management
- Trust Among Distributed Autonomous Parties
- Discovery and Analysis of Security Properties and Vulnerabilities
- Secure System and Network Response and Recovery
- Traceback, Identification, and Forensics
- Wireless Security
- Metrics and Models
- Law, Policy, and Economic Issues

Participants in the R&D Agenda development process frequently noted the importance of needs which did not result in identification of research gaps, but are important considerations for researchers who seek ultimately to affect the practice of cyber security. These include education, training, and awareness; quality assurance methodologies; information sharing and coordination; practicable procedures; and physical security. Similarly, technology transfer was frequently identified as problematic. Potentially valuable research results may not be reflected in products or systems, and research may not be informed by operational considerations. Participants in the R&D Agenda development process highlighted the desirability of further attention by cyber security R&D programs to strategies for improving the flow of ideas and technologies between researchers, product developers and system integrators, and end user organizations. In particular, pilot demonstrations, exercises, and experiments that span systems and organizations could serve to drive better problem definition as well as technology transfer.

The Agenda, together with its supporting information, is intended to aid researchers in identifying problems and R&D program managers in defining program directions. The R&D Agenda will be updated by continuing in-depth research into the state of the art and practice in cyberspace security, by revisiting some of the stakeholder groups who contributed to this effort, and by evaluating the research products made possible by this R&D Agenda. The I3P also recognizes that the information infrastructure is truly global, and that a large portion of the R&D in this field is done overseas. It is the I3P's intent to expand its focus to the international community in 2003, to broaden its understanding of the cyber security problem, and to receive the benefits of the insights of the global research community.

Appendix A: I3P On-Line Resources

To serve as an information clearinghouse on the status of R&D efforts for information infrastructure protection and to foster collaboration among cyber security R&D efforts in academia, industry, and government, the I3P offers a Web portal (http://www.thei3p.org) and a digital archive.

The I3P Web portal serves as an information-sharing mechanism among I3P consortium members and provides an Internet presence to support the education and outreach objectives of the I3P. I3P publications are posted on the Web portal, including the survey and analysis documents as well as concept papers. The I3P Web portal provides discussion forums, enabling the general public to comment on and discuss these publications as well as other topics related to the operations of the I3P. I3P events are announced on the portal, and works in progress (e.g., workshop minutes and briefings) are made available on a need-to-know basis.

The goals of the I3P Digital Archive are to (1) provide a centralized repository of cyber security information, (2) enable faster dissemination/access by a broader community to cutting-edge cyber security research, and (3) facilitate maintenance and revision of the I3P National R&D Agenda. Because there will be different types of users with different requirements for information, some provision may need to be made for restricting access to sensitive information. At a minimum, the Digital Archive will be a repository for electronic copies of publications from the information security community

Appendix B: R&D Agenda Development Process

Unique Characteristics

The I3P process identifies priority research and development areas for information infrastructure protection that are either not funded or under-funded by the collection of private sector and government-sponsored research. This Cyber Security R&D Agenda, and the activities and analyses which support it, is intended to serve the research community and to promote collaboration and information sharing among academia, industry, and government. By identifying and prioritizing high-leverage cyber security research "gap" problems of national importance, this Agenda also serves as the basis for the I3P's program planning and research funding and evaluation processes. The I3P process operates in parallel with government and industry efforts. It nevertheless differs from many such efforts because it enjoys the following combination of features:

- *Resources.* The I3P has committed substantial resources to implement and sustain this process. Teams of experts from the MITRE and RAND corporations and the Institute for Security Technology Studies at Dartmouth College have been engaged to lead all phases of the Agenda development effort.
- *Infrastructure.* The I3P is also establishing a number of permanent mechanisms to facilitate development, validation, and revision of the agenda; these mechanisms include the I3P Digital Archive, the I3P Information Portal, and I3P Cluster Groups.[71]
- *Outreach.* Rather than relying on a relatively small group of technical experts, the I3P seeks input from the full spectrum of stakeholders. The I3P process engages not only technologists, but also representatives of critical infrastructure domains, industry solution providers, government, and specific problem domain experts. As part of the process leading to this first iteration of the Cyber Security R&D Agenda, the I3P hosted or supported more than a dozen events dedicated to obtaining input from outside experts. Workshops, supported by Web-based surveys, elicited input from the following sectors: Banking and Finance, Energy, Chemical, Water, Information Technology and Telecommunications, Emergency Response, Manufacturers and Vendors, Not-for-Profit Researchers, Government, and Transportation. The I3P supported the formation of cluster groups in the areas of Enterprise Security Management, Wireless Network Security, and Legal, Policy, and Economic Issues.
- *Up-to-Date.* The I3P will update its Cyber Security R&D Agenda on an annual or biannual basis.
- *Procedural Rigor*. The I3P has adopted a transparent and procedurally rigorous process that can be reproduced (and improved) during subsequent iterations.

[71] A cluster group is a community of experts and researchers at different institutions working on related problems. See the I3P Portal, http://www.thei3p.org, for the I3P Cluster Group Concept Paper.

- *Research Funding*. Funding permitting, the I3P national research agenda will guide a process that funds research to fill gaps in the nation's cyber security R&D portfolio.

The Agenda Development Process

The I3P used the process illustrated below to develop this Cyber Security R&D Agenda. While the elements of the process are shown in order, the process developed with many steps running in parallel. As this is not a one-time effort, feedback and lessons learned from each iteration of the process will inform subsequent iterations and help make the process stronger and more useful.

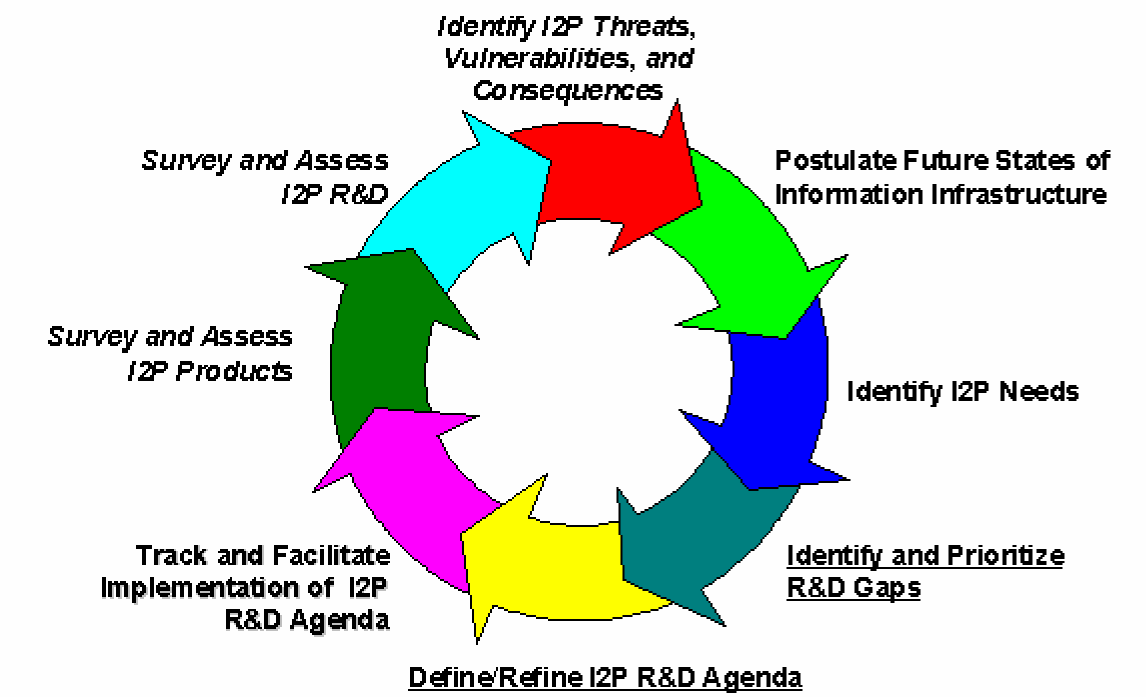
Identify I2P Threats, Vulnerabilities, and Consequences
Postulate Future States of Information Infrastructure
Identify I2P Needs
Identify and Prioritize R&D Gaps
Define/Refine I2P R&D Agenda
Track and Facilitate Implementation of I2P R&D Agenda
Survey and Assess I2P Products
Survey and Assess I2P R&D

Establish Baseline

The I3P sought to establish a baseline for the current and anticipated state of information infrastructure protection. That baseline includes the cyber security marketplace, existing R&D, and stakeholder perceptions of threats, vulnerabilities, and consequences.

The I3P surveyed and assessed cyber security products and research that could be applied to information infrastructure protection. The resulting analyses are living documents, intended to elicit further information and to provoke discussion, as well as to provide input to the agenda process. The I3P Web Portal serves as a venue for discussion and for submission of additional information. The I3P will update the analyses as necessary to reflect inputs from consortium members, researchers, planners, and R&D funding agencies. Stakeholder perceptions were elicited through the workshops and Web-based surveys.

Needs Assessment

The I3P elicited statements of current and anticipated uses of the information infrastructure, and the associated cyber security needs, from stakeholders through the workshops and Web-based surveys. To provide a common starting point for discussion of anticipated needs, the I3P postulated states of the information infrastructure at different points in the future. These portraits of future states, available from the I3P Portal, reflect information technologies, architectures, and processes; how information infrastructures are (or are expected to be) used; and the policy, legal, and regulatory environment.

The I3P also collected and assessed analyses of future needs contained in relevant agenda-setting and strategic-planning documents.[72] The I3P will seek validation of its needs assessment from Consortium members, area experts, and stakeholder representatives on an ongoing basis.

Gap Analysis and Agenda Construction

The I3P identified cyber security technology and decision support needs based on inputs from participants in sector and stakeholder workshops, respondents to I3P surveys, domain experts participating in I3P-sponsored cluster groups, expert opinion from Consortium members and advisors, and such documents as technology roadmaps and research agendas. Workshop and cluster group participants assigned relative priorities to the needs they identified.

For purposes of the gap analysis, high-priority needs were initially grouped by areas identified by workshop or cluster group participants. These areas were primarily functional. Specific needs, which either defined or implied a set of related research problems, were identified and assessed as possible research gaps. The gap analysis considered factors such as the relative priority of different needs, the likelihood of R&D success, and the timeframe in which R&D occurs. Those problems, which were determined to be major gaps, of crucial importance to information infrastructure protection or partial gaps, providing an opportunity for synergy and leadership, and of high importance to information infrastructure protection, were then regrouped into research areas. The definition of research areas is intended to provide a more uniform level of characterization and to facilitate development of research programs.

This initial R&D Agenda incorporates analyses of related agenda-setting and strategic-planning activities. This Agenda addresses identified gaps, encourages promising activities and programs, and identifies areas suitable for new research funding. Should funds become available, the Agenda will guide a grant-making process addressing unmet needs in information infrastructure protection R&D.

[72] These technology surveys and supporting documents are available at http://www.thei3p.org.

Track and Facilitate Implementation

The objectives of this step are to ensure that the overall process leads to useful results and to act as a feedback loop for the next iteration of the process. The I3P will broadly disseminate these findings and track changes in assumptions, risks, vulnerabilities, products, R&D and requirements that might require revision or reassessment of one or more facets of the Agenda and the Agenda development process. If funds permit, the I3P will also implement mechanisms to support research that fills the identified gaps in the nation's information infrastructure protection R&D portfolio.

References and Supporting Documentation

I3P History

The following documents are available at the I3P Web portal, http://www.thei3p.org:

- PCAST Letter on Critical Infrastructure Protection, President's Committee of Advisors on Science and Technology, 10 December 1998, http://www.ostp.gov/html/pcastcip.html
- *A National R&D Institute for Information Infrastructure Protection (I3P)*, Institute for Defense Analyses, April 2000, IDA Paper P-3511
- White Paper on the Institute for Information Infrastructure Protection, Office of Science and Technology Policy, 11 July 2000

I3P Publications

The following documents are available at the I3P Web portal, http://www.thei3p.org:

- *Towards a National Cyber Security Research and Development Agenda: The I3P Process*, Version 1.0, 15 May 2002
- *I3P Cluster Group Concept Paper*, 3 June 2002
- *Analysis of Products, Research, and Roadmaps: Context and Pervasive Issues*, Version 1.0, 9 September 2002
- *Survey of Products, Tools, and Services*, Version 1.0, 9 September 2002
- *Survey of Research and Development*, Version 1.0, 9 September 2002
- *Survey of Related Roadmaps and R&D Agendas*, Version 1.0, 9 September 2002
- *Gap Analysis*, Version 1.0, date to be determined

Summaries of Stakeholder / Sector Workshops and Supporting Surveys:

- Memorandum: I3P Workshop: Toward a National Cyber Security R&D Agenda, Washington DC, 15-16 April 2002
- Memorandum: I3P Workshop: Industry Perspectives on the Process for Developing a National Research Agenda, Washington DC, 3 May 2002.
- Memorandum: I3P Workshop: Banking and Finance Sector, Washington DC, 24-25 June 2002
- Memorandum: I3P Workshop: Energy, Water and Chemical Sector, Washington DC, 15-16 July 2002.
- Memorandum: I3P Workshop: Telecommunications and Information Technology Sector, Washington DC, 29-30 July 2002.
- Memorandum: I3P Workshop: Emergency Responders Sector, Washington DC, 13 September 2002
- Memorandum: I3P Workshop: Manufacturing and Vendor Sector, Santa Monica CA, 24 September 2002

    - Memorandum: I3P Workshop: Not-for-Profit Research Sector, Washington DC, 2 October 2002
    - Memorandum: I3P Workshop: Government Stakeholders, Washington DC, 16 October 200
    - Memorandum: I3P Workshop: Transportation Sector, Washington DC, 23 October 2002

Related Roadmaps, R&D Agendas, and Studies

- Publications of the Computer Science and Telecommunications Board of the National Research Council:
    - *Cybersecurity Today and Tomorrow: Pay Now or Pay Later*, 2002. <http://www.cstb.org>
    - *IDs – Not That Easy: Questions about Nationwide Identity Systems*, Committee on Authentication Technologies and Their Privacy Implications, 2002. < http://books.nap.edu/html/id_questions>
    - *Youth, Pornography, and the Internet*, Committee to Study Tools and Strategies for Protecting Kids from Pornography and Their Applicability to Other Inappropriate Internet Content, 2002. < http://bob.nap.edu/html/youth_internet>
    - *Information Technology Research, Innovation, and E-Government, Committee on Computing and Communications Research to Enable Better Use of Information Technology in Government*, 2002. <http://books.nap.edu/books/0309084016/html/index.html>
    - *IT Roadmap to a Geospacial Future*, Workshop on the Intersection of Geospatial Information and Information Technology, forthcoming. <http://www7.nationalacademies.org/cstb/project_geospatial.html>
    - *Embedded, Everywhere: A Research Agenda for Networked Systems of Embedded Computers*, 2001. <http://www7.nationalacademies.org/cstb/pub_embedded.html>
    - *Trust in Cyberspace*, 1999. <http://www7.nationalacademies.org/cstb/pub_trust.html>
    - Realizing the Potential of C4I: Fundamental Challenges, 1999. <http://www7.nationalacademies.org/cstb/pub_c4i.html>
    - *Improving Surface Transportation Security: A Research and Development Strategy*, Committee on R&D Strategies to Improve Surface Transportation Security, National Materials Advisory Board, 1999. <http://bob.nap.edu/books/0309067766/html>
    - *For the Record: Protecting Electronic Health Information*, 1997. <http://www7.nationalacademies.org/cstb/pub_fortherecord.html>
    - Cryptography's Role in Securing the Information Society, 1996. <http://www7.nationalacademies.org/cstb/pub_crisis.html>
    - *Computers at Risk: Safe Computing in the Information Age*, 1991. <http://www7.nationalacademies.org/cstb/pub_computersatrisk.html>

- Interagency Working Group (IWG) publications:
  - *High Confidence Software and Systems Research Needs*, IWG High Confidence Software and Systems Coordinating Group, January 10, 2001. <http://www.hpcc.gov/pubs/hcss-research.pdf>
  - *Report of Workshop on New Visions for Large-Scale Networks: Research and Applications*, Large Scale Networking (LSN) Coordinating Group of the Interagency Working Group (IWG) for Information Technology Research and Development (IT R&D), March 12-14, 2001. <http://www.hpcc.gov/iwg/pca/lsn/lsn-workshop-12mar01/3.pdf>
  - *Report of the Workshop on New Visions for Software Design & Productivity: Research & Applications*, IWG Software Design and Productivity (SDP) Coordinating Group, December 13 - 14, 2001. <http://www.isis.vanderbilt.edu/sdp/SDP_Wrkshp2-5-6-02.doc>

- Reports to the President's Commission on Critical Infrastructure Protection: <http://www.ciao.gov/resource/pccip/pccip_documents.htm>
  - Legal Foundations series, 1997
    1. Legal Foundations: Studies and Conclusions
    2. The Federal Legal Landscape
    3. The Regulatory Landscape
    4. Legal Authorities Database
    5. Infrastructure Protection Solutions Catalog
    6. Major Federal Legislation
    7. Adequacy of Criminal Law and Procedure (Cyber)
    8. Adequacy of Criminal Law and Procedure (Physical)
    9. Privacy and the Employer-Employee Relationship
    10. Legal Impediments to Information Sharing
    11. Federal Government Model Performance
    12. Approaches to Cyber Intrusion Response
  - *Commercial Perspectives on Information Assurance Research*, 1997.
  - *Economic Impacts of Infrastructure Failures*, 1997.
  - *Government Incentive Tools: A Study*, 1997.
  - *Incentives to Encourage Infrastructure Assurance Investments*, 1997.
  - *Liability and Insurance: Infrastructure Assurance*, 1997.
  - *Opinion Survey of Infrastructure Owners and Users*, 1997.
  - *Regulating the Internet*, 1997.
  - *Threat and Vulnerability Model for Information Security*, 1997.
  - *Toward Deterrence in the Cyber Dimension*, 1997.

- A Security Agenda of the NSF Directorate for Computer and Information Sciences and Engineering (CISE), consisting of the program announcements in 2001-2002.
  - The *Data and Applications Security Program*, announced 23 May 2002. <http://www.nsf.gov/pubs/2002/nsf02132/nsf02132.htm>
  - The *Trusted Computing Program*, announced 6 September 2001. <http://www.nsf.gov/pubs/2001/nsf01160/nsf01160.html>

  - The *Networking Research Program*, announced 2 May 2002. <http://www.nsf.gov/pubs/2002/nsf02123/nsf02123.htm>
  - The *Highly Dependable Computing and Communication Systems Research (HDCCSR) Program*, announced 5 April 2002, jointly funded with NASA. <http://www.nsf.gov/pubs/2002/nsf02114/nsf02114.htm>

- A Security Agenda for the DARPA involving programs announced in 2000-2002.
  - Advanced Technology Office (ATO) program agenda, specifically:
    - The Composable High Assurance Trusted Systems (CHATS) Program. <http://www.darpa.mil/ato/programs/chats.htm>
    - The Cyber Panel Program.
    - The Dynamic Coalitions Program. <http://www.darpa.mil/ipto/research/ftn/index.html>
    - The Fault-Tolerant Networks Program. <http://www.darpa.mil/ipto/research/dc/index.html
  - Information Awareness Office (IAO) program agenda <http://www.eps.gov/spg/ODA/DARPA/CMO/BAA02-08/Attachments.html)>
  - Information Processing Technology Office (IPTO) program agenda, specifically the Organically Assured and Survivable Information Systems (OASIS) program and the OASIS Demonstration and Validation (DEM/VAL) program <http://www.tolerantsystems.org>
  - Other related programs:
    - Advanced Technologies for Information Assurance and Survivability (ATIAS) program, operated by the Air Force Research Laboratory (AFRL), for the former ITO. <http://www.afrlsn.afrl.af.mil/index.html)>
    - The Dynamic Assembly for System Adaptability, Dependability, and Assurance (DASADA) program, operated by IPTO. <http://www.darpa.mil/ipto/research/dasada/index.html>
    - The UltraLog program, operated by the Tactical Technology Office (TTO). <http://www.darpa.mil/tto/programs/ultralog.html>

- *National Scale INFOSEC Research Hard Problems List (draft)*, by INFOSEC Research Council (IRC), September 1999; *The INFOSEC Research Council*, John Davis, Mitretek Systems, November 2001. <http://www.infosec-research.org/>

- *Research and Development Program Development Plan: Information Security for Critical Infrastructure Protection (draft)*, by PCIS WG#4, February 2002 <http://www.pcis-forum.org/>.

-

- *Annual Reports to the President and Congress of the Advisory Panel to Assess Response Capabilities for Terrorism Involving Weapons of Mass Destruction,*[73] 1999, 2000, 2001, 2002. <http://www.rand.org/nsrd/terrpanel/>

[73] These are also known as the Gilmore Commission Reports.

- *Preliminary Research and Development Roadmap for Protecting and Assuring the Information and Communications Infrastructure,* part of the report entitled *Preliminary Research and Development Roadmap for Protecting and Assuring Critical National Infrastructures*, by the Transition Office of the President's Commission on Critical Infrastructure Protection and the Critical Infrastructure Assurance Office, July 1998.
  <http://www.ciao.gov>

- *Research and Development Priorities for Communications and Information Infrastructure Assurance,* by four Department of Energy National Laboratories, June 1997.
  <http://george.lbl.gov/security/PCCIP/C+I_Report.html>

Acronyms

CA	Certificate Authority
CCEVS	Common Criteria Evaluation and Validation Scheme
CERT	Computer Emergency Response Center
CERIAS	Center for Education and Research in Information Assurance and Security
DARPA	Defense Advanced Research Projects Agency
DDoS	Distributed Denial-of-Service
ESM	Enterprise Security Management
FFRDC	Federally Funded Research and Development Center
FTC	Federal Trade Commission
FY	fiscal year
I&A	identification and authentication
IAM	Information Assurance Measurement
I3P	Institute for Information Infrastructure Protection
IDA	Institute for Defense Analyses
IDS	intrusion detection system
IEEE	Institute of Electrical and Electronics Engineers
IETF	Internet Engineering Task Force
IP	Internet Protocol
ISP	Internet Service Provider
IPTO	Information Processing Technology Office
ISTS	Institute for Security Technology Studies
IT	Information Technology
LAN	local area network
NCSC	National Computer Security Center
NIST	National Institute of Standards and Technology
NSC	National Security Council
OMB	Office of Management and Budget
OSTP	Office of Science and Technology Policy
P2P	peer-to-peer
PCAST	President's Committee of Advisors on Science and Technology
PCIPB	President's Critical Infrastructure Protection Board
PKI	public key infrastructure

R&D	research and development
RFC	Request for Comments
SAML	Security Assertion Markup Language
SATAN	Security Administrator Tool for Analyzing Networks
SCADA	Supervisory Control and Data Acquisition
WAN	wide area network
XML	eXtensible Markup Language