

COMMERCIAL PERSPECTIVES ON INFORMATION ASSURANCE RESEARCH

Report to the
President's Commission
on Critical Infrastructure Protection
1997



This report was prepared for the President's Commission on Critical Infrastructure Protection. The report represents the opinions and conclusions solely of its developer, IDA. The Commission has not made a judgment on, nor should the publication of this document be taken to represent an endorsement by the Commission for the content of the material contained herein.

IDA Paper P-3359

**Commercial Perspectives on
Information Assurance Research**

William T. Mayfield
Ron S. Ross
Stephen R. Welke
Bill R. Brykczynski

October 1997

Prepared for

The President's Commission on Critical Infrastructure Protection

INSTITUTE FOR DEFENSE ANALYSES
1801 N. Beauregard Street, Alexandria, Virginia 22311

H 97-002996/1

Preface

This document was prepared by the Institute for Defense Analyses for the President’s Commission on Critical Infrastructure Protection (PCCIP). The work, sponsored by the National Security Agency (NSA), was performed under the task order entitled “Development of DoD’s Information Assurance (IA) Program Management Processes.” The document responds to specific tasking from NSA to assist the PCCIP by:

- Providing an assessment of commercial IA research and development funding, and
- Determining where commercial technologies providers are currently investing and where they think investments should occur in the future.

The results contained in this document also contribute to the overall task objective, which is to provide technical analyses to support the development of IA programs for the Department of Defense. Contributing participants and reviewers of this document are listed separately in the Acknowledgments page.

Acknowledgments

The authors wish to acknowledge all the individuals who have contributed to the development of this report. We include all the technology providers; Mr. John Davis, Commissioner of the President's Commission on Critical Infrastructure Protection (PCCIP); external reviewers; and the staff at the Institute for Defense Analyses (IDA).

Technology Providers

- AT&T - Mr. John Kauza
- CISCO - Mr. Dan Scheinman, Ms. Laura Ipsen
- Gemini Computing - Dr. Tien Tao
- Haystack Labs - Mr. Steve Smaha
- Hewlett-Packard - Mr. Jim Schindler, Mr. Fred Luiz, and Mr. Hal Ableson
- Intel - Dr. David Auksmith
- IBM - Mr. Charles Palmer, Mr. David Yaun, and Mr. Bill Whitehurst
- Lucent Technology - Mr. Mel Cohen, Mr. Bill Coran
- Microsoft - Dr. George Spix
- Motorola - Mr. Bob Firth
- Novell - Dr. Glenn Ricart
- Oracle - Ms. Louanna Notargiacomo
- Raptor - Mr. Lance Urbas
- Secure Computing - Dr. Thomas Haigh, Mr. Spence Minear
- Security Dynamics/RSA - Mr. Art Coviello, Dr. Burt Kalisky
- Spyrus - Dr. Russell Housley
- Sun Microsystems - Mr. Joe Alexander, Mr. John Leahy, Mr. Pierce Crowell, and Mr. Tom Hassing
- Sybase - Dr. Thomas Parenty, Mr. Gene Thurston, and Mr. Somansundaram Shanmugam
- Trusted Information Systems - Mr. Steve Walker, Ms. Marty Branstad
- Wheelgroup - Mr. Lee Sutterfield, Mr. Scott Olsen
- 3COM- Mr. John Hart, Dr. Dan Nasset

External Reviewers

- NSA - Mr. Dick Schaeffer, Dr. Robert Meushaw, Mr. Grant Wagner, Ms. Chris McBride

IDA

- Computer and Software Engineering Division: Dr. Richard Ivanetich, Dr. Al Brenner, Ms. Katydean Price, Ms. Joyce Walker, Ms. Leslie Norris, Ms. Helen Robertson, Dr. Dennis Fife, Dr. Ed Feustel
- Cost Analysis and Research Division: Dr. Tom Frazier, Dr. John Bailey
- System Evaluation Division: Dr. Robert Turner

Table of Contents

Executive Summary	ES-1
Section 1. Introduction	1
Section 2. Status of Current IA Research	7
Section 3. Major Drivers of Commercial IA Research Investment	11
Section 4. Commercial Information Assurance R&D Investment Trends	17
Section 5. IA Areas Needing More Funding and/or Emphasis	23
Section 6. Government Roles in IA Research	33
Section 7. Conclusion & Observations	37
Appendix A. Key Technologies	A-1
A.1 Basic Research in IA Fundamentals	A-3
A.2 System-Level Engineering	A-14
A.3 Individual Component Development	A-23
Appendix B. IA Research Investment Estimate	B-1
Reference List	References-1
Acronym List	Acronyms-1

Figures

Figure 1. Interview Template	5
------------------------------------	---

Tables

Table 1. Technology Providers Interviewed	3
Table 2. Annual Overall R&D Expenditures	18
Table 3. Framework for IA Research	24
Table A-1. Top IA Research Needs	A-2
Table B-1. Revenues, Overall R&D Expenditures, and Est. IA Research Funding	B-7

EXECUTIVE SUMMARY

OVERVIEW

This report responds to the following tasking from the President's Commission on Critical Infrastructure Protection:

- Provide an assessment of commercial information assurance (IA) research and development (R&D) funding, and
- Determine where commercial technologies providers are currently investing and where they think investments should occur in the future.

The results documented in this report will assist the Commission in developing its own set of recommendations for a national IA research agenda, including recommendations for government funding, as part of the Commission's final report to the President.

A small research team from the Institute for Defense Analyses (IDA) conducted a set of interviews with twenty-one telecommunications and computer technology providers. These providers included large companies with significant information technology markets and niche security technology companies that provide specialized solutions.

The time allotted to gather data and write this report was constrained, thus precluding the participation of several other such companies that were initially identified by the Commission as willing to participate. The interview technique was thought to be the best approach at getting the data requested by the Commission within the time allotted. While consolidating the set of responses, which varied widely in content, the IDA team also provided a snapshot assessment that included the providers' viewpoints and their identification of those IA research areas that have gaps or are critically needed. Findings, conclusions, and appendices containing supporting data and context are provided.

FINDINGS

Finding 1 The U.S. commercial information assurance R&D activity is fairly robust in breadth, but is lacking in depth.

In general, the commercial IA research and development activity covers the breadth of telecommunications and computing technologies. It is expected that such research will continue to broaden as these technologies evolve. What appears to be lacking is depth in several critical areas of research (e.g., security management, operating systems, and database management systems).

Finding 2 Industry believes that it “owns” the commercial IA technology problem and should spend to solve it.

The responsibility for incorporating information assurance into commercial products and infrastructure systems belongs to both those industries providing the products and those industries using the information technologies in their infrastructures. Critical infrastructure providers can now begin to manage their risks through a continuous improvement process that includes the appropriate buying of available IA-enabled products. Technology providers cannot wait for customer demand—they must help “shape” that demand with methods that increase awareness of possibilities and sharpen awareness of need.

Finding 3 U.S. commercial information assurance R&D investment is focused on satisfying customer demand, especially electronic commerce.

Commercial technology providers are specifically focusing their IA research on electronic commerce. However, their customers’ widely varying perceptions of IA technology needs and benefits have significantly (often adversely) impacted demand. Legacy systems of their customers provide major barriers to new technology integration. Customers want to achieve financial savings from employing new security technology, not added costs. Commercial providers can change these perceptions through better listening to customer views and being prepared to take advantage of the constantly changing consumer market.

Finding 4 All the companies interviewed indicated that their R&D investments in IA technology were increasing and that, for most companies, this trend should continue for the next few years.

There are many factors contributing to this finding, including (1) the increased demand ensuing from networking via the Internet, (2) increased market competition, and (3) electronic commerce as the current driver. The industry is reacting to increased demand via increased overall R&D as well as increased advanced technology partnerships, alliances, consortia, and other associations. Technology start-ups are also increasing. These companies are facing global opportunities and competition. Their ability to compete in the global market will play a significant role in any increase in investments.

Finding 5 There are important areas of IA research that either are not being pursued by commercial technology providers or else require additional emphasis and funding.

In order to organize and convey the relative importance of what commercial industry believes is needed, the authors developed an IA research framework consisting of three categories: (1) basic research in IA fundamentals, (2) system-level security engineering, and (3) individual component development. The most significant gap in pursued IA research is system-level security engineering, particularly in the area of system-level security architectures. System-level security engineering must be further supported by basic research in IA fundamentals, particularly in the areas of availability and integrity. Among the most critical needs in individual component development are security management and intrusion detection. More work in assurance technology is also needed across all three categories of IA research.

Finding 6 Technology transfer remains a significant problem.

The lack of strong ties between industry, academia, and government makes it difficult for technology to be transferred from research to implementation and use. Partnerships and alliances are occurring with greater frequency to facilitate technology transfer and will predominate in the near term. Applications programming interfaces and standards will also be key factors in promoting technology transfer.

Finding 7 Export control policy is perceived to be the biggest barrier to further commercial IA investment, thereby reducing the capability to protect our critical infrastructures.

Most technology providers had a strong opinion about what constituted the greatest obstacle to further investment in information assurance: U.S. government policy on export controls of strong cryptography. Most providers believe that the “equity” issues between protection and access (e.g., key length, escrow, or key recovery) have been overcome by the fact that cryptography has become globally pervasive—it is perceived to be “out of the box!” Significant pent-up venture capital is believed to be available for information assurance R&D should this barrier be removed. Commercial providers believe that the global markets are key to their competitive survival. The unintended consequence of current U.S. government policy is that it is primarily working against protecting our own systems because it restricts large U.S. telecommunications and computing technology providers who must compete in a global market.

Finding 8 Government-funded research, leadership, and vision can make a difference.

All the respondents believe that government-funded research is important in achieving IA goals over the long term (e.g., a goal of continuous improvement). Further, they believe that government-funded university research and curricula are important tools in producing much-needed, technically qualified human resources, which are currently very scarce.

**CONCLUSION &
OBSERVATIONS**

The authors arrive at one basic conclusion from the findings and provide three observations for the Commission’s consideration based on the findings and this conclusion.

Conclusion *Commercial industry believes that it must solve the IA problem for critical infrastructures. However, industry will not take the necessary actions—to the degree required—without government leadership, facilitation, and motivation.*

Commercial industry will meet many of the IA needs of our critical infrastructures as it meets the needs of its customers, especially in the area of electronic commerce. However, meeting those needs is not going to be sufficient. Government

leadership will be required to overcome industry inertia, to define the problems to be addressed, and to motivate solutions.

Observation 1 Any new government-sponsored information assurance R&D should be a long-term, continuous investment activity that should not be expected to play a significant role in protecting the nation's critical infrastructure in the near term (i.e., within three years).

Given today's commercial product cycles, it is unlikely that any new government-sponsored research will produce protection results that can be transferred within three years into our critical infrastructures. Current IA-enabled products must now be used within the critical infrastructures, with a plan to continuously improve the protection utility of these products. It is important that new government research efforts take a long-term view of continuous improvement and that the government direct its research funding accordingly.

Government should start focusing through university research on IA fundamentals and emerging telecommunications and computing technologies. These fundamentals and technologies are needed for improved system-level security engineering.

The government must also begin to look at the need for future global solutions and examine its export control policies accordingly.

Observation 2 The model for future information assurance R&D should be multi-disciplined and collaborative, involving industry, academia, and government in well-focused, mutually supportive partnership efforts.

Most providers agreed that government should play a significant role in information assurance R&D and in protecting our critical infrastructures. Such a role they believe should be to facilitate and motivate infrastructure protection, and to provide funding to sponsor private-sector information assurance R&D, especially university R&D.

The providers believe that increased use of partnerships between government, universities, and industry will be needed in light of scarce expertise. They believe that pre-competitive technology development and large-scale experimentation should be considered.

The Chief Information Officer Council and/or the INFOSEC Research Council are two possibilities for organizing and

coordinating the intra-governmental efforts. Some relief from regulatory relief may be required as well as developing creative approaches to partnerships.

Observation 3 The use of a national information assurance research agenda (NIARA) is an appropriate vehicle for communicating the vision of what is needed to address information assurance as a component of critical infrastructure protection.

Government vision backed with resources was seen as a critical component in motivating industry to action. A long-term government initiative will be needed to kick-start critical infrastructure protection and provide sufficient momentum to ensure that the effort can become self-sustaining (i.e., protection must be thought of as a continuous process).

The NIARA could be executed in a decentralized fashion but with oversight coordination through the government's recently established Chief Information Office Council or through an expanded Federal INFOSEC Research Council. Since the framework developed by the authors for Finding 5 (page ES-3) covers the spectrum of research needs for information assurance related to infrastructure protection, it could be used as a starting point for creating the NIARA. However, to ensure that individual infrastructure sectors are receiving appropriate attention, the three categories of the framework should have two focal points: one that is technology aligned and a second that is sector aligned.

**FINAL
REMARKS**

Of the three categories identified in the authors' IA research framework, system-level security engineering is the most pressing overall need. However, information assurance R&D investment is impacted by more than a technical research agenda and levels of investment funding. Government policy, commercial competition, customer demand, and the lack of qualified human resources all have a significant impact on information assurance R&D investment. If the Commission is to have a significant and long-lasting effect on our nation's critical infrastructure protection through information assurance R&D, it must take into consideration all of these factors in its deliberations for recommended actions.

Section 1.

Introduction

PURPOSE This report provides a commercial perspective on current and future information assurance (IA) research needs to support critical U.S. infrastructures.¹ It includes a snapshot assessment reflecting the nature of security-relevant telecommunications and computing research in the commercial sector. The results reported herein will assist the President's Commission on Critical Infrastructure Protection, which is charged with developing a national strategy on critical infrastructure protection. Recommendations for evolving a national IA research agenda and for further government-funded IA research will be partially derived from this report and incorporated into the Commission's final submission to the President of the United States.

PROBLEM Information is indispensable to all aspects of our nation's critical infrastructure operations. Without the protection that ensures the availability, integrity, and, in some cases, the confidentiality of that information, the information-based processes underlying infrastructure operations will fail. Without the ability to accurately exchange information among components within an infrastructure, such operations will also fail.

Today, the information environment supporting such critical infrastructure operations is dramatically changing. More and more, the nation's critical infrastructures are turning from analog controls to digital automation and networking to gain competitive efficiencies. Due to their growing information dependence, these infrastructures now increasingly require protection from cyber and physical attacks. What is significantly different today is that these infrastructures and their potential cyber threats share much of the same networks. Cyber-oriented protection is the focus of IA research (i.e., IA research must

¹ Eight critical infrastructures have been identified by the Commission: telecommunications, electrical power systems, gas & oil storage and transportation, banking and finance, transportation, water supply systems, emergency services, and continuity of government services.

address the problem of assuring the integrity and availability of such automated system and networking resources, and of the information that is manipulated, transmitted, or stored within them).

ORGANIZATION In this report, we document the status of commercial IA research, the drivers and trends of commercial IA research investment, the IA research areas identified by commercial information technology providers as needing more emphasis and/or funding, and the opinions of commercial providers regarding government policy and needed government participation in IA research. From the findings, we have drawn one basic conclusion and provided three additional observations to help focus the Commission in preparing its final report.

We provide two appendices to the report which contain additional details on (1) the areas that commercial providers identified as needing more emphasis and/or funding, and (2) the level of commercial IA research investment.

TASKING The Commission tasked the National Security Agency's Chief of Information Security (INFOSEC) and Technology Research to provide an assessment of research related to information assurance in both the government and private sectors. The NSA, in turn, tasked the Institute for Defense Analyses (IDA) to assist in this assessment. NSA gave IDA direct liaison authority to work with the Commission on the private sector portion of the assessment.

APPROACH The Commission and IDA jointly determined that an interview approach, as opposed to sending out and following up on a survey, would be the better tool in terms of time and being able to directly interact with the technology providers.

A two-member research team from the Computer and Software Engineering Division of IDA interviewed twenty-one commercial information technology providers. This team consisted of Mr. Terry Mayfield, Assistant Director, and Dr. Ron Ross, Research Staff Member. Table 1 lists the twenty-one companies, organized by type (large company and niche company).

Table 1. Technology Providers Interviewed^a

Large Companies	Niche Companies
IBM * †	Secure Computing Corp.
Hewlett-Packard * †	Security Dynamics
Sun Microsystems †	Raptor
Novell	Haystack Computing
3COM †	WheelGroup
CISCO * †	Trusted Information Systems
Lucent Technologies * †	Gemini Computing
AT&T * †	Spyrus
Intel * †	
Motorola * †	
Oracle * †	
Sybase	
Microsoft * †	

a. *Business Week*, July 7, 1997, pp. 52-97. An asterisk (*) signifies “Top 100 Global Competitors.” A dagger (†) signifies “Global 1000 (U.S. Country Composite).”

The interviewed technology providers consisted of large companies with significant market shares of software and hardware technologies for telecommunications and computing, and niche information security companies that provide specialized solutions to securing the infrastructure. Company Presidents, Chief Technology Officers, Chief Operating Officers, Product Development Managers, and Security Architects served as interviewees.

Eleven of these companies, ranked by market *value*, are in the top two hundred U.S. companies reported in the “Business Week Global 1000.” Nine of these companies are ranked in the “Top 100 Global Companies.”²

² *Business Week*, July 7, 1997, pp. 52-97.

Constraints The limited time frame in which the interviews took place (May through June 1997) imposed a constraint on the number of technology providers that could participate in the process. The Commission had identified twenty-seven technology providers as possible candidates for interviews, but for various reasons only twenty-one agreed to participate within the time frame available.

The Commission recognized from the beginning that getting quantifiable data would be difficult at best. This possibility proved to be true with respect to gaining commercial IA research investment data. Such quantified investment information was not made available to us by the industry participants. The interview data did not allow us to precisely or generally quantify the amount of commercial IA research being done, either in terms of numbers of people involved or in terms of dollars expended or budgeted. This absence of quantified data led to an estimate being independently calculated by the authors. Our estimate comes primarily from calculations performed on other forms of industry research data that were publicly available.

Interview Template In conjunction with the Commission, IDA created an interview template containing eight questions and used it to guide the (often) free-form discussions and to pull specific quantifiable information from the persons being interviewed. Figure 1 contains the interview template.

INTERPRETING THE REPORT

We have used the respondents' words to the maximum extent possible, with some interpretation and editing to bring similar responses into a more cohesive flow. We also have some added organizational structuring and provided some additional context to make the report coherent.

Our findings are derived primarily from the set of twenty-one interviews. Additional, publicly available source material was incorporated to expand the context of the finding wherever appropriate. Collectively, the findings provide a representative assessment of current private sector investment in IA research.

1. What are the major thrusts of INFOSEC and Assurance research investments? [Elaborate as much as possible; quantify and qualify categories to the maximum extent possible (include national and international). Identify area(s) of principal concentration.]
 - Hardware Platforms (e.g., Switches to PCs)
 - Software
 - Protocols
 - Architectures (platform, system)
 - Management and Control (DBMS, Networking, Certificate Authorities, Security Management Tools, Software Maintenance)
 - Analysis and Monitoring (Intrusion Detection Systems, Performance)
2. What is the investment trend relationship of information security and quality assurance research to overall research?
3. What do you project these trends will be in three years? in seven years?
4. What areas do you perceive as **NOT** being researched in the protection of information?
5. What critical technologies do you depend upon to be successful? [e.g., trusted OS's, security management tools]
6. What is your view of the adequacy of research across the spectrum of technologies to enhance assurance?
7. Where would you recommend additional resources be brought to bear, and what are the top two areas you want resources applied?
8. Do you have any additional comments you want to provide to the Commission regarding INFOSEC and Assurance related research issues?

Figure 1. Interview Template

Subfindings are highlighted within each finding for contextual emphasis. We also included subfindings that emphasize a concern or issue identified by the technology providers.

The responses from these technology providers often were wide ranging, which the reader will detect and may perceive as conflicting statements. This wide variation—some collectively showing strong ambivalence, others completely unique (including not being able to provide an answer or willing to offer an opinion), and still others collectively having much more unanimity—has been captured herein so that all of their viewpoints are represented. To be faithful to the respondents, the report conveys such mixed messages.

The selected sample set of technology providers was too small and too varied within the set to create data of any statistical significance. Additionally, the sampled industry members and the authors' interpretations of their responses undoubtedly result in some bias in the reported results. Further, while very cooperative in participating, seldom did the participants answer any particular question directly. All were resistant to providing quantified data, although some offered anecdotal quantifications.

Section 2.

Status of Current IA Research

This section of the report addresses the general status of IA research within the commercial sector (Finding 1) and provides an industry perspective on responsibility for acquiring needed technology to protect the nation’s critical infrastructure (Finding 2).

FINDING 1 **THE U.S. COMMERCIAL INFORMATION ASSURANCE R&D ACTIVITY IS FAIRLY ROBUST IN BREADTH, BUT IS LACKING IN DEPTH.**

Subfinding 1.1 *Research generally covers the breadth of the telecommunications and computing technology spectrum.*

Subfinding 1.2 *Research will continue to broaden.*

Subfinding 1.3 *Research activity is lacking in depth.*

Breadth
in research

Subfinding 1.1 discussion. There is commercial R&D activity going on in just about every facet of information assurance in the telecommunications and computing technology spectrum. The providers interviewed described their own research as well as pointing to research being done by others. Developmental research to integrate many of these various technologies, especially cryptography, into existing product lines was emphasized. The information assurance R&D identified as ongoing included:

- Base hardware (i.e., microprocessor), including incorporation of cryptography and other protection-enabling mechanisms (e.g., identification and authentication, reliable time).
- Operating systems protection, including the incorporation of cryptography, identification and authentication, domain type enforcement, and kernel enforcement mechanisms.
- Network protocols, including the implementation of security in the newest version of the internetworking protocol (IPv6), cryptography over asynchronous transfer mode (ATM), and Quality of Service (QoS).

- Directory services protection.
- Security management.
- Identification and authentication technology, including smart cards and biometrics.
- Firewall and guard security enhancements.
- Applications, including workflow-based authorization, database security, portable application code, secure payment mechanisms, and web-browser technologies.
- Intrusion detection systems.
- Applications programming interfaces (APIs), including cryptographic, smartcard, intrusion detection.
- Intellectual property rights protection, including encryption for Digital Versatile Disks (DVD).
- Cryptography, including key recovery and integration into various system layers.

Research
broadening

Subfinding 1.2. discussion. Most of the providers believed that such IA research would continue to grow as computing and telecommunications technologies evolve. New opportunities would arise both from new ideas about security and from new technologies that need security.

Research
missing
depth

Subfinding 1.3. discussion. What is missing is the needed depth of research in any given area. The current depth may only be five or six companies deep in some cases, or it may be an alliance or consortium that is working on a single alternative. While this gives rise to much needed consensus-engineered or *de facto* standards, it doesn't promote the plethora of ideas that could arise from increased depth of research in any given area.

FINDING 2 **INDUSTRY BELIEVES THAT IT “OWNS” THE COMMERCIAL IA TECHNOLOGY PROBLEM AND SHOULD SPEND TO SOLVE IT.**

Subfinding 2.1 *The responsibility for incorporating information assurance into commercial products and infrastructure systems belongs to both those industries providing and those using the information technologies.*

Subfinding 2.2 *Critical infrastructure providers can begin now to manage their risks through a continuous improvement process that includes appropriate buying of available products.*

Subfinding 2.3 *Customer demand will require industry “shaping.”*

Industry
responsibility

Subfinding 2.1 discussion. Needed research in information assurance is an industry problem and industry should spend to solve it. This finding is probably most true for continuing near-term needs (two to three years). The requirement to incorporate IA features into commercial computing and telecommunications products must be accomplished by the providers of those products. The requirement to assemble systems composed of those products belongs largely to the private sector companies providing much of the nation’s critical infrastructure (e.g., electric power systems, telecommunications). However, if the perception of cyber-based risks remains low within the companies providing critical infrastructure, then the demand for such features will not be generated and the technology providers will be much less likely to provide them.

Employ
available
products

Subfinding 2.2 discussion. The key to private-sector infrastructure owners increasing the information assurance within their critical infrastructures is to start employing the products that are available. The strategy that should be employed is not to seek a complete “risk avoidance” solution but rather to put a process in place that aims for managing risks through continuous improvements in information assurance.

A technology that is currently receiving the most attention through such a process is firewall technology. It is a critical start towards building network enclave protection. Another technology that is beginning to emerge is intrusion detection. This technology provides

important operational monitoring capabilities that are needed in nearly every system. These technologies will require operational deployment to gain feedback necessary to improve feature quality and advance them to more capable products.

The quality of IA features in such products can vary considerably and integrating such products into a large-scale system can be daunting, especially when the products have not been thoroughly analyzed individually or within composed systems, and when the composed system-level definition of assurance is not available. Nonetheless, without deployment, these technologies will not be advanced. This is an industry problem that must be solved by industry.

“Shaping”
customer demand

Subfinding 2.3 discussion. Part of the technology-providing industry’s problem is “shaping” customer demand (i.e., increasing awareness of possibilities and sharpening the awareness of need). Networking has become the key element in IA demand “shaping”—understanding the possibilities of networking has given rise to increased awareness of needs regarding information assurance. Networking can change organizational possibilities, but “safe” networking will require IA technologies. The rising awareness of the need for information assurance is beginning to have an effect on the technology providers. Demand for a wider variety of IA technologies is beginning to emerge within the private sector. Correspondingly, a wide variety of IA products is emerging in the marketplace with the depth of providers among such varieties slowly increasing. These products provide sets of IA features needed to protect critical infrastructures.

Section 3.

Major Drivers of Commercial IA Research Investment

This section of the report identifies the major drivers of the commercial industry's research investment in information assurance (Finding 3). It identifies the effects customers have on IA technology providers.

FINDING 3 **U.S. COMMERCIAL INFORMATION ASSURANCE R&D INVESTMENT IS FOCUSED ON SATISFYING CUSTOMER DEMAND.**

- Subfinding 3.1* *Specific emphasis is on electronic commerce.*
- Subfinding 3.2* *Customer demand varies between perceived need and benefit.*
- Subfinding 3.3* *Legacy systems impose a barrier to IA investment.*
- Subfinding 3.4* *Security must give the customer financial savings.*
- Subfinding 3.5* *It is essential to obtain customer views directly.*
- Subfinding 3.6* *The consumer market is constantly changing.*
-

Electronic
commerce

Subfinding 3.1 discussion. The primary driver for information assurance R&D investment in the private sector is electronic commerce. One of the faster growing information assurance R&D areas has been electronic payment technology. The full potential for global electronic commerce can only be enabled if information assurance is provided. This potential is driving the global competitiveness of commercial technology providers. Today's marketplace is truly becoming global and it must have its information protection provided accordingly. A protection solution that includes the global use of standardized, strong cryptography is needed. This protection cannot be individual government mandates as that would not allow reciprocal trust among nations. In general, this will require increasing openness in the development of IA mechanisms.

It can be anticipated that dramatic changes will occur in the nature of commerce and the roles of governments in commerce as electronic commerce takes hold in the ensuing decades. Global electronic

commerce will require governments working together to foster such commerce—and this implies working closer together on policies that enable the protection of electronic commerce information. Strong cryptography is perceived as the foremost electronic-commerce protection technology requirement that must be addressed through inter-government cooperation.

Variation in
Customer
Perceptions

Subfinding 3.2 discussion. Major customers across the infrastructure sectors being addressed by the Commission have significantly different views on their needs to secure their systems and information. Economic reasoning predominates.

The Department of Defense (DoD) has been using its Trusted Product Evaluation paradigm³ over the last two decades to foster widespread availability of trusted products. Although some progress has been made in advancing trust technology, DoD has not really been interested in high assurance products, even when it sometimes states such an interest in its system acquisition requirements statements. Evidence of this assertion abounds—high assurance products are not being purchased and, in many cases, the government is purchasing upgraded versions of original evaluated equipment that have not been re-evaluated. The timing of high-assurance evaluations vs. acquisition and the rapid obsolescence of evaluated products are often given as reasons for not buying such evaluated products.

The financial services sector is perhaps the leader in applying IA technologies. It has long been involved in information security and internal controls, so it doesn't have a great distance to go in determining its current and future IA needs. What is helping this particular sector is a changing viewpoint that security technology can be viewed as an enabler that can be costed against the lost business opportunity costs (i.e., what business cannot be pursued without a

³ *Department of Defense Trusted Computer System Evaluation Criteria* (TCSEC), DoD 5200.28-STD, 1985 (also known as the “Orange Book”). The trusted product paradigm is built on four fundamental blocks: (1) security policy (access control), (2) accountability (identification & authentication; audit), (3) assurance (operational and life cycle), and (4) documentation. Such trust has been oriented largely toward the operating system and its Trusted Computing Base (TCB).

secure system). It is clear to most segments of financial services that there is a significant business potential enabled by security.

The electric power utilities sector, on the other hand, defines “security” as stability in the system. It is most concerned with ensuring that power generation can be brought on and off line safely, that energy “wheeling”⁴ within the grid is performed safely, and that operational faults (e.g., loss of a generator, transmission line, substation or transformer) are contained to preclude cascading power failures across the system. The utilities sector has long been involved with large-scale system monitoring, fault tolerance, and the concept of availability, but not heavily involved in information system protection. This sector has a further distance to go than the financial services sector in recognizing and dealing with its IA needs.

Legacy System
Barriers

Subfinding 3.3 discussion. The capital investment in home-grown legacy systems which are optimized to a particular sector’s applications may also impact the willingness of infrastructure owners to pursue IA technology, especially in a highly competitive market. What is critical in today’s market (for most infrastructures) is getting more performance out of their installed base. This is especially true in telephony and power. Information assurance may lie below the economic investment line for many of the companies within such infrastructures. Replacement and upgrade costs may play a significant factor. Commercial off-the-shelf (COTS) products may not integrate within or easily replace home-grown system components. Longer-term capital replacement to accomplish information assurance through COTS or home-grown upgrades is likely to be the strategy in such situations. Enclave protection, wrapper technology,⁵ and operational monitoring may encompass the most appropriate solution in the short term.

⁴ “Wheeling” is the rerouting of power transmission from one regional grid to another.

⁵ “Wrapper technology” is software that can encapsulate some part of a legacy system—without modification to the legacy part—and provide new IA properties.

Savings Through
Security

Subfinding 3.4 discussion. For some in the marketplace there is a changing perspective on the costs of security. Savings can result from applying security. Take, for example, a large energy-infrastructure company who, through the help of a niche technology provider, has been able to achieve significant savings. These savings came about by eliminating dedicated telecommunications lines and applying firewalls and strong end-to-end encryption which allowed the company to connect to the Internet (i.e., virtual private networking). The dollar trade-off was significant with relatively little change in risk. This is cost-effective “enabling through security.”

The IA marketplace, though rapidly evolving, is still not very well formed. From a technology provider’s perspective, what investment is right? What security standards should be implemented? There is currently a very long list. What makes most sense in this rapidly changing environment? Some companies must be positioned to meet the large variety of these standards (or at least the most prevalent). Which standard among competing ones does a company back with its limited resources? In such cases, it is difficult to break out and lead the pack.

Products are at the center of standards, and standards lead to the types of successes desired by all companies. They provide needed stability to larger companies and, if they are *de facto* standards, they provide strength to those companies that enabled such standards. Today, all companies require greater cooperation in deploying technology. There are more alliances, partnerships, and licensing agreements being established each month to deploy security technology with the hope that their technology will set the standard.

Obtaining
Customer
Views

Subfinding 3.5 discussion. Many respondents emphasized meeting with customers to obtain customer views on what they thought security was and what they needed. One meeting indicated the need for authentication of system parts. There was also a need to address identical vs. unique parts. Another meeting emphasized the need to protect rich intellectual content while recognizing that near-term protection probably cannot stop “Copy Houses” abroad.

One respondent reported that its customers do not want extra authentication. The respondent indicated that the issue of Certification Authorities is most difficult. X.509v3⁶ certificates are available, but they are “rootless”⁷ or exercised with multiple roots. That is why Secure Multipurpose Internet Mail Extensions (S/MIME) schemes have not yet become viable in commercial systems. There is a lack of “trust” chains. Instead there are multiple Certificate Authorities (e.g., VeriSign, Entrust, RSA). Today, both corporate and personal trust are essentially formed by knowledge of an individual or organization from an existing or previous set of circumstances.

There is a strong desire *not* to have the government participate as a central Certificate Authority. Rationale for this desire includes reasons of personal privacy and the issue of not wanting a national identifier. The Pretty Good Privacy (PGP) decentralized approach of creating trust chains from the bottom up (i.e., by one individual at a time) is perceived by some as a good starting approach to creating trust chains. However, it is insufficient. Trust hierarchies will be needed and here the government could have a role.

Changing
Consumer
Market

Subfinding 3.6 discussion. The consumer market itself is changing because of networking and such changes will require enabling security technology. In particular, privacy and infrastructure reliability will be more urgently required. In electronic commerce, one change is the notion of a virtual mall. Such a mall will change the way people interact in purchasing a variety of items. For example, the ability to try on clothes electronically to see how a change in color or design looks will require the submission of more personal information (including digital images of the individual). Such information will require protection at the end points and in transit. As another example, many trust issues emerge in health care as information is increasingly shared

⁶ X.509 Directory Authentication Service is part of the CCITT X.500 recommendations series defining a directory service. X.509 provides a framework for using digital signatures and certificate-based public key authentication. Certificates are created by some trusted Certificate Authority. X.509v3 is the latest version.

⁷ Lacking a common parent key for hierarchically formed keys.

via networks. Telemedicine will require dramatic changes in infrastructure reliability as it moves into the operating theater.

Section 4.

Commercial Information Assurance R&D Investment Trends

This section of the report discusses investment trends of commercial information assurance R&D. The respondents would not quantify their current levels of IA research investment because such information was either proprietary or not specifically collected. All respondents, however, indicated that their investments would be increasing and anecdotally illustrated the reasons why this increase should occur (Finding 4). This trend is discussed in more detail in Appendix B.

FINDING 4 **ALL THE COMPANIES INTERVIEWED INDICATED THAT THEIR R&D INVESTMENTS IN IA TECHNOLOGY WERE INCREASING AND THAT, FOR MOST COMPANIES, THIS TREND SHOULD CONTINUE FOR THE NEXT FEW YEARS.**

- Subfinding 4.1* *Customer demand is increasing because of the Internet.*
- Subfinding 4.2* *Global market competition is increasing.*
- Subfinding 4.3* *Overall R&D funding within technology providers has been increasing.*
- Subfinding 4.4* *Advanced technology partnerships are increasing.*
- Subfinding 4.5* *Technology start-ups are increasing.*
- Subfinding 4.6* *Alliances, consortia, and associations are increasing.*
-

Customer demand *Subfinding 4.1 discussion.* The telecommunications and computing markets have grown significantly due to the Internet. Companies have increased their R&D funding largely in response to consumer demand resulting from the explosive growth of the Internet, the advent of electronic commerce, and the growing awareness of customers regarding their need for protection in an interconnected environment.

Evidence of increased IA awareness is emerging as customers generate more specific IA demands on the technology providers. More customers are realizing that security is an enabler for their businesses and their privacy in an interconnected environment. As awareness continues to increase, so will the demand for technology that can

provide or support information assurance. Thus, while not having specific quantification, the investment trend forecast for the next few years is for continued growth.

Global market
competition

Subfinding 4.2 discussion. Part of the increase in R&D funding can be attributed to global market competition and the need for product differentiation. Providers need to continuously innovate to keep their products fresh and competitive. An increase in advanced development to devise and integrate new security features into existing product lines is envisioned for this reason alone.

All the technology providers indicated their need to grow more capability in information assurance. Their ability to do so will vary. One large technology provider lost its critical mass of security research capability in a company split. This provider is now trying to rebuild in a very difficult hiring market where a limited supply of security expertise is in high demand.

Overall
R&D
increase

Subfinding 4.3 discussion. As indicated in Table 2, the overall annual R&D expenditures within the sampled set of major technology providers have been increasing for the last three years.

Table 2. Annual Overall R&D Expenditures^a

Company	1996 (\$M)	% Increase/ (Decrease)	1995 (\$M)	% Increase/ (Decrease)	1994 (\$M)
Oracle	389	49	261	32	197
Sybase	165	8	152	33	114
IBM	4,654	12	4,170	(4)	4,363
HP	2,700	18	2,300	14	2,000
Sun	657	17	563	13	500
Microsoft	1,432	67	860	41	610
Novell	276	(25)	368	6	347
3COM	233	40	166	65	101
Cisco	399	89	211	98	107
Lucent	1,838	10	1,672	(30)	2,385
Intel	1,808	40	1,296	17	1,111
Motorola	3,152	15	2,743	11	2,461
TOTAL	17,427	18	14,762	8	13,696

a. Source: Company 10K's.

For most of these providers, 1996's R&D growth was a significant increase over previous years; aggregate R&D funding growth in 1996 over 1995 was 18%, with one company's R&D growing by 89%. It is anticipated by the respondents that 1997 will continue to reflect large increases in R&D (e.g., Microsoft has announced intentions to spend \$2.1 billion for research in 1997, a 48% increase over 1996's funding). Competition is driving technology providers to develop new products and to add more features to their existing products in response to customer demand and product differentiation. Competition is also driving producers to innovate in ways to reduce costs. These drivers should continue to fuel R&D for the foreseeable future.

Such R&D increases also can be seen in an analytical report of U.S. industry R&D spending patterns produced by the U.S. Department of Commerce's Office of Technology Policy (OTP).⁷ The industry basis of the report was U.S. publicly traded, R&D-conducting firms. As described in an article in *New Technology Week*,⁸ the OTP report indicated two key developments.

[C]ompanies that are 'more R&D intensive' now account for a larger portion of the U.S. economy, having replaced 'more traditional businesses'; and U.S. companies overall are spending a larger percentage of sales on R&D than they did at the beginning of the last decade. A 'crossover' occurred in about 1980, when industry funding for R&D overtook government funding. The total share of U.S. private-sector R&D carried out by companies in the electronics and information fields jumped from 32 percent in 1981 to 41.9 percent in 1988, then moved up to 43.6 percent in 1995. Electronics [is a sector] where federal R&D spending has been significant. 'Sustained Government investment over a long period of time...has created a base and placed the U.S. in a

⁷ Office of Technology Policy, *Globalizing Industrial Research and Development*, 1997, PB96-119201NB.

⁸ Ken Jacobson, "Industry R&D Spending Patterns Shifting," *New Technology Week*, May 12, 1997, pp. 6-7. He quotes the Assistant Secretary of Commerce for Technology, Graham Mitchell, regarding the results contained in the OTP report.

strong competitive position ... it has, in addition, attracted R&D spending as the private sector has seen opportunities for commercialization.'

This same OTP report shows that overall industry R&D intensity, as measured by percentage of R&D to sales over the period 1977 to 1995, increased from 1.8% to 3.77%. Wide variation within industry sectors is recognized. The OTP report provides alternative scenarios for 1995 industry R&D spending that show a range of between \$57 billion (with 1981 R&D intensity and portfolio) and \$107 billion (current 1995 OTP estimate). Using the current OTP estimate, the U.S. 1995 R&D spending in electronics (43.6%) would have been about \$47 billion.

The overall increase in R&D funding leads by implication to a conjecture that information assurance R&D funding will also increase, though not necessarily in proportion to the overall increase. In general, most providers felt that their R&D funding for information assurance was adequate and would, in most cases, be increasing. However, they were resistant to quantifying such increases by any specific amount or percentage.

Partnerships

Subfinding 4.4 discussion. Some of the R&D funding increase is coming about through advanced development technology partnerships and acquisitions. These relationships enable one partner's IA technology to be integrated into the other's product line. Identification and authentication technology has been a major player in this approach. Encryption technology is another major player in such efforts. Intrusion detection technology is one that is emerging along with firewalls. Security-enabling smart cards will soon be another one to emerge with network computers, personal computers, and workstations.

New technology partnerships and acquisitions are an ongoing feature of today's research and advanced technology acquisition environment and are expected to increase in the next few years. Opportunities for such partnerships and acquisitions exist on an international scale. One provider thought we would see a flurry of IA technology emerging over the next three years and then stabilizing out into the future. Others see the rapid emergence of IA technology in the near term with a steady but

modest growth rate into the future. Partnerships and acquisitions will most likely increase during the rapid emergence of new IA technologies from niche providers, particularly to gain market share and to establish barriers to entry by other providers.

Technology
start-ups

Subfinding 4.5 discussion. A portion of this R&D funding increase will be seen through start-ups. Technology start-ups in the area of information assurance are expanding in number. This expansion has been especially true in the areas of firewalls and intrusion detection. The R&D funding increase is likely to continue to expand as new technology ideas allow innovators to obtain funding for new start-ups.

Such start-ups are increasingly coming from university research operations and may include both faculty and students. The old theme of “publish or perish” is increasingly being changed to “technology start-up or perish.” Downsizing by the corporate “giants” has also contributed to the growth of start-ups as “downsized” technologists with entrepreneurial aspirations begin their own companies. These new niche technology start-ups are the “feedstock” partnerships or are available for acquisition by major telecommunications and computing technology providers.

For many niche technology companies, a significant portion of their start-up costs is in advanced development to mature their product to “industrial strength” and to create “barriers to entry” for other companies who may want to provide similar technology. A current strategy with some start-ups is to literally give away an initial product version of technology to create a large market share and a *de facto* standard, and then sell or license the more robust and advanced product versions of that technology.

An opposite, and perhaps more strongly held, viewpoint is that openness created by government R&D investments in pre-competitive technology can counter such entry barriers and provide for significant start-up opportunities. The DARPA firewall toolkit is an example where government R&D investment (contracted commercial research) provided a base pre-competitive technology that rapidly spawned an industry of firewall providers to fulfill a market need. Increasing

government funding of contract and university research, leading to open technologies that fulfill a specific market need, could have dramatic effects in further spawning and increasing the IA technology industry.

Alliances,
consortia,
associations

Subfinding 4.6 discussion. Technology development associations, alliances, and consortia are another way in which the increasing trend in research funding is being realized. Among the many examples are the Financial Services Technology Consortium, The Smart Card Industry Association, and The Open Group, where the sharing of research risks, funding, and resultant IA technology is occurring. These efforts may be devoted to pre-competitive technologies, technology transfers, and technology standardization.

New associations, alliances, and consortia are continuing to be established. Most of the technology providers interviewed participate in or, in some cases, advise such associations, alliances, or consortia. Many providers participate in several of these organizations. Both standardization and competition interplay as crucial motivators to participation. How those two factors interplay determines the speed of progress and the effectiveness of results

Section 5.

IA Areas Needing More Funding and/or Emphasis

This section of the report presents industry's views on the IA research areas that need more funding and/or emphasis (Finding 5), and a process area (technology transfer) that needs to be emphasized to advance information assurance technology (Finding 6).

FINDING 5	THERE ARE IMPORTANT AREAS OF IA RESEARCH THAT EITHER ARE NOT BEING PURSUED BY COMMERCIAL TECHNOLOGY PROVIDERS OR REQUIRE ADDITIONAL EMPHASIS AND FUNDING.
<i>Subfinding 5.1</i>	<i>The most significant gap in pursued IA research is system-level security engineering.</i>
<i>Subfinding 5.2</i>	<i>The most critical system-level security engineering area is system-level security architectures.</i>
<i>Subfinding 5.3</i>	<i>The two most critical IA fundamentals are the protection concepts of availability and integrity.</i>
<i>Subfinding 5.4</i>	<i>The two areas in greatest need of individual component development are security management and intrusion detection.</i>
<i>Subfinding 5.5</i>	<i>More work in assurance (confidence-building) techniques is needed across all three categories of IA research.</i>

The information gathered from the interviews covered many IA areas that need to be addressed by researchers. In order to organize and convey relative importance of what commercial industry believes is needed, the authors constructed a framework consisting of three categories: (1) basic research in IA fundamentals, (2) system-level security engineering, and (3) individual component development. Table 3 summarizes the framework, with the topics in each category organized according to the relative importance articulated by the respondents. The important items needing to be addressed with some sense of urgency are further articulated in this finding. Additional

details on the topics captured in the framework are contained in Appendix A.

Table 3. Framework for IA Research

Basic Research in IA Fundamentals	System-Level Security Engineering	Individual Component Development
Protection Concepts & Principles	System Architectures	Security Management
System Complexity Issues	Heterogeneous Component Integration	Intrusion Detection
Vulnerability Analysis	Secure Interoperability & Evolvability	Identification & Authentication
Trust Concepts	Applied Engineering Research	Smart Cards
—	System Assurance	Networking
—	Standards	Applications
—	—	Secure Operating Systems
—	—	Applied Cryptography
—	—	Hardware-Based Security

Security
engineering
gap

Subfinding 5.1 discussion. The consensus view among the respondents is that the principal gap (i.e., research area not being pursued) and most critical problem in IA research is at the system level. System-level security engineering must be addressed now, but it will require critical support from research on fundamentals. Among the many issues affecting system-level security engineering are the following (unordered list):

- **Scale:** How is information assurance in large-scale, interconnected systems best understood? How are their descriptions captured? What is the best way to deal with the issues of complexity? How is trust scaled? What principles and theories are needed? This aspect of information assurance is relatively unexplored. One respondent suggested that the National Research Council's *Realizing the Information Future*, along with its *Computers at Risk*, be the starting points for IA researchers to pursue such an understanding.⁸

- **Integration:** How can security components be effectively and efficiently inserted and composed? What is the best way to accommodate legacy components? How is secure interoperability achieved with heterogeneous components? How is assurance of the whole achieved? Understanding policy composition will be a critical element in any integration solution.
- **Design trade-offs:** What are the trade-offs in the design space? Where and how should security mechanisms be allocated? Measurement is the key to enabling trade-offs—more must be done in the area of measuring. There are demands for both sharing and separation of data. These demands provide the key “tensioning” in the IA design space of large-scale information systems.⁹ A range of protection policies must be developed to address design spaces (i.e., domains) stretching across both demands. The economics of security technology and its associated management (e.g., human resources) must also be understood as this technology is allocated throughout the system design space.
- **Effectiveness:** There is an economic demand for efficiency of mechanism and an assurance demand for robustness of mechanism. How are these demands best accommodated? What does it mean to provide cost-effective defense-in-depth in large-scale systems? (That is, there could be many different and potentially independent approaches to security with each approach contributing layers to the overall depth of defense.) What are the costs to ownership and system performance of such defenses?
- **Adaptability:** Large-scale systems must preclude an adversary from attacking via a common weak mechanism. Adaptability will

⁸ National Research Council, *Realizing the Information Future: The Internet and Beyond*, National Academy Press, 1994; *Computers at Risk: Safe Computing in the Information Age*, National Academy Press, 1991.

⁹ By “tensioning,” we mean the conflicting demands placed upon a designer resulting from applying opposing design goals to a single design space. Choosing one goal over another to some degree (e.g., min-max) provides one domain within the design space. Compromise of both goals will be required to achieve a domain which maximizes both goals.

be required at various layers and component interfaces to meet this need. How should this be most effectively accomplished? How is adaptability managed efficiently and securely?

- **Evolvability:** Fielded systems must be economically and technically evolvable over time. What are the principles of large-scale secure system design that enable such evolution?
- **Multi-disciplined requirements:** System-level security must be recognized as a multi-disciplined field. It touches all aspects of a system, including human-computer interfaces, applications, system software, hardware, and networks. Collaborative, multi-disciplined approaches will be needed.
- **Experimentation:** Experimental research involving real protocols, real systems, testbeds, and experimental simulation facilities will be a critical part of what is needed to address information assurance in the large. System testbeds and simulation facilities are needed to allow various concepts and products to be tried out in a large-scale system context. Such experimentation will enable the evolution of a system-level security engineering discipline. The Next Generation Internet (NGI) is a potential government-funded research program where this experimentation could occur. More significantly, using the NGI could enable experimental results that build in security from the inside out and result in a new, secure Internet that could serve as the core for the protection of the common elements of our critical infrastructure.

Security
architectures

Subfunding 5.2 discussion. System-level security architectures that enable secure interoperability among heterogeneous components are critically needed. While many security architectures exist, none quite meet all the requirements for architectures that have been articulated. These architectures must be network-service oriented, information-centric, and founded upon the concepts of secure interoperability *and* secure evolvability. They must address legacy components, scale, flexibility, adaptability, etc. Such architectures should encompass privacy, security, and intellectual property rights. Traffic controls, walls, fences, and other means of disciplining access in large systems

are needed. Several sets of minimal essential infrastructures (MEIs) within the context of system-level architectures will need to be addressed. Each of the sectors being looked at by the Commission would likely have its own unique MEI in addition to a common MEI. More research is needed to address the concepts of layering and security services. Much more work on system integration principles is required. An understanding of where and how to place trust in critical elements is needed. Trusted Computing Bases (TCBs) may not be needed everywhere. Collaboration will be needed to achieve this (set of) architecture(s).

There is also an increasing trend toward adaptivity (i.e., dynamic re-configuration). This trend is evident in virtual organizations, computing in the network, and active networks. IA researchers should examine each sector's protection needs and system-level architectures through this dynamic viewpoint. The ease of dynamic change is promoting organizational use of networking, and networking is promoting dynamic change in organizations. This idea of dynamic change will give rise to entire new areas of protection research: it provides both new opportunities and new challenges.

The issue of knowledge of change (i.e., what is changed) vs. control of change (i.e., where, when, how, and by whom it is changed) needs to be addressed. The ever increasing amount of "hidden" control that is built into a system to make it more user oriented must be understood. With such hidden layers of control, it is now much easier to subvert such systems at multiple levels in the structure. A better understanding of dynamic feature interaction will be required. The analytic tools to promote this understanding must be developed. Systems with such layering, especially as "middleware"¹⁰ evolves, will require new vulnerability analysis and penetration analysis tools as well as modifications to intrusion detection sensors, analysis engines, and

¹⁰ "Middleware" is a term used herein to collectively identify general-purpose service software with common software interfaces and protocols sitting in a layer between a computing platform (i.e., hardware plus the operating system) and its user-specific applications. Examples include object request brokers, distributed computing services, message translators, directories, and database management systems.

response mechanisms. These modified mechanisms must, in turn, be incorporated into the system-level architectures.

Availability &
integrity

Subfinding 5.3 discussion. Availability and integrity were identified as the two biggest technology problems urgently requiring both research and product development. The issue of availability is one area that needs many things to be addressed (e.g., more work on Class of Service and Quality of Service protocol constraints, efficient resource-reservation protocols, and fault-response policies). An understanding of what it means to write applications that work well with various impedance mismatches and unreliable components is needed. The application of fairness and market-driven resource allocation algorithms should be pursued further. This suggestion includes an operating system model for resource reservations to support video and audio channel mix (i.e., a better resource manager). Similar resource managers will be required for mixed-media communications. More work on robust architectures is needed to include improved component and system restart and recovery capabilities.

Security
management &
intrusion detection

Subfinding 5.4 discussion. Security management is the most costly burden of implementing security in today's systems, a burden that is increasing as systems become more interconnected. Infrastructure, tools, and procedures are needed to reduce this cost of ownership. Infrastructure is required to support "trust authorities" (e.g., certificate authorities). Secure end-system and network administration tools, particularly those tools that promote secure remote administration, must be developed to reduce the number of and workload on system administrators. Finally, standardized procedures and supporting mechanisms must be developed to provide secure associations between communicating parties to pass appropriate security information (e.g., access control permissions, policy enforcement requirements, certificates).

A great deal of R&D has been performed in the area of intrusion detection. This work has produced useful tools, but the tools are point-product solutions that have trade-offs between what they can detect and how many false-positives they generate. Most providers believe more

work is needed in large-scale network security monitoring and intrusion detection. Better integration of security administration and system monitoring tools will be required. Some believe that a consortium of companies working on intrusion detection systems (similar to what has been done for cryptographic public key management) is needed.

Assurance
techniques

Subfinding 5.5 discussion. Assurance is a topic that received ambivalent treatment from these technology providers, though all expressed some requirement for more work. On the one hand, they wanted more assurance, and, on the other, they wanted to reduce the burden of high assurances since no one is demanding such assurances through their buying activities. Many have incorporated some form of quality assurance into their product development process and believe this approach to be sufficient for developmental assurance.

Many felt that the assurance requirements imposed by DoD were no longer relevant. One hardware vendor pointed out that formal design validation is becoming increasingly important as complexity is overtaking testing and there is insufficient time to test once full production begins.

One thought is that the fundamentals and wisdom contained in more than thirty years of work in assurance should be picked out and re-applied. Several believe that policy, accountability, and documentation (especially assurance evidence) remain core elements that should be required. Others believe that operational assurance (i.e., monitoring) is most important and that it will be required to ensure immediate detection and response, including confinement of damage. Most agreed that we, as an IA community, need more work in assurance that includes:

- Defining assurance with more precision, including more work in risk assessment.
- Developing and applying scientific analyses, including the ability to express requirements for trustworthiness and approaches for evaluating such trustworthiness.

- Incorporating robustness through fault avoidance (reliability) and fault tolerance (a significant aspect of availability that includes reconfiguration).
- Adding structure in our systems—which comes from policy. How should assured domain (enclave) interconnection be achieved?
- Providing support for integration and operational certification.

FINDING 6

TECHNOLOGY TRANSFER REMAINS A SIGNIFICANT PROBLEM.

- Subfinding 6.1 Strong ties between industry, academia, and government are missing.*
- Subfinding 6.2 Partnerships and alliances will predominate technology transfer.*
- Subfinding 6.3 Application Programming Interfaces (APIs) and standards can promote technology transfer.*

Missing strong ties

Subfinding 6.1 discussion. One perspective on much of the academic research is that it is out of touch. More practical operational experience is needed to make information assurance relevant to the technology product needs of today. In this perspective, academic-driven research lags behind the niche security companies in the commercial world. Another view is that there is a big disconnect between the commercial companies and security research. What the technology providers perceive as missing are the strong ties between major industry players, academia, and government that existed in the 1970s and early 1980s. Such ties, while mainly based on government funding, were also highly collaborative.

Partnerships & alliances

Subfinding 6.2 discussion. The use of partnerships and alliances is facilitating the incorporation of niche IA technologies into mainstream computing and telecommunications products. This approach to technology transfer is occurring with increasing frequency. Niche companies usually can move faster to develop a particular technology. In general, they are “hardening” technology that came from academia or contract research. Niche companies are also helping to create the market for such technologies in larger companies providing mainstream computing and telecommunications products. Partnerships

with value-added resellers provide additional sales channels and increased integration with other products.

APIs &
standards

Subfinding 6.3 discussion. Standards are critical, and the lack of them is viewed as a significant technology problem. The ability to incorporate new technologies into existing platforms and software is enhanced when APIs and standards are developed. These standards can be consensus engineered or *de facto*. One approach to technology transfer and to standardization is the use of open “reference” technology. Research that culminates in demonstrable reference technology prototypes may aid in technology transfer. Whatever is done in applied research should include a transition path to incorporate the resulting technology—and, in general, that path must evolve from existing technology and systems.

Section 6.

Government Roles in IA Research

This section of the report addresses the commercial viewpoints regarding government’s role in IA research. It discusses the commercial providers’ perspective of the biggest barrier—export control—to furthering their IA research investment (Finding 7), and suggests where government should strengthen its roles in IA research (Finding 8).

FINDING 7 **EXPORT CONTROL POLICY IS PERCEIVED TO BE THE BIGGEST BARRIER TO FURTHER COMMERCIAL IA INVESTMENT, THEREBY LIMITING THE CAPABILITY TO PROTECT OUR CRITICAL INFRASTRUCTURES.**

Subfinding 7.1 *Pent-up venture capital could be released to support global market if the government policy is changed.*

Subfinding 7.2 *Creation of a sustainable market is important—global markets are key.*

Venture capital/global markets *Subfinding 7.1 discussion.* It is a strongly held opinion by most of the interviewed providers that the “equity” debate surrounding current U.S. export control policy is hurting their global competitiveness. Cryptography is “out of the box”—it is pervasive! One source indicated that the world-wide market in software technology has moved from about \$100 billion in 1994 to \$250 billion in 1997, and that U.S. export controls need to become better aligned with this market if security is to be included in mainstream products.

It is also the opinion of several respondents that this area could “explode” and become replete with funding¹² if the export control policy is changed. One key perspective—but not yet a consensus one—is that pent-up venture capital could potentially be available for information assurance R&D should the federal government’s policy on export controls related to this technology (i.e., cryptography) be lifted.

¹² One large company, currently hesitant to invest much effort in security, believes that there will be at least a \$2 billion per year industry potential once export controls are lifted.

The likelihood of such policies being changed is currently in question. A major focus on this policy controversy is now within the U.S. Congress. For example, the Senate Commerce Committee recently passed a bill to back the current Administration's position on export controls. There is separate legislation on encryption policy proceeding in the House of Representatives—strongly endorsed by the Business Software Alliance—which has a contrary position to the Administration's. It is not clear what the outcome will be between the competing positions.

Additional perspectives on the issues and options of government policy on cryptography export controls can be found in *Cryptography's Role in Securing the Information Society*.¹³

Sustainable
markets

Subfinding 7.2 discussion. As with other technologies, the creation and growth of a sustainable market (i.e., buyers and suppliers) for IA technologies will cause continued (and potentially increasing) R&D investments in these technologies. Global markets are a key to growth in these industries and the U.S. technology providers are facing increasingly competitive foreign providers (e.g., a German cryptography technology company's flouting of U.S. export controls is cited as the reason for its success). It is also becoming clear that some U.S. technology providers are moving some R&D capabilities off shore to get around current government export controls (e.g., Sun Microsystems's move to use Russian cryptography, RSA's subsidiary in Japan). IA technology is clearly emerging in a global market that could become quite lucrative to those technology providers who get there first with the needed set of technologies.

There is also evidence that foreign governments are willing to put a substantial R&D investment in place where they cannot acquire the needed IA technologies from U.S. providers (e.g., the Singapore government's \$80 million cryptography research initiative¹⁴ that started because it could not import strong U.S. encryption technology).

¹³ National Research Council, *Cryptography's Role in Securing the Information Society*, National Academy Press, 1996.

¹⁴ This funding is assumed to be over a five-year period.

The result of such foreign actions and the continued policies of the U.S. government could mean a reduction in the potential global market share of U.S. technology providers and, therefore, a reduction in their investments in IA technology.

FINDING 8 **GOVERNMENT-FUNDED RESEARCH, LEADERSHIP, AND VISION CAN MAKE A DIFFERENCE.**

Subfinding 8.1. Government-funded research is important in achieving IA goals over the long term.

Subfinding 8.2 Government-funded university IA research and curricula are important in producing technically qualified human resources.

Long-term
research

Subfinding 8.1 discussion. One perspective is that government can enhance the form and levels of research in this area by leadership, vision, motivational incentives, and the commensurate funding necessary to achieve common IA goals. Several respondents believe that government cannot predict where research dollars will have commercial success. Thus, government should not try and select particular telecommunications and computing research for primary focus with the goal of driving security into existing commercial products. Rather, the government should spread a portion of its research funding across the base of computing and telecommunications technologies, focusing on strategic, pre-competitive, and fundamental research with long-term IA goals (as technology challenge problems) to be met.

More long-term, high-risk, pre-competitive research with government backing is needed to achieve order-of-magnitude changes.¹⁵ This area is where government funding over a sustained period of time can make a significant difference. Opportunities for new ideas that push the

¹⁵ In general, “long term” refers to a continuously sustained effort. However, within such a continuous effort there will exist many specific, often overlapping, time-based projects that were each envisioned as lasting between three and five years. This is not unlike other engineering disciplines that have sustained research efforts with some level of continued government funding sponsorship.

technology envelope should be explored. More scenario-based motivation for IA research should be provided.

The government may be able to define some programmatic leverage points that, if proven successful, could make orders-of-magnitude difference in the results of IA research (e.g., more protection-enabling capabilities built into hardware, robust and efficient assurance-enabling protocols, protection-enabling personal access devices, self-protecting information objects, and the technology base to provide international trust infrastructure(s)). Increased collaboration is needed, which the government can facilitate. Encouraging innovative partnerships with academia and industry to get involved in large-scale experimentation will provide a vital strength to this effort.

University research
and curricula

Subfinding 8.2 discussion. Human resources that have the necessary technical qualifications to expand the area of information assurance are a scarce commodity. The government should fund university IA research and curricula that produce smart people for industry to hire.

Government-funded university research was strongly encouraged by the sampled technology providers. They felt that such research not only leads to new ideas, but it produces the critical human capital needed within industry to enable it to evolve and grow. Several respondents perceived the need for a predictable commitment of funding to academia (five years plus). They also thought it desirable that such funding be provided with greater flexibility in its use and with a reduced administrative burden.

At least one respondent remarked that the United States does a poor job of educating our university students in security. Information assurance must become more widespread throughout university curricula—the primary focus needs to be on technology (e.g., computer science, electrical engineering), but IA also needs to be part of the curricula for future users (e.g., business, other engineering and science disciplines).

Section 7.

Conclusion & Observations

This section provides a single conclusion and three observations to further assist in the Commission deliberations regarding government information assurance R&D investments. The material contained herein is based on the accumulated interview information, our derived findings, and personal knowledge. We conclude our work with some brief final remarks.

CONCLUSION

COMMERCIAL INDUSTRY BELIEVES THAT IT MUST SOLVE THE IA PROBLEM FOR CRITICAL INFRASTRUCTURES. HOWEVER, INDUSTRY WILL NOT TAKE THE NECESSARY ACTIONS—TO THE DEGREE REQUIRED—WITHOUT GOVERNMENT LEADERSHIP, FACILITATION, AND MOTIVATION.

Most of the technology providers are being driven by electronic commerce. This technology driver, while supportive of many of the needs for critical infrastructure protection, does not address many other protection needs that are specific to critical infrastructures (e.g., application-oriented internal controls, availability, and minimal essential infrastructures (MEIs)).

Several of the commercial infrastructure sectors are becoming increasingly deregulated and more competitive. This aspect of their environment will make individual companies within a particular sector highly motivated to remain competitive and less motivated to worry about the overall IA health of their particular sector. The companies will be motivated to take localized actions to improve their protection, but they will not seek to define and establish either the protected MEI necessary for their individual sectors or a protected MEI that is common to all.

Government leadership will be required to define such MEIs and facilitate their establishment. Three issues that contribute to this conclusion are discussed in the following paragraphs:

- Legacy systems that must be overhauled to provide adequate information assurance;
- Multiple, large-scale trust hierarchies for Certificate Authorities; and
- System evolution.

Legacy Systems Infrastructure legacy systems that must be replaced or significantly modified to achieve improved information assurance will not be adequately protected against many threats in the near term. Protection mechanisms for both external and internal threats are required. For example, insertable software protection is being researched through DARPA's work on wrapper technology. This technology could, at some time in the future, provide an economic alternative to full replacement and could serve to aid in the modification of legacy systems. However, it is a technology that today only holds promise—and it remains unproven.

The current best available insertable technology is the firewall. Firewalls will need to be improved not only in their capacity to provide enclave isolation protection but in their efficiency and operational assurance. The advent of widespread IP security (IPv6) and the use of end-to-end encryption will positively add to the network protection dynamics, but these still will not solve the total protection requirements of the nation's critical infrastructures. Such infrastructures will require such things as:

- Isolation mechanisms in critical servers,
- Improved security association protocols that incorporate availability and trust hierarchies, and
- Application-oriented protection mechanisms (especially intuitive, automated internal controls).

Multiple, Large-scale Trust Hierarchies Multiple, large-scale trust hierarchies¹⁴ will be required to support the MEIs. They must be capable of interoperating to facilitate the common MEI. *Such interoperation is a major technology and policy issue.* These trust hierarchies will form part of a border guard structure for

information assurance of the common MEI (i.e., to operate within the common MEI will require interoperable standards, policies, and procedures among Certificate Authorities).

Application programming interfaces for security association will be needed to support these trust hierarchies. The belief that such hierarchies must be built from the bottom-up means that the government must help facilitate the building of standards and policies that will accommodate interoperable, hierarchical structures and their associated key management. Further, not every company will believe in the common MEI; some will believe that they can go it alone. This independence will cause added difficulty in achieving large-scale interoperable trust hierarchies to support individual sector and common MEIs.

Large-scale trust hierarchies continue to be a research problem as issues of management and assurance at various scales and operational efficiency remain ill defined. The trust hierarchy problem must be solved to achieve the necessary protection of critical infrastructures on a large scale.

System Evolution

System evolution is an issue that relates to having an evolvable scheme of improving infrastructure systems vs. complete system replacement. Infrastructures must be provided with an evolvable technology base. How easy and affordable is it to evolve securely? How will these companies evolve their operating system base to increase the level of security in critical servers? These companies cannot afford to replace equipment or software on a large scale within a short time frame. They must have a means to upgrade in a secure evolutionary form that is based on a technology strategy for continuous improvement of information assurance. Information assurance currently does not have such a continuous improvement path in most of the technologies. There is currently no architectural scheme that is oriented towards secure

¹⁴ Trust hierarchies are layers of infrastructure for multiple forms and numbers of cryptographic key-validation entities that can vouch for the association of a key with organization or individual to guard against impersonation. There are two basic forms: one that is organization-centric (i.e., certificate authorities) and the other is user-centric (i.e., a web of trust).

system interoperability *and* evolvability—this issue remains as a research topic. However, to improve the current situation while awaiting additional research, the government could look at providing IA technology investment incentives (e.g., tax credits) for critical infrastructure sector companies that take specific steps to continuously improve their IA posture, and for technology providers that have demonstrated a continuous path of products containing upgraded IA capabilities.

OBSERVATION 1 **ANY NEW GOVERNMENT-SPONSORED INFORMATION ASSURANCE R&D SHOULD BE A LONG-TERM, CONTINUOUS INVESTMENT ACTIVITY THAT SHOULD NOT BE EXPECTED TO PLAY A SIGNIFICANT ROLE IN PROTECTING THE NATION’S CRITICAL INFRASTRUCTURE IN THE NEAR TERM (I.E., WITHIN THREE YEARS.)**

While information assurance R&D has been ongoing for more than thirty years, it is only within the last four to five years that products possessing IA capability have entered the marketplace with proven commercial viability. Existing products can begin to play a necessary role in the information assurance of our critical infrastructures. The commercial infrastructure owners must invest in and use these existing technology products to their best advantage. In using these products, infrastructure owners must continue to assess the sufficiency of improvements that can be achieved with this existing technology base and to provide appropriate feedback to the product vendors.

*Fundamentals &
Emerging
Technologies*

However, this activity is only a start. Information assurance should not be thought of as achieving an ultimate solution; rather, it should be viewed as an evolving process. This process includes risk management, continuous improvements, and sound investment. A part of this process is providing for the future through R&D. Long-term investments in such information assurance R&D should be focused on fundamentals and on emerging technologies rather than on existing technologies. Such a focus will move information assurance from an implementation craft to an engineering science. It will also promote a better path for technology

transition as information assurance will be addressed much earlier in any given technology's evolution.

The majority of any new research investment by government should be in fundamentals of information assurance, especially those related to large-scale systems. The government should take a leading role in initiating a long-term research investment in IA fundamentals. This sort of research is necessary to develop and maintain a *systems security engineering discipline*. Information assurance must be taught as a fundamental element across various information technology and application domains. Long-term fundamental research should be multi-disciplined so that various technology developers and applications developers have the necessary system-level appreciation and knowledge to build in appropriate capabilities or “hooks” for information assurance. The government should expect that the long-term fundamental research necessary to build and maintain a systems security engineering discipline will be continuous and that early results of newly initiated research could emerge within the mid-term (i.e., four to seven years).

Industry is now primarily focused on applied research and advanced product development, with fairly robust and expanding efforts across a spectrum of technologies and products. However, current industry's applied research efforts are still not sufficient to meet the overall protection needs of our critical infrastructure today. More fundamental research will be needed first (e.g., system-level architectures, availability, and integrity).

Further, it is unlikely, given today's commercial product cycles, that any IA technology research not already underway could affect the protection of critical infrastructures in the near term. It is more likely that results from any new research initiatives would begin to emerge in the mid-term. The majority of any new development investment should be in applied research to develop technology transition paths that incorporate increasing IA capabilities in base technologies. The government should play a limited supporting role in such developmental

investments (e.g., pre-competitive technology development; large-scale, advanced-technology demonstrations).

*Export Control
Changes*

The government can also play a role by changing export control policies. While it is understood as an equity issue, the Administration’s current position on export control of strong cryptography could remain a barrier to significant increases in commercial funding of information assurance R&D (i.e., asserted pent-up venture capital would not be released).

OBSERVATION 2

THE MODEL FOR FUTURE INFORMATION ASSURANCE R&D SHOULD BE MULTI-DISCIPLINED AND COLLABORATIVE, INVOLVING INDUSTRY, ACADEMIA, AND GOVERNMENT IN WELL-FOCUSED, MUTUALLY SUPPORTIVE PARTNERSHIP EFFORTS.

*Government’s
Role*

It is clear that the sampled technology providers believe that government has a significant role in information assurance R&D and in protecting our critical infrastructures. However, they view it as constrained, to be primarily a leadership role to facilitate and motivate infrastructure protection and a funding role to provide government R&D sponsorship and resources to the private sector. The providers do not believe that the government should play a role that tightly controls or regulates IA technology.

Partnerships

It is also fairly clear from our discussions with these technology providers that there is a critical shortage of expertise in information assurance. Thus, leveraging these scarce resources will be required in any newly initiated information assurance R&D efforts. Government-industry, government-academia, and industry-academia partnerships will be needed to improve the protection of our critical infrastructures.

*CIO Council/
INFOSEC
Research Council*

The government itself must set up cooperative arrangements among its departments, agencies, and laboratories. Coordination, control, and accountability must be established within government. An emerging organization that could facilitate intra-governmental coordination, as well as external coordination, is the Chief Information Officer (CIO) Council. Another organizational structure that could facilitate such

coordination is the recently established INFOSEC Research Council (IRC). An expansion of the IRC Charter and funding would be required.

*Innovative
Approaches*

Innovative partnership approaches will be needed. Some of these partnerships will be technology aligned, others will be application or sector aligned. Specialized government procurement authority may be needed to enter into some of these partnerships. Specialized regulatory relief may be required for some industries to collectively participate. Specialized relief may also be required for the government to partner with a specific technology provider.

Innovative approaches to expanding educational capacities for information assurance should be encouraged. Such efforts should include curricula development, courseware, remote delivery of lectures, and university alliances to share expertise. A mixture of government grants and government-sponsored contract research should be pursued. Leveraged funding opportunities can come through matched funding from the private sector, including commercial industry, state governments, and the universities themselves. Initial funding should be applied to those universities already performing information assurance to broaden and deepen their research programs. Broad Area Announcements should continue to be used as a primary means of soliciting university proposals for research projects. The innovative use of Cooperative Research and Development Agreements (CRADAs) to leverage the national laboratories and university research should also be pursued, especially on sector-aligned research and on issues of scale where simulation facilities can prove useful. The existing CRADA approach will need to be examined for adequacy in supporting such partnerships.

*Pre-Competitive
Technology
Development*

The government should consider using laboratories and agencies to support large-scale experimentation with technology provided by commercial industry. Pre-competitive technology development through government-sponsored research should be fostered between academia and industry. Integrated product team approaches may be highly appropriate in some situations.

Lines of Demarcation To appropriately determine how best to provide mutual support, sharing and control relationships must be established. For example, a demarcation of what a particular commercial infrastructure sector will pursue on its own and where it needs government support must be established through continuous industry-government interaction.

OBSERVATION 3 THE USE OF A NATIONAL INFORMATION ASSURANCE RESEARCH AGENDA (NIARA) IS AN APPROPRIATE VEHICLE FOR COMMUNICATING THE VISION OF WHAT IS NEEDED TO ADDRESS INFORMATION ASSURANCE AS A COMPONENT OF CRITICAL INFRASTRUCTURE PROTECTION.

Government Initiative Government vision backed with resources was seen as a critical component in motivating industry to action. A long-term government initiative will be needed to kick-start critical infrastructure protection and provide sufficient momentum to ensure that the effort can become self-sustaining¹⁵ (i.e., protection must be thought of as a continuous process). Such an initiative must include an organizational component for IA research that is responsible for developing and executing a long-term NIARA.

Oversight Coordination NIARA could be executed in a decentralized fashion but with oversight coordination through the government’s recently established CIO Council. The CIO Council charter may need to be amended to incorporate this new responsibility. The details necessary for an initial long-term NIARA, including funding levels and priorities, would need to be assembled across all Departments for CIO Council coordination and subsequent Departmental budget submissions.

NIARA Components The framework developed by the authors for Finding 5 (page 23) could be used as a starting point for creating an NIARA. Such an NIARA would comprise three categories: (1) basic research in IA fundamentals, (2) system-level security engineering, and (3) individual IA component development. It would cover the spectrum of research needs for information assurance related to infrastructure protection.

¹⁵ The authors believe that such an initiative must last at least a decade to become a self-sustaining process (i.e., one that is inculcated (imprinted) into the culture of the technology providers and infrastructure owners). It should be founded on a three-fold strategy of risk management, continuous improvement, and sound investment.

However, to ensure that individual infrastructure sectors are receiving appropriate attention, these categories should have two focal points: one that is technology aligned and a second that is sector aligned. The technology-aligned focus would ensure that such elements common across sectors are properly addressed. The sector-aligned focus would ensure that such elements unique to a particular sector are also satisfied. Strong interaction among the participants and wide-spread information sharing will be vital to the successful execution of this agenda.

FINAL REMARKS

Information assurance R&D investment is required in three major categories as suggested by the framework previously presented. Research that supports system-level security engineering is the most pressing overall need and will require support from fundamental IA research. However, information assurance R&D investment is impacted by more than a technical research agenda and levels of investment funding. Government policy, commercial competition, customer demand, and the lack of qualified human resources all have a significant impact on information assurance R&D investment. If the Commission is to have a significant and long-lasting effect on our nation's critical infrastructure protection through information assurance R&D, it must take into consideration all of these factors in its deliberations for recommended actions.

APPENDIX A.

KEY TECHNOLOGIES

This appendix elaborates on the IA research framework identified in Finding 5 (page 23). The appendix identifies key technologies that commercial providers believe are important areas of IA research and/or areas not addressed by previous or current research. Most of these providers addressed information assurance from the perspective of their specific technologies and only generally addressed the broad needs of protecting the nation's critical infrastructures. Their responses did not target the needs of any specific sector's infrastructure.

Most of the technology providers believe that government will not be able to pick the research that will be most relevant to any specific commercial products, especially their own products. Therefore, in setting forth their views on IA research, the providers generally believe that government can be most beneficial in defining the hard problems to be solved over the long term and in soliciting participation and/or proposals to gain and fund the best set of available ideas and capabilities to address these problems.

Top research needs identified by these technology providers have been organized under three categories shown in Table A-4:

- Basic research in IA fundamentals, where availability and integrity are the critical areas to move information assurance from a craft to an engineering discipline;
- System-level security engineering, the most critical need; and
- Individual component development, where security management and intrusion detection were the most immediate needs.

Table A-4. Top IA Research Needs

Basic Research in IA Fundamentals	System-Level Security Engineering	Individual Component Development		
Protection Concepts & Principles (page A-3) <ul style="list-style-type: none"> • Availability • Integrity 	System Architectures (page A-14) <ul style="list-style-type: none"> • Existing Architectures • Global Architectures • Enforcement Allocation 	Security Management (page A-23)	Applications (page A-34) <ul style="list-style-type: none"> • Databases • Distributed Directory Services 	
System Complexity Issues (page A-9) <ul style="list-style-type: none"> • System Dynamics & Adaptability • Composability • Security Economics • Intuitiveness 	Heterogeneous Component Integration (page A-15) <ul style="list-style-type: none"> • Applications Perspective • Seamless Integration • Supporting Management Tools 	Intrusion Detection (page A-23) <ul style="list-style-type: none"> • Attack Taxonomies • Correlation • Adaptivity • Faster Deployment • New Locations • Tiered Structuring 	Secure Operating Systems (page A-35) <ul style="list-style-type: none"> • Fine-Grain Object Support • Better Server Engines • Distributed Authentication Support • Domain Support • Label Support • "Trusted" Operating Systems • Directory Systems 	
Vulnerability Analysis (page A-12) <ul style="list-style-type: none"> • Protocol Analysis 	Secure Interoperability & Evolvability (page A-16) <ul style="list-style-type: none"> • Evolvable Integration • Understanding Changes 	Identification & Authentication (page A-26) <ul style="list-style-type: none"> • Anonymity • Middleware • Biometrics 	Applied Cryptography (page A-36) <ul style="list-style-type: none"> • Algorithms & Protocols • Usability & Trust • Scalable Certificate Authorities • Key Recovery • Encryption Chips • Autoimmunity • Crypto-seals • Long-term Key Management 	
Trust Concepts (page A-13) <ul style="list-style-type: none"> • Defining Trust • Risk Management 	Applied Engineering Research (page A-17) <ul style="list-style-type: none"> • Robustness • Protocols • Security Mechanisms • Composability • Integrated Security Analysis 	Smart Cards (page A-27) <ul style="list-style-type: none"> • Cryptography in Smart Cards • Optics • Biometrics • Tamper-proof Devices • Secure operating system 	Hardware-Based Security (page A-40) <ul style="list-style-type: none"> • Component Design • Virtual Machines • Power vs. Speed • Design Verification • Manufacturing 	
	System Assurance (page A-18) <ul style="list-style-type: none"> • System-level Assurance • Operational Assurance • Quality Approaches • Scale of Assurance 	Networking (page A-29) <ul style="list-style-type: none"> • Active Network Component • Programming the Network Fabric • Protocol Policies • Service Associations • Robust Error Recovery • Establishing Trust Across Networks 		<ul style="list-style-type: none"> • Nomadic Computing • Virtual Private Networks • Non-disclosed Protocols • Optical Networking • Advanced Firewall Technology
	Standards (page A-20) <ul style="list-style-type: none"> • Enabling Standards • Cryptographic Standards 			

A.1 Basic Research in IA Fundamentals

Wherever there is the potential for large financial payoffs, industry will address the problem. Basic research is perceived by this sample of technology providers as addressing hard IA problems. They also see it as more oriented on doing something novel in information assurance (e.g., the DARPA Firewall Toolkit) as opposed to figuring out the Java structure and engineering a secure solution for it. It is the belief of many technology providers that the nation will win more through basic research over the long haul.

The technology providers believe there is a need to broaden the array of fundamental research in information assurance. With the massive changes continuing to occur in telecommunications and computing, gaps in security technology will emerge and must be addressed. New opportunities will also emerge that can be exploited to enhance security.

Also needed is a thorough review against new base technology of what has been tried before to achieve information assurance. A security “genealogy” that links together the evolution of security technology could prove useful. The community should look to past successes and failures in approaches to security and apply the lessons learned to current situation. Ideas that were ahead of their time 15 or 20 years ago may now be the right ones for today, e.g., virtual private networks. The change in technology performance today may enable such ideas to become fully workable.

Protection Concepts & Principles

Pursuit in establishing fundamental protection concepts (e.g., isolation, secure association) and principles (e.g., least-privilege, mutual suspicion) for information assurance must continue. These should drive system-level research, component research, and the products that use that research. Fundamental research will require:

- Increasing collaboration among practitioners in various fields as information assurance must become a multi-disciplined field,
- More of a system focus,

- Work towards improving the capture of operational security requirements, and
- More scenario-based motivation in IA research.

Interconnected systems have changed the paradigms for using computers securely. Enterprise-critical processing now implies the urgent need for availability and strong integrity. More research is needed in these two areas as they are now a critical part of system-level security engineering.

- Availability is a hard problem that needs to be addressed. It encompasses reliability, resource allocation, and recovery. Not having it will be extremely detrimental as we continue to become more dependent on networked computing. A sense of urgency should be assigned to availability research.
- Integrity is a very significant problem area that should be given a high priority for additional research as it is needed for enterprise-wide computing. Critical mass to address it is lacking. Little or no research work is going on to get wide-spread solutions.

Availability An extremely vital area of needed research is in the availability of data and information services. It is an area that is not getting sufficient research attention. Confidentiality and integrity issues can be handled by users through a variety of mechanisms. But having data and information services available when needed (“on demand”) is becoming a critical problem for businesses.

Protocol
service
components

Future protocols will require a “Quality of Service” (QoS) component. We may have many different networks, each providing a different “Class of Service” (CoS). Speed and performance will be key drivers in this area. There may be a need for prioritized service for crisis situations. Increased service may cost more than routine service. As in a toll road example, there is an expectation of better driving conditions on a toll road than on a typical city street. What does it mean to write applications that work well with various impedance mismatches and unreliability of other components? We need availability mechanisms which include

fault avoidance [reliability], fault tolerance [reconfiguration], and recovery mechanisms.

Reliability Reliability, as a component of availability, is generally well researched. However, reliability in large, networked systems with secure and heterogeneous environments has not been well researched. There is a strong demand for reliability in telecommunications. More research is needed, especially with regard to scale and heterogeneity. Mixing and matching of components make availability more and more difficult. Design of very complex systems with different technologies makes it hard to understand all the interactions that can occur. One provider is doing some work in examining feature interactions in software (e.g., unanticipated interactions). To fix such problem interactions, however, ideally requires complete control (not practical in heterogeneous environments) or improved collaborative approaches. At least one provider is doing some work on fingerprinting to help identify cause and effect.

Reconfiguration Reconfiguration, as a component of availability, includes operational optimization dynamics, fail-over, fail-safe, and degraded operational continuity.

- Operational dynamics are an aspect of normal and contingency operations. Such mechanisms continuously (re)allocate and/or (re)schedule system resources based upon static and/or dynamic policy (e.g., CoS or QoS, shared-resource user agreements, or contingency-based degraded operational service priorities). This aspect of availability requires significant research as assurance with dynamic behavior is not well understood. Some of the providers have one or more initiatives in QoS research. These include adding QoS to routers. RSVP (Estrin)—originally thought to be a significant part of the answer—is now thought to be too heavyweight for practical solutions. One provider is also working QoS in association with mobile and cellular research.
- Fail-over provides operational replacement spares for continued, non-degraded operations (e.g., non-stop behaviors) through added robustness (redundancy) of components in a system.

- Fail-safe provides the capacity to do no harm through controlled shut-down or reconfiguration of a system or a portion of its components.
- Degraded operational continuity is achieved through contingency-based (re)allocation of insufficient operational resources to meet some minimal operational capability.

Recovery Recovery, as a component of availability, is “response to catastrophic failure” and requires more research. How do we deal with recovery issues after catastrophic failure of a large system, system of systems, or network? We need more research in denial-of-service issues and the area of priority messaging. There is a need for improved response management to provide restart and recovery capability. This capability should be incorporated in the new “sealed PC environments” as part of lowering the cost of administration and maintenance.

Information availability Availability research is defined by some in a slightly different way. They view a person’s inability to find needed information on the Internet as an availability assurance problem. Intelligent Internet search engines are needed to give individuals a greater degree of assurance that they can find and access critical information in a timely manner. Current presentation protocols do not handle queries as well as the state-of-the-art databases do.

Outsourced assured sites Another availability viewpoint needing further exploration is the concept of “Distributed Leased Secure Storage” provided by a third party. Having replicated data and information services at “assured sites” may provide the ultimate goal. Secure fault-tolerant computing using distributed “leased” resource pools may be an approach toward achieving this assurance.

Integrity Integrity is a major research issue that needs to be addressed with a high priority. Data and systems integrity remain critical, unsolved problems for information assurance.

Critical problem Much has been done but much remains to be done. This is especially true as we proceed to do more over networks. One of the overlooked yet more important issues is how to (efficiently) discover what has actually

occurred in an attack (e.g., imagine random damage throughout a critical database and its backups over an extended period of time) and to safely recover.

Cryptographic-
based
integrity

Cryptography has become an important enabling technology for data integrity. The technology exists for digital signing and sealing of data to preclude use of invalid or disrupted data. The technology exists for non-repudiation of origin and receipt. Some of this technology simply needs to be deployed while other parts still need advanced research.

Some providers have incorporated research results in cryptographic-based signing to increase systems integrity. One hardware provider includes signing of components that come with pieces of equipment and those that are added dynamically. The provider will use this as part of active fault management.

Internal
controls

One critical area is the need to increase research on automated mechanisms to enforce the variety of mandated or regulated integrity requirements via internal controls. While currently most applicable to financial services, such controls also are required in many industry processes, including nuclear power. Bridges from various applications to common enforcement mechanisms may be provided by workflow or similar process control systems with embedded integrity policies. Such approaches need further investigation.

Secure
software
distribution

The ability to “lock-down” code on licensed software on a particular machine is needed as is secure software distribution. Distribution could be done easily by Net Layer Modules (NLMs). However, we do not have a secure loader. A rogue NLM from one provider has demonstrated that another provider was insecure in using this method. We need improved digital watermarks—attribute information that is saved regardless of the state of communications.

Intellectual
property
protection

A similar critical need in the commercial world is intellectual property protection. Intellectual property protection is now beginning to have some integrity technology results that are promising (e.g., watermarks, tracking traders). More research work is needed. It should include involvement of legal experts and law enforcement officials. Such research could include:

- Integration of multimedia (more than what has been provided via Secure Multipurpose Internet Mail Extensions (S/MIME)).
- Self-protecting content.
- Better anti-piracy solutions.
- Electronic distribution support for the concept of “copy ok but to play any copy will cost.”
- High-value Digital Versatile Disk (DVD) players with assurance that recording at end-point is not occurring.
- Fully signed applications and components (not only Active-X type components but entire configurations must be signed).
- Document control and licensing with a “breakable” key (i.e., one that can be destroyed) which could expire all copies and ensure that they are destroyed.
- “E-Control” (i.e., the capability to send back or destroy documents).

Anti-virus

Integrity research includes the need for more anti-virus research. Auto-immune systems are one approach being investigated to guard against viruses and hacking attacks. This approach makes it possible to detect new (previously unseen) viruses and hacks. Another approach is cryptographic sealing.

Archival
integrity
controls

One integrity area that may be overlooked is archive controls. An interesting research topic is the long-term archiving and life cycle management of digitally signed documents. What happens when the certificate expires? Another is media integrity. The loss of valuable data or privacy can occur when changing out the media being used for archiving. An example of U.S. Social Security data being available within Hong Kong was used to illustrate the concern regarding the loss as a result of not having such controls. Discarded, scrapped, or resold media were not cleaned appropriately before it was disposed by the original owners.

System Complexity Issues

An orthogonal view of critically needed fundamental research is the area of system complexity. System complexity touches all aspects of information assurance. We need research on understanding and dealing this complexity. It is vital to being able to perform system-level security engineering. System-level security has never been well defined. There are hard problems to be attacked. Understanding scale is critical. What can you do at various sizes of networks? When do the views of the problems change? This research encompasses system dynamics and adaptability, composability, security economics, and intuitiveness, and includes:

- More fundamental research on security from a system's view to clarify what is meant by security in large systems.
- Understanding aggregate behavior of a large number of systems (i.e., a "system of systems" perspective).
- Development of new behavioral models and analysis tools.

System Dynamics & Adaptability

There is a need to do more work on system dynamics and adaptability. An increasingly important facet of this analytic need is more work on safe executable content and "plug-n-play" software that interoperates with high confidence.

Dynamic security policies

Adaptable policies are an essential area of work where new research should be focused. There are many interesting opportunities emerging. Dynamic security policies will be needed to handle the variety of contingencies organizations will face in doing business over the Internet. Requirements include:

- Pluggable policy support. Different customers have different needs (e.g., adaptable or flexible solution sets).
- A single (adaptable) solution to meet any customer's need. They need a common product to buy. Adaptable modules enable us to meet such needs.

- The capability to provide such solution sets more easily. We need better approaches to adaptable policy representation and to understanding policy composition and interaction.

Adaptable multi-level security

Several technology providers thought that there will ultimately be a huge demand for multilevel secure (MLS) capable systems. This is both a policy and technology issue. For the Department of Defense, coalition warfare will require increased emphasis on “adaptable” MLS. It will require sharing of information at different sensitivities in an efficient manner. It must be able to adapt to changing circumstances rapidly. This sort of adaptability is certainly not available even as state of the art and will require significant R&D efforts.

Composability

Significant research is needed in policy composition. This research should include:

- Work on theoretical and practical composability.
- Addressing the issue of policy and enforcement allocation. We must understand how to allocate and integrate security mechanisms with convincing evidence that the result is right. Different components enforcing different policies require this allocation and composability analysis.
- Work on policy representation and models. What does a system-level policy look like? How do we assess policy composability? How do we apply security principles and objectives to such assessments? What are the trade-offs that should be emphasized?
- Work on a technical view of policy as applied within a system. What are all of the components to an enterprise-level security policy? What are the broad areas that must be fused together?

In addition, some of the possible long-term research that ought to be conducted in this area include dynamic security policy specifications and policy support for dynamic nets (wherein packets are routed dynamically based on content or other specified parameters).

- What are the security properties desired in a shared infrastructure? (For example, adaptable, parameterized common mechanisms in network protocols, hosts, and lower layers.)
- Where does commonality of enforcement mechanism allow for use of adaptive policy parameters?

*Security
Economics*

A significant area of system complexity research is security economics where extensive fundamental work is needed. There is no “silver bullet” for information assurance. System security is a collection of security components that includes security at the perimeters, the servers, and the desktop. We must be able to add in and integrate more security with a good understanding of the achievable risk reductions. How to trade off security and performance must be understood. This requires a better understanding of the economics regarding where we distribute our security among these components. We need to have more quantified data on performance and risk reductions. Establishing the economic understanding is part of doing a better job in making security intuitive. An understanding of what to measure is important. Process activity costs must become visible. Economics are required for addressing the cost of security technology ownership. Example questions:

- How to cost? How to reduce?
- How much does a hacking attack cost?
- How much does recovery cost?
- How much does an investigation cost?
- What technologies radically improve the efficiency and effectiveness of security procedures?
- What IA technologies are the most costly?
- Where can IA process change occur?
- What technologies can leverage lower-skilled personnel? Human resources with significant security expertise are not and likely will not be available.

- What technologies can reduce the requirements on the numbers of personnel needed to manage and monitor security in large-scale systems? There has been little investment directly made into security tools.

Intuitiveness Another important area of system complexity research is the issue of human factors and information assurance. Security must have more intuitiveness. This is a fundamental research issue that must be addressed both at the system-level and within individual IA components. We must remove as many steps required of people as possible to improve the strength of assurance. An analogy of physical keys being left in doors serves to illustrate the need to reduce the process steps. Entry combination-key pads and proximity or swipe devices are technologies that remove that hazard by changing the steps required.

Vulnerability Analysis

Fundamental research is needed in system vulnerability analysis. We need to answer questions such as how to assess, how to describe, how to share. Also needed are analysis tools and supporting databases. An insufficient amount of time has been spent looking at vulnerabilities. This area needs more attention. Much more work is needed in vulnerability analysis as part of system integration. We must understand how to place our trust in critical elements.

Protocol Analysis

Research is needed in protocol analysis and should include:

- Developing and understanding the models for secure protocols at various layers (e.g., applications layer, middleware¹ layer, communications layer). We need to go up a level of abstraction in protocol analysis to look at classes of protocols (mail, file transfer, crypto, etc.) and to develop threat models. The examination of protocol classes should be through more rigorous analysis over larger sets of threats and classes of protocols.

¹ “Middleware” is a term used herein to collectively identify general-purpose service software with common software interfaces and protocols sitting in a layer between a computing platform (i.e., hardware plus the operating system) and its user-specific applications. Examples include object request brokers, distributed computing services, message translators, directories, and database management systems.

- Methods and analysis tools for secure protocols. We need analytical tools that are more oriented toward secure interoperability and secure composition of protocol stacks for communications.
- Engineering principles on how to build protocols in secure functions.

Trust Concepts

Defining Trust Defining and understanding trust is fundamental to information security. We are striving for a trustworthy relationship between machine and humans. Decisions having dynamic paths are generally made by humans; machines generally execute fixed or rule-based instructions with predetermined decision paths (i.e., decisions possessing regularity). What are the required elements of trust in a system? As components become increasingly dynamic, where are decisions made and policy enforced? What are the parameters of interest in establishing trust in these decisions and their enforcement? How do the values of such parameters create different levels of trust? What are the elements of distributed trust?

Risk Management Measurement is key to enabling management. It is vital to achieving trust. We must do more in the area of measuring protection levels and associated risks. Risks must become well understood. Security in commercial technology is being driven by electronic commerce. However, the security community and financial communities view things a little differently. In the security world, things are most often viewed as either secure or insecure. In the financial community, the organization looks at how much financial exposure it has through various forms of risk (e.g., credit, operational, systemic) and makes necessary system investments and/or procedural adjustments to ensure that such exposure remains within a given threshold. While being careful of analogies that do not hold, the security community should attempt to understand the risk models used by the financial and business communities (i.e., models employing cost-benefit analyses).

A.2 System-Level Engineering

System Architectures

Existing Architectures

Many respondents believe that the bulk of work needed is in the area of architectures. While there are several different types of security architectures in existence, it is not yet clear that they are sufficient to cover the foundational need for secure system interoperability and evolvability.² Better paradigms for securely interoperating are needed. Intel’s architecture—Common Data Security Architecture (CDSA)—has been so recently released that many have not had a chance to evaluate it. One of the more experienced technology providers who looked at Intel’s CDSA stated that it is going to be extremely important to future system security.

Most providers believe that more architectural research is critically needed. Information assurance architectural work being pursued out of DARPA’s Information Systems Office was cited by another provider as an important example of the government pursuing such research. This work is examining all available architectures and attempting to utilize the best approaches from each to lead towards a demonstrable prototype system architecture.

Global Architectures

Global architectures are the most critical aspect of architectural work to be pursued. These architectures include multiple trust hierarchies for the global employment of cryptography through Certificate Authorities. This means more must be done in architectural policy issues and standards as well as in system-level security engineering.

Enforcement Allocation

Continued architectural work is needed on how to distribute and allocate security mechanisms in networked, distributed systems. Research is needed to address the concepts of layering and security services, including a focus on advancing the “trusted subsets” and “trusted partitioning” work that was done by the Department of Defense as part

² For example: Intel’s Common Data Security Architecture (CDSA); Open Group’s Common Object Request Broker Architecture (CORBA) Security Architecture; and DoD’s Goal Security Architecture (DGSA).

of the Trusted Computer System Evaluation Criteria (TCSEC)³ and its interpretations.

Heterogeneous Component Integration

<i>Applications Perspective</i>	Within system-level security engineering, there is a need to take more of an application view of security integration. What needs to be monitored within an application? What needs to be monitored within clusters of workstations? Non-stop behavior (e.g., fault tolerant, continuous operation) is an area that needs larger investments. Improved hardening ⁴ of the machine is important to applications. This includes the ability to restore to a particular state, no non-replicated state, and large amounts of locally cached state.
<i>Seamless Integration</i>	The need for more seamless integration is perceived as critical. We must know how to glue components together securely, with a small footprint. Current “tubes of glue” (e.g., Open Group’s Distributed Computing Environment (DCE)) are perceived to have too large a footprint. Today’s system architectures need to couple together guards and firewalls, remote authentication, and web-filtering products to enforce an enterprise-wide security policy. The architecture must support operational security monitoring. Key focal points for integration of such monitoring are network devices and protocols. The architecture needs to be easy enough to use that a network manager can use it without significant burden.
<i>Supporting Management Tools</i>	This implies a need to provide integrated supporting management tools. Simplicity is key (e.g., the analogy of automobile dashboard instrumentation). CORBA and the Distributed Common Object Model (DCOM) are becoming more important. Type-enforcement support is an important low-level mechanism. Other needs include separation mechanisms and type enforcement put into object request brokers (ORBs). Also needed are application-specific integrity and smart filters

³ *Department of Defense Trusted Computing System Evaluation Criteria*, DoD 5200.28-STD, 1985.

⁴ Hardening could include adding new availability or integrity monitoring and enforcement mechanisms to strengthen the properties of system availability or integrity.

on uniform resource locators (URLs). With the HyperText Transfer Protocol (http) presenting internetted computers as a virtual machine, applets will become like “toasters”—ubiquitous. Efficiently managing their security will be critical.

Secure Interoperability & Evolvability

Evolvable Integration

The basic theme of new research in system-level security engineering and in component development should be *evolvable integration* (i.e., engineering integrated system security in a fashion that enables security component replacement and/or upgrade with minimal effect on the overall integration of security). Standardized, adaptable, and flexible interfaces and protocols are key to evolvability and integration. We need more security integration at all levels of technology. The concept of evolvable integration should become fundamental to our system-level security architectures.

Understanding Changes

This aspect of system-level engineering requires understanding where things are changing and what technological base best positions a system to evolve. Common extensible and adaptive mechanisms will be required to evolve securely. Insertable security or adaptation mechanisms (e.g., wrappers⁵) will be required where legacy components must be continued. An example of such engineering requirements stems from changes in communications:

- Communications speeds are increasing. What can be done with firewalls? What cannot be done?
- Communications are moving increasingly into mixed media (e.g., land-line, cellular radio, satellite). How can these various media be advantageously used? What are the security problems such mixed media present?
- Communications will be pervasively enabled by low-level hardware components (i.e, digital signal processors co-located with microprocessors) and adaptive software components. How will such components be used or upgraded securely?

⁵ Wrapper technology is software that can encapsulate some part of a legacy system—without modification to the legacy part—and provide new IA properties.

- Communications will be providing content-based routing. How will more complex content (e.g., multicast video conferencing) be securely managed? How will secure electronic distribution of software be efficiently provided? How will misuse detection be handled? Are adaptable policies for firewalls and their successors going to be required?

Applied Engineering Research

This is applied work, not conceptual work. We need more prototypes that demonstrate workable solutions, not research papers. Fundamentally, without demonstrations, the research will not transition easily. We need to have system testbeds that allow various concepts and products to be tried out in a system context. It may be possible to use the Next Generation Internet (NGI) as a vehicle for such demonstrations. DARPA and NSA also have the ability to field such security technology demonstrators. Results from systems-level security engineering research should include producing system security guidelines that can be used by others to build and securely integrate system and security components. Five significant areas needing applied engineering research are:

- | | |
|-------------------------------------|--|
| <i>Robustness</i> | <ul style="list-style-type: none">• Robustness of mechanism (applied reliability engineering, e.g., fail-over, load-balancing in firewalls, highly available secure communications). |
| <i>Protocols</i> | <ul style="list-style-type: none">• Protocols (developmental engineering of secure protocols at various layers (e.g., applications layer)). |
| <i>Security Mechanisms</i> | <ul style="list-style-type: none">• Placement of security mechanisms (performance assessments of global architecture). |
| <i>Composability</i> | <ul style="list-style-type: none">• Composability (policy-driven, practical application of composability-analysis theory). |
| <i>Integrated Security Analysis</i> | <ul style="list-style-type: none">• Integrated security analysis (identification and rigorous assurance of those components that matter). |

System Assurance

Most of the niche providers interviewed believe more work in assurance is needed. More precision is needed in the definition of assurance. Rigorous scientific analyses should be used as a means to provide high assurance. A long-term, important research area where the government may want to invest is in assurance metrics for computing products and in large-scale systems. Nearly all of the technology providers want research on how users can gain needed assurances without today's costly product evaluation.

The Trusted Product Evaluation paradigm⁶ that has been in use for two decades is no longer realistic. Although the intent of the product evaluation program was to evaluate a product once and then reuse that product's evaluation assurance evidence in certifying many different system implementations, the actual results have not proven to be as effective as intended. Most product providers believe that continuing to evaluate a single product is irrelevant today because that product often must be adapted to participate as part of a larger system. Distributed systems are highly complex and the current criteria used for evaluation are outdated. The current costly product evaluation approach does not adequately support system-level needs. Most technology providers assert that the process must be changed if the government is to have any continued partnership role in assessing the assurances of commercial products.

System-level Assurance

Extensive product evaluation to gain assurance before system installation may be premature. Perceived system-level needs include:

- The ability to objectively determine what a product does (e.g., claim validation) in a process that is highly repeatable and relatively impervious to a different interpretation.
- Reasonable tests for the presence of security functionality.

⁶ TCSEC, DoD 5200.28-STD, 1985. The Trusted Product Paradigm is built on four fundamental blocks: (1) security policy (access control), (2) accountability (identification & authentication; audit), (3) assurance (operational and life cycle), and (4) documentation. Such trust has been oriented largely toward the operating system and its Trusted Computing Base (TCB).

- A basis for finding security and system behavioral properties.
- An understanding of the security contributions of various components and how to compose them.
- An understanding of the “weakest link” within the system at the component level. If we have strong components throughout most of the system but have an exploitable weak link, then one has to ask what is the worth of all that strength.

Operational Assurance

Operational assurance will be increasingly important. We need more penetration-analysis efforts. The government should continue to partner with the critical infrastructure providers to help beef up their infrastructure security from an operational point of view.

High assurance issue

Many providers believe high assurance is still at the starting point. Both MLS and the DoD approach to high assurance have peculiar solutions. One provider pointed out that the differences in formal design specifications vs. what actually happens in the implementation have shown the “foolishness” of the TCSEC high assurance approach. Formal verification still requires more expert people than tools. Existing tools have not led to more cost-effective formal method solutions in the area of software. There is not enough assurance done at a level of scale tied to real product development. Exceptions include Rushby’s work with SRI International’s Prototype Verification System (PVS). Formal methods technology, in particular model checking, has been very useful in the area of hardware.

Integrating assurance into engineering

Research is needed in assurance technologies, tools, and processes. We must bolster our engineering teams and engineering processes to make assurance methods a natural component of hardware and software engineering. Engineers need better design analogies and better integration concepts that support assurance argumentation. These all lead to the idea of rigorous design languages that allow the formal expression of a concept and/or behavior. Most still believe that the key to high assurance lies in formal methods integration into software development tools.

*Quality
Approaches*

More research is needed on applying quality approaches to software and hardware engineering processes. Evidence exists that those technology providers applying various quality approaches are beginning to achieve significant cost and assurance payoffs.⁷ Software production needs more attention. Some providers believe that we need to better integrate the software development Capability Maturity Model work of Carnegie Mellon University's Software Engineering Institute into our engineering processes. The use of computer-assisted software engineering (CASE) tools, quality standards, standard procedures (including quality analysis), and software reuse in product migration add to the quality of software. Several providers are spending more resources on such tools and procedures, and most believe much more capability must be developed.

*Scale of
Assurance*

Establishing assurance in large-scale systems is a hard problem. Such assurance must come from quality engineering, including design and test, functional components that enforce some property of assurance, and operational monitoring. One provider emphasized that results exist in globally distributed centers from more than thirty years of research work on assurance and that the nuggets of wisdom should be plucked from this body of work and reapplied to today's needs. An increasingly important area of research within the issue of scale is developing tools to deploy large numbers of computing resources with assurance.

Standards*Enabling
Standards*

Standards are key to achieving information assurance in our critical national infrastructures. Standards are also key to the economics of deploying IA technology. There is a need to attack this area with some urgency. Applied research efforts to establish much needed standards across a spectrum of IA technologies were deemed important by all the respondents. Their customers need standards (*de facto* or consensus engineered). Standards provide a baseline for the community and are a mirror of a *strategic* way to solve a problem. The technology providers asserted that we need better "enabling" standards (e.g., X.509v3, IPv6).

⁷ Christopher Fox and William Frakes, "The Quality Approach: Is It Delivering?," *Communications of the ACM*, Vol. 40, No. 6, June 1997, pp. 24-29.

Without such standards, electronic mail cannot be protected and IP security would be impractical—precluding the ability to protect routing on the Internet. Many such enabling standards are needed for successful protection deployment. Heterogeneity will be a problem when we have a wider array of products. Without common standards, we will not be able to properly exchange security attributes, and enforcement will be much more expensive for administration and performance.

*Cryptographic
Standards*

Standards in cryptography will be key to technical success in information assurance. Standardized cryptographic authentication is critically needed for tying things together at the system level. Some providers really believe that it is critical to have *one* standard for cryptographic authentication. All these vendor-produced “standards” create more confusion. We should not have more vendors going off and producing yet even more such *ad hoc* standards. Leadership in standards making is needed. For example, there is controversy on who is in charge of the civil sector public key infrastructure (PKI). The General Services Administration claims ownership.

Public key
infrastructure

One example of needed standardization is a PKI for both the military and civil sectors. The PKI requirements are different in each case. Common approaches need continued research. Additional research in Simple Distributed Security Infrastructures (SDSI) may also prove fruitful (similar to the work being conducted by Lampson at Microsoft and Rivest at the Massachusetts Institute of Technology). More work is needed in the area of Simplified Public Key Infrastructures (SPKI) as well to support the growing demand for secure electronic commerce.

Emerging
standards

There are standards close to completion to help solve the PKI problem: IETF-PKIX, ANSI X9.57, and the X.500 Series. Among emerging standards involving cryptography that are important to electronic commerce is the Secure Electronic Transaction (SET) standard. Another important standard is X.509 digital IDs binding a person to a public key. To facilitate one provider’s objective of supporting role-based authentication, the X.509 standard needs to be expanded to include the ability to bind a person to a role. Standards are needed, in general, for enterprise-wide roles.

Due care standards An interesting and important area that should be considered is “licensing” or providing standards of due care and conformance parameters. The National Research Council study, *For the Record*, provides a more detailed discussion of this concept in dealing with the protection of electronic medical records.⁸

⁸ National Research Council, *For The Record: Protecting Electronic Health Records*, National Academy Press, 1997.

A.3 Individual Component Development

Security Management

An important and often ignored area is security management (administration and monitoring). Security management is one of the most critical areas of component development to be pursued. It needs a great deal of attention. Cost of ownership is significantly affected by our technical solutions to security management. The important research topic is in the area of scalability of security management to very large systems. For example, issuing more than 1,000 root passwords for systems controlling large power grids and distribution centers is a major vulnerability. Research must focus on:

- Reducing the cost curves and raising the effectiveness curves of security management. Security management tools and procedures must address scale in reducing operational costs.
- Secure approaches for remote security administration and monitoring.
- Engineering development of security management tools and databases. A critical component of security management is policy management, which requires operating at the right levels of abstraction with tools to compile stated policy into enforcement mechanisms or their supporting security management information bases (SMIBs).
- Integration of the local enterprise security management environment with the network intrusion detection monitoring environment and the network management environment.

Intrusion Detection

An important area for additional government R&D investment is in intrusion detection. There was a very strong endorsement from the technology providers for this area of research. This is an area that many technology providers will not invest in research but could benefit from any research results. Research that takes a more comprehensive view of intrusion detection is needed. Remember it is not just the lock, it is the

“applications” around it that provide security. More than point-product solutions are needed—firewall “solutions” are insufficient. A security policy that supports this larger network monitoring viewpoint is required. Most providers believe more work is needed in large-scale network security monitoring and intrusion detection. Some believe that a consortium of companies working on intrusion detection systems (similar to what has been done for cryptographic public key management) is needed.

*Attack
Taxonomies*

Organizations must be able to react to real-time events such as the “Morris Worm” and other malicious software attacks. Attack taxonomies are needed; for example:

- Distributed coordinated attack (N -> 1)
- One against many machines (1 -> N)
- Distributed coordinated attack against many machines (N -> N)

Some providers are beginning to identify new exploits as well as continuing to understand variations in buffer overflow attacks. More research in hacker techniques similar to those of Gene Shultz of SRI International (I-4 program) is needed. Schultz is working on how people hack and how to prevent and/or respond. We need intrusion detection research laboratories serving as the “thinkers” as well as laboratories that are more pragmatic (hands on)—and these laboratories must interact.

New capabilities are needed in intrusion detection systems: correlation, adaptivity, faster deployment, new locations, and tiered structuring.

Correlation

Current intrusion detection systems (IDSs) are not particularly effective for large-scale systems. They lack correlation capabilities, use console-level type tools, and are hard to integrate into systems management.

Adaptivity

More adaptive, signature-oriented IDSs are required. Several IDS technology providers are going after such adaptive IDS. Some believe they are still about two years away from such a capability and that more research is needed.

*Faster
Deployment*

How do we get operational IDS capability faster? Primarily it will occur through more writing of software by those who understand the problem at hand. For example, there is a difficulty in writing attack signatures. Poor signatures can cause more false positives. To write network-based signatures, one must know TCP/IP inside out. This requires more training, equipment, and better deployment approaches. To respond faster to new attack capabilities, the IDS technology provider's overall resource demands on software vulnerability research must be lowered. This requires:

- IDS software written in a modular fashion that is easily upgradable.
- More work on shared vulnerability databases that can provide information on both how current software is being exploited as well as how to defend against such attacks.
- More research on IDS software that works with networking devices.

New Locations

Some providers believe intrusion detection should be merged with the capabilities of switching to send data to network intrusion monitors. However, it is unlikely that promiscuous-mode devices will be available on switches. Speed will be an issue. Placing host-based intrusion detection on key servers will be required. There is also a need to develop data-mining techniques for monitored data from various locations.

*Tiered
Structuring*

Tiered structuring and management of intrusion detection systems will be needed to achieve scaling. At one tier will be management of hosts. This is detection to/from hosts and doesn't get into the network. At the next higher tier will be firewalls or management of enclaves. A third tier above enclaves will be network infrastructure monitoring systems. This illuminates the intrusion scanning issue: large scale, secure control, network security discovery. How do you do this correctly and efficiently? How do you establish responsibilities and priority of navigation and observables for the system IDS Director, the Network-Remote Server, and a local area network (LAN) of workstations?

Identification & Authentication

Authenticated identification is important. It provides the linkage that ties accountability to behavior. Most technology providers believe that there needs to be additional research in the area of authentication, including:

- Approaches to drive the costs of this technology lower and increase its reliability.
- Expanding the application of this technology throughout the system.
- Going beyond the traditional concept of authenticating an individual person and moving into the area of authenticating “roles” within an enterprise or an organization.
- Composite identification schemes. A bottom-up and not a top-down approach to identification is needed to cast multiple personalities; also needed is a smart way of connecting an individual to a signature.
- Applying a combination of security mechanisms for authentication. Smart cards, cryptography, biometrics, and plug-n-play hardware and software components (e.g., proximity devices, proxy agents) will be technologies that will support authentication.
- Authentication of system parts at various levels. Hardware may need to authenticate specific components based on uniqueness or on identical commonality. Hardware support may become a central component for efficient authentication. Applications will need mutual authentication with services in a middleware layer. Clients will need to authenticate to servers in distributed systems. Certification Authorities will require cross-boundary authentication.
- Developing additional mechanisms for “name-space” (personal identity, machine, domain) and “time-space” correlations. This includes ordering, clock synchronization, and secure network time or pseudo-time.

Anonymity Some providers believe that stronger authentication will be available but will still not be widely used. They believe that the stronger demand in the

business world is for anonymity (i.e., membership based vs. personal identification). Credit cards are such a membership instrument—they place individuals into credit classes. With respect to risks, checking has low losses and is highly personal while credit cards have a fairly high margin of loss (paid for by membership fees at both the consumer or merchant ends). Credit card companies still prosecute fraud at a certain threshold but not all fraud. It is clear that the convenience of credit cards outweighs the risks. That is the business case.

Middleware Middleware is a place to focus. First, we must do no harm! We will continue to have tensions on compromising personal privacy. Middleware can be compared to service-oriented merchants. Merchants, like the hotel concierge, can serve us in a much more tailored and satisfactory manner if they know more about us. This does not necessarily equate to personal identification, but rather to membership (e.g., roles and groups). There is a need for more capability to provide a membership number that is not directly associated with name and address, but that can be indirectly matched to complete the service application. Several providers believe that public key technology to do this is too expensive—a shared key that expires daily, or more frequently, is more desirable. A distributed authorization mechanism, such as Kerberos is desired but with lighter weight (i.e., improved performance). There should be no user choice in passwords. Best practices must be used.

Biometrics Biometrics remains a potential growth area and needs additional research. We need improved biometrics integrated within identification and authentication sub-systems in a practical operational manner.

Smart Cards

Ambivalent opinions still exist regarding the usefulness of smart card research. Many believe that one of the biggest gaps in U.S. security research today is in smart cards. Foreign research dominates this smart card technology, though not necessarily the security technology that could be placed within them. Currently the Europeans are leading in producing and using smart card technology. Many of the technology providers believe that this technology is key to building bottom-up

solutions to system-level information assurance and they recognize that more research remains to make IA enabled smart cards viable. Smart card technology needs applied research as well as education. Educated people are needed to implement this technology.

Other technology providers believe that smart cards are not the solution—they believe we will leap over such technology and be wearing smart devices or carrying proximity tag-alongs.⁹ Most agree that smart card technology built into keyboards is attractive in the near term, but some still see this technology as remaining below the line of needed IA investments. The questions are, who is going to adopt such technology, how is it going to be adopted, and when and where will its adoption be most abundant?

A likely progression can be postulated. First, we must begin to change over from magnetic stripe reader devices to smart card enabled devices. The public must get used to smart cards. To gain this familiarization, the smart card must first come to the PC. Next, it must be able to hold money—then we can address the card holding multiple personalities. Taking a view of things from a pragmatic level, a smart card with an individual's picture will provide near-term physical security. Proximity/Smart card reader technology will be readily available. Smart card readers will be in all PCs/laptops within two to three years. These cards will be PIN based. Cascading (single sign-on) public key technology will be supported via smart cards containing private keys.

There will be a need in the future to have more than one key in smart cards. RSA has suggested that three key pairs will be sufficient: one for identification and authentication, one for non-repudiation (signatures), and one for communications. Such cards may need the ability to hold old keys for some period of time. We will more likely have multiple cards and multiple algorithms for the near to mid-term. This is where the independence of the cryptography technology provider can bring value to the marketplace. Multiple cards will create a problem for the

⁹ An example of such tag-alongs is the small electronic keying device for opening car doors.

customers (too many cards). Enterprise customers entering into the smart card business are wondering how to support multiple cards.

*Cryptography
in Smart Cards*

One way to achieve user-transparent, ease-of-use of cryptography in systems is the use of smart cards. Smart card technology will make it easier to attain individual entitlements through simplified individual authentication. Smart cards provide an extendable capability that could include smart purses and basic identification functions as well as specific access IDs. Several technology providers are planning to interface to these types of products; some already have APIs established (e.g., Microsoft SmartCard API, Sun JavaCard API).

Several areas for continued IA research in smart cards include optics, biometrics, tamperproof devices, and a secure operating system.

Optics

- Optics. Integrating optical components for information assurance is just beginning. Examples of such technology already exist, such as *Active Card*, developed by the French, which provides an optically read card with multiple (four) personalities. It has a duress mode (which is export controlled by the French).

Biometrics

- Biometrics. Integration into smart cards remains an unsolved research issue. Biometrics are still not marketable—too many false assumptions built into the technology (e.g., the 1994 Olympics usage where the hand reader did not account for people with missing digits).

*Tamperproof
Devices*

- Tamperproof devices. There is an increased need for tamperproof or tamper-responsive devices. At least one provider is now working on tamper-resistant hardware.

*Secure
Operating
System*

- Secure operating system. Several providers believed that developing a secure operating system in a smart card is a good idea for research.

Networking

There is a technology trend moving toward massive networks of heterogeneous systems in the next three years. Seven years out is impossible to predict in the dynamic world of information and network

technology. Networking is becoming pervasive within companies today and will become pervasive in people's homes in the near future. Network security and systems management continue to be the significant problems demanding additional attention. Network security protocols will be the key enabler of our future networks. Even with the rush to employ new technologies, there is a sense that the average person does not trust the network to be "secure."

The key networking paradigm shift described by one networking technology provider was the movement toward the "three-dimensional network." Traditionally, service providers have been concerned with two critical dimensions of access to telecommunications services—speed and distance. However, there is a third dimension of network access that must be considered in the future—policy.

The policy dimension provides a component to address "Quality of Service" and security issues in addition to distance and speed. Policy "attributes" can also set the stage to have varying degrees of services offered by the network (i.e., class of service). For example, a policy may define a very "basic" type of service which would be the moral equivalent of flying "coach" on an airline. Changing the policy may generate a different (possibly higher) level of service (better quality, more secure), analogous to upgrading to fly "first class." Requesting and receiving a higher level of service would likely be more costly. Companies (and individual users) may find that traditional services such as electronic mail run very nicely on "standby" type networks (the lower grade, "coach" type service) which guarantees two-second response time. Higher quality service or additional security services may be needed for such things as video teleconferencing, thus requiring a different policy that would map to a higher quality service.

*Active
Network
Component*

Part of the architecture needed to support the 3-D model and 3-D networking described above (i.e., policy, speed, distance) will cause networks to change. Key technologies will be needed for this new 3-D world. The desktop computer will become an active component of the network. The desktop computer will request service (possibly pre-defined by a policy) and the network will react to provide the type of

service called for in the policy. The desktop computer must have the appropriate technologies in place to be able to authenticate with the network before asking for service. There could be a host of prearranged policies that have been designed by the company for its individual and corporate needs. The network is aware of these policies and activates upon authentication.

*Programming
the Network
Fabric*

Accomplishing this paradigm shift requires that we do more in scheduling (programming the use of) the networking fabric. We need to rethink networking between entities. Scalability is important. We need to re-center nomadics.¹⁰ We need to work more on secure, robust networking protocols. Redirection is required to get around our current add-on solutions. The Internet Protocol (IP) is a 30-year-old extended-life technology. IP security (IPv6) is currently being deployed to provide much needed security capabilities. Although becoming increasingly pervasive, IPv6 is already perceived as outmoded and needs to be replaced. It is time to re-invent the second level of protocols (more along plastic flow vs. packets in a ping-pong game). This includes a “smart” fabric that uses a new set of protocols. Additional research investments are needed in self-routing packets, intelligent switches, and improving on our multicast capabilities.

*Protocol
Policies*

There is almost no work in this area and probably very few ideas on how to do this research. Research is needed in setting up soft layer-three circuits (RSVP) or establishing layer-one circuits where the connection is more secure. Latency and jitter problems are better handled at layer one and it is also more secure. Certain types of transactions need to be handled differently. Network computers will require higher bandwidths. We need policy-oriented routing (different than RSVP or RTP protocols of today). Smart networks must deal dynamically with policy and network context. We must be able to “juggle” net responsibilities. We must have dynamic vs. static policies. Bandwidth is not a constant! We

¹⁰ Users will become increasingly mobile (hence the term “nomadic”), having computational power wherever they are at the time they need it. Access may be via their own portable computing devices and/or remote (public or private) fixed computers. Their secure access to such computing resources anytime from anyplace will be critical.

have bursty¹¹ nets that require resource management. Example “net-busy” management options might be to “e-mail to disk” (i.e., save the file as a stored mail queue and autonomically retrieve and send at earliest “net-free” opportunity).

*Service
Associations*

QoS research still needs more emphasis. CoS will add another level of complexity. We will see an increase in the mix and match of networking media (land-line and radio). We will have Active Content¹² at high speeds. Maintaining QoS through security mechanisms will be required. We will be forced into misuse detection via logging of activity. We will not be using virus detectors; instead, we will have signed code, valid guarantees (bonding certificates, knowledge of sources) attached to code (promises), and we will have “proof-carrying” code.

*Robust Error
Recovery*

We should have a policy that “time-outs are considered harmful.” This requires dealing with error in an absolute manner. One networking technology provider pointed to a new proprietary protocol as an important step towards moving beyond IP. This protocol incorporates very robust error recovery. This new protocol also serves as the base layer of this particular company’s network architecture. It was described as working similar to the electrical protocol in a backplane. The company has demonstrated protocol exchanges at 10 nano-seconds in an Application Specific Integrated Circuit (ASIC) hardware implementation. This protocol potentially will be introduced as an Internet Engineering Task Force (IETF) Request for Comments (RFC).

*Establishing
Trust Across
Networks*

Changes in communications will provide new opportunities but at the same time could potentially introduce new vulnerabilities. Larger-scale networking is such a change. A major research thrust in the future should be on finding ways to extend the enterprise—that is, using the network to securely extend the traditional boundaries of the enterprise. How do you effectively establish levels of trust between people who are

¹¹ This phenomenon is the result of distributed control—sometimes portions of the network will be relatively quiescent and, at other times, they will be overloaded. These periods may be relatively short and random.

¹² “Active Content” refers to the information contained in a network packet that is used for making dynamic routing decisions (i.e., it is content information used as parameters in enforcing network policies).

engaging the enterprise from a variety of geographically distributed locations? Examples of such locations include nomadic individual to enterprise, home to enterprise, within an enterprise, and enterprise to enterprise.

*Nomadic
Computing*

Mobile systems are now viewed as necessary for business. In time, nomadic computing will make it hard to distinguish between mobile and fixed sites. Business will more naturally distribute assets and/or resources to provide more efficient services. Secure communications is key to this distribution. The ultimate objective is seamless, secure communications within the enterprise corporate structure and across the Internet. Mobile communications significantly magnifies the security problems that must be dealt with in this case. The application of mixed communications media (e.g., radio, land-line) requires protection at the interface. The intended applications for mobile devices will require secure transaction capability. Communications providers must now ensure that security is incorporated into the equipment being provided. These are not necessarily universal capabilities but rather suites of capability that can be tailored to meet different types of transactions. Research is necessary to provide these suites of capability. Transaction IO Schemes will also be required.

*Virtual
Private
Networks*

The concept of virtual private networks (re-introduced in 1994) was originally conceived in 1986 by DoD in its Blacker Project and by the CIA in a similar project. The ideas of VPNs should continue to be pursued vigorously.

*Non-disclosed
Protocols*

There is still the issue of non-disclosed protocols that will need to be secured. Cable TV and some phone company protocols are examples of protocols that are not published (except through patents or via the underground). This must change if we are to use these various media securely and should be researched further.

*Optical
Networking*

Optical networking provides another important research area for information assurance. There is still much to research.¹³ In addition to technology advances, optical nets present the issue of information aggregation vs. disaggregation. One provider thought that end-to-end encryption does not appear to require the speed of optical switches—

encryption prior to aggregation will provide most of the protection in high-bandwidth pipes.

*Advanced
Firewall
Technology*

One potential R&D area in which some industry technology providers are interested is advanced firewall technologies in environments of very high bandwidth and very high transaction rates.

Applications

Databases

There is almost no work being done in secure databases. The potential government and commercial funders of such research perceive that the problems are solved; thus, funding for such research has been largely eliminated. However, the area of multi-level database security still needs more work. The advent of increasing parallelism in computing cannot be dealt with in today's MLS approaches. For example, we cannot do MLS distributed locking. This limits large-scale applications that can use parallel processing. Today, most users are ignoring security and using the parallelism provided in products in a system-high security context. Database front-ends are changing. Where should security mechanisms go as DBMS front-ends become more powerful? The answer: currently unknown. The issues of intelligent databases and data mining present additional IA issues that need research.

*Distributed
Directory
Services*

At least one provider is heavily engaged in database research to obtain efficient extraction of data resident in many different locations. This effort is aimed at both general databases as well as those tuned for network management. With content highly distributed, directory services will become increasingly important. Improved protected database technology will be needed. One driver is the government mandate (by CY98) to provide "local number portability" (e.g., 1-800 numbers have a directory mapping to a local number). In the short term, this effort will be accomplished out of band, but it will have access control in-band in the longer term. Operating systems will continue to share a major role in the security of such databases.

¹³ See *Workshop Report: The Role of Optical Systems and Devices in Security & Anticounterfeiting*, edited by Bahram Javidi. The workshop was sponsored by NSF, DARPA, and AFOSR, and held at the Institute for Defense Analyses on February 26-28, 1996.

Secure Operating Systems

Some providers believe other operating systems do not matter much—Microsoft’s NT is too dominant. They believe computer users have “lost the desktop” in the context of security (i.e., Microsoft’s Win95 and NT predominate and have not set the security bar high enough). Most believe that we still need more research on secure operating systems. There is a strong belief that we should widely explore new ideas in operating system research—something useful will show up later. We need to get out of the near-term research view. That view is clearly being pursued in the engineering improvements that are being made as part of product development.

More security support is needed in operating systems. One of the biggest issues is making secure operating systems work well with applications. Research requirements include:

*Fine-Grain
Object Support*

- Fine-grain object support. Operating systems will need to support much finer-grain control (e.g., the fine-grained objects that are dealt with by applications and application utilities such as databases).

*Better Server
Engines*

- Better server engines. This would include providing isolation of services, more constrained security policies (i.e., physical isolation for non-discretionary access controls in servers and support of user-oriented discretionary access controls in clients), and a better resource manager (e.g., an operating system model for resource reservations to support video and audio channel mix).

*Distributed
Authentication
Support*

- Distributed authentication support. This would include additional features such as Remote Procedure Call (RPC) authentication with impersonation vs. delegation.

*Domain
Support*

- Domain support. This would include multiple untrusted applications with some mutual trust relationship.

Label Support

- Label support. Within the operating system, we need the capability to handle sensitivity marking (e.g., security labels, tagged architectures, crypto-seals). The crypto-seals can provide releasable qualification (i.e., a file is not releasable without it).

*“Trusted”
Operating
Systems*

Some believe that the government should not be investing in “trusted” operating systems. The entire area is very fluid at this time. The Java programming language and programming environments are raising many important issues with regard to local vs. non-local control.

DoD’s MLS “trusted assurance” paradigm “busts the economics” needed to get products into a highly competitive market in a timely fashion (i.e., the evaluation process is too long and costly with little to no marginal returns). Many providers still believe there will be a broader market for MLS. This will happen as we move back from the perimeters to the individual elements (e.g., servers). However, this will not happen very fast. MLS-based operating system security is needed—but with more flexibility in providing it.

*Directory
Systems*

Operating systems cannot handle trust issues alone. Robust directory systems are needed along with strong authentication. Corporations must be able to dynamically configure their systems to be able to collaborate with other companies on mutual projects in a secure manner. (Project-oriented partnerships, alliances, and subcontracting are examples.) When collaboration is complete, the systems must be reconfigured so the sharing is no longer possible. This dynamic association will be made possible with the advent of smart switches; programmable, low-level highly configurable mechanisms; and virtual private networks with embedded crypto of which the user is entirely unaware. Companies and individuals could, through appropriate policy definition, buy any needed level of service (including security). Smart cards may very well carry configuration policy and time-based security information that would be communicated to the network.

Applied Cryptography

We need faster, more robust, higher assurance, well-understood (openly published) cryptography in commercial systems. Electronic commerce will be relying on public key cryptography, not layer-one security. Information assurance in a globally networked environment depends on the critical technology of encryption. With respect to the global economy, any U.S. government-funded R&D investment in this technology will be considered U.S. centric unless industry participates

as a full partner. Industry involvement will promote rapid technology transfer. The current Administration's position on export control and cryptography remains the most significant barrier to securing distributed systems on a global scale. Removing this barrier will result in a higher priority within technology providers for enhancing their product line with appropriate cryptographic functionality.

*Algorithms &
Protocols*

The technical issues associated with key size, key signatures, and non-repudiation are fairly well addressed. We need continued research sponsorship for the successor to the Digital Encryption Standard (DES). We need more work in computational number theory. Efficiency of key generation and algorithms are important. One of the biggest gaps in U.S. security research today is in the development of more efficient cryptographic algorithms. We need more cryptographic protocol research—new time-based and token/crypto integration is needed. In general, although the communication protocols may change, cryptographic-based security related to them may stay relatively constant.

*Usability &
Trust*

One major issue is usability. How does a user have a sense of trust that cryptographic components perform their operations correctly? There is a gap between technology and trust. There will be a cultural change needed to close this gap. Individuals need to become familiar with key management. The problem with key-sharing is that an individual can easily forget what the keys are for and when to use them.

Two important areas needing additional exploration are Certification Authorities and key recovery agents. Developing useful schemes that separate these two functions should be pursued.

*Scalable
Certificate
Authorities*

While there is probably enough government- and commercial-sponsored research going on in Certificate Authorities, research is still lacking in addressing the scalability issues of Certificate Authorities. Research is needed to examine the issues of availability and turn-around time. Public key cryptography provides strong authentication. Work needs to be done on the issue of multiple Certificate Authority hierarchies. We will have Certificate Authorities within an enterprise and in trust channels above the enterprise. Such multiple hierarchies of trust need to be efficient and

easy to use (i.e., the user shouldn't need to be highly knowledgeable in cryptography). This is one of the most critical research topics—examining how we deal with distributed, global trust issues. What are the required scales and layers of trust and how do we efficiently and effectively provide them?

Government deployment of Certificate Authorities will be needed to address part of the scale issue. The government has a legitimate role in certifying Certification Authorities. A potential research approach is the partnering of industry with government as the Next Generation Internet (NGI) evolves. This could allow the placement of industry products into an experimental effort that would result in widespread dissemination of knowledge. To make this work, we must change the adversarial relationship between government and industry. We need appropriate mechanisms to make this happen.

Key Recovery Many issues essential to key recovery systems are not solved and will require more research. Several technology providers have spent and are continuing to spend a lot of money on this area. Many providers believe the approach to key escrow/key recovery must be changed. If the government wants the benefit of cryptographic protection in our critical infrastructures—while protecting the use of cryptography from adversaries—it must be less “big brother” about it. The government should look for new approaches to key recovery that will improve recovery from inept employees, including system administrators. Several technology providers believe that current key recovery schemes must be modified to take a networking viewpoint. The networking viewpoint changes the nature of this problem—it is not satisfactory to have many unique approaches. We need new work on how to satisfy national priorities through network-oriented common mechanism technology. Several providers indicated that the Key Recovery Alliance¹⁴ is not moving in this direction.

¹⁴ The Key Recovery Alliance is made up of about 62 companies. It was formed in 1996 to develop exportable, strong encryption having *key recovery solutions* that could both meet business requirements and ease the restrictions on import/export controls world-wide.

Encryption Chips Several technology providers depend on cryptographic technologies for their success, specifically encryption chips. A key research area in chip development is focusing on using special types of currents to do fast exponentiation. The new elliptic curve algorithms rely on fast multiplies vs. exponentiation. We need ciphers that behave. Cryptography is moving into the processor, including fast, partial encryption. A major issue is the reliability and trust of random number generators—we don't know when these mechanisms fail. We need quality tests (on board) that provide assurance of correct operations. We must also decide how to deal in large collaborations with random number generators.

New uses of applied cryptography include autoimmunity and crypto-seals:

Autoimmunity Research is needed to use cryptography in new and interesting ways to achieve security. For example, autoimmune research by one industry technology provider includes the concept of self-encrypting software to provide anti-virus capability.

Crypto-seals Cryptography can provide the capability to communicate with protection in both the “private to private” context and the “public to private” contexts. We need “protection at a distance” (i.e., the ability to execute remotely and securely). One provider offered the idea of crypto-sealing, which he has incorporated into his products, as part of the solution set for such protection. Crypto-sealing allows one to match the seals prior to opening the encapsulated packet and executing it. It allows for the building of enclave measures that will only allow appropriately crypto-sealed entry. This can work for mobile code, for dynamic access at network layer, and for the network transport layer. System entry history at guard/firewall can be maintained with crypto-sealing. This is done on a file basis and allows change detection as well as precluding entry through the firewall. Research along these lines should be continued.

Long-Term Key Management Issues of long-term archiving and retrieval of encrypted documents should be examined (see the integrity discussion on page A-8).

Hardware-Based Security

Hardware has some key IA research issues. Hardware has long been attractive in terms of providing basic support for information assurance. Many believe that we will achieve significant economic and protection gains by migrating “trust” into hardware. Cryptographic support, integrity support, and isolation support can all be provided through hardware. Insertable, specialized protection-hardware can economically support important IA needs of our systems. Tamper-responsive hardware adds considerable assurance to our hardware base and to overall system assurance. Some key capabilities are coming out of hardware research, especially cryptographic capabilities supported in base hardware at the microprocessor and motherboard levels of integration. These capabilities include the ability to have signed executables that can be verified at load and run-times; ability to have a secure boot process through a digitally signed Basic Input/Output System (BIOS); and the ability to build a system domain from the boot loader and operating system loader (all signed).

Component Design

Hardware-based security component design is an issue. Software technology providers are looking for improved availability and integrity in hardware. Such providers may perform one to two reference ports per product. A key to these providers’ software quality is that they must generate code for all compilers (e.g., for all X86 versions). Efficient cryptographic algorithms are needed. Component source authentication may be needed. Efficient hardware support for complete (isolated) virtual machines will be important.

Virtual Machines

The virtual machine concepts that evolved out of the MIT work on MULTICS should be revisited as a basis for part of the system-level security architecture. More “complete” virtual machines are needed. The multi-state machine that MULTICS provided was from 8 to 16 states using the ring mechanism. Gemini computers support 10 states. The X286 supported only four states and these are still in use today. The Intel x432 Capability Machine and the DG8000 supported more than eight states. One technology provider thought that four states could be sufficient (mandatory controls, discretionary controls, applications

security, and public (no security)), but added that these may need refining.

*Power vs.
Speed*

Power is an issue! It is getting harder to distinguish 0's and 1's. This is the issue of small voltage and higher clock speeds. To gain the higher speeds at low voltages, new design approaches to circuit integration and power efficiency, including new substrates, are going to be required. Clearly, basic research in physics and materials will become increasingly important in advancing the integrity of hardware components and should be included as part of that component's research. Alternative substrates may be part of the integrity solution in just a few years. Alternative "systems on chip" designs may also be part of these solutions.

*Design
Verification*

Hardware design verification is a key issue. Lead time is important. Some designers are currently specifying chips as far out as the year 2003. The driving factor of this issue is *design verification*. Currently, there are more nodes than a verification condition generator can run. Thus, there is a turn by hardware designers to model checking which provides approximate analysis and supports compatibility assessments, failure analyses, and sample and design testing.

Manufacturing

Hardware manufacturing is becoming an issue. By the year 2000, one provider will be pushing eight processors per second off the fabrication line. How will random number generator get tested? How will the public/private key pairs be burned-in (off-line)? How will key generation be tested?

Manufacturers must turn more to design validation. Their chips for the year 2000 are already designed. All the time between now and high-volume manufacture is devoted to test and analysis. Some simulation is done—this works best where the state of the machine cannot be predicted, and visualization of internal machine interactions can allow observation of unpredictable phenomena.

Appendix B.

IA Research Investment Estimate

This appendix addresses quantification issues of commercial IA investment and provides an estimate calculated by the authors as representative of what such a commercial investment might be.

Commercial IA Research Investment Not Uniquely Identified

Ideally, we would have liked to have acquired quantified funding for categories within IA research. However, IA research funding, as a more general category, is itself not uniquely identified or captured. Much of IA research is directly incorporated as part of product development. In the case of a niche information security company, assuming all its R&D is devoted to information assurance (primarily advanced product development and contract research), all the funding for such R&D efforts potentially can be captured from public financial information¹ or from government contract information—but both must be captured to ascertain the commercial funding portion. For most telecommunications or computing technology providers, this information assurance R&D funding can't be broken out. Such research is generally a part of direct product improvement and is not separately identified from other product development. Even in separately identified research budgets, the large companies do not make a unique distinction of IA-related research.

¹ Four of these niche companies (Security Dynamics, Secure Computing, Raptor, and Trusted Information Systems) are rapidly growing, publicly owned enterprises having combined total revenues in 1996 of \$157 million dollars. Based on their individually stated R&D as a percentage of annual revenues, they expended approximately \$28.1 million (18% on average) for R&D. It is undetermined how much of that expenditure was for advanced product engineering and how much was government-supported contract research. Data source: Collen Frye, "Software 500," *Software Magazine*, July 1997, pp. 41-67.

*Proprietary
Information*

It is not possible to make distinctions where a particular technology research funding emphasis is being placed over a particular timeframe without the explicit cooperation of the company. Research categorized by area of investigation and by project investment timeframes is not publicly available. Categorized, as opposed to overall, research funding is generally proprietary. Technology providers of any size, who consider their information assurance R&D funding as proprietary, will not make it available.

*Contract
Research*

Some of the IA research funding was identified as coming from U.S. government contracts. A number of technology providers, primarily niche information security companies, perform a portion of their research as part of government contracts (consulting services and, specifically, contracted research). These contracts may be from the DARPA, NSA, or other government agencies. Parts of the research results have been used in the company's product offering. While the data for this funding can be captured, it should be captured as government-funded IA research and not as commercial IA research.

*Anecdotal
data*

Several companies did volunteer their staff-year data that was devoted to some part or all of their R&D efforts in information assurance. Others hinted at funding levels. The data was anecdotal, providing inconsistent parameters across the set of respondents. Its value was, therefore, limited. Among the numbers given for staffing and funding:

Staffing. Niche technology providers:

- One niche company reported 440 engineering staff-years over 15 years (more than 29 staff-years per year) are assigned to develop the basic technology for its product family.
- Another reported about 4% of more than 40 engineering staff members are assigned to advanced IA research.
- A third niche company reported about 5% to 10% of the engineering effort is spent on contract and internally funded research, of which 90% of that is currently government contract research.²

Staffing. Large companies:

- A very large provider with an engineering staff of several thousand people indicated it has about 40 people assigned to IA research in two security laboratories: 28 people within a U.S. lab and 12 in a lab off shore. Both laboratories are supported by corporately funded external research, particularly university research and research off shore.
- Another large provider said it had about 3% of its engineering staff working on security but they weren't fully dedicated. Some portion of the remaining staff also contributed at times to some aspect of security, but only part time. Another large provider related a similar situation with about 1.5% to 2% working primarily on security.
- A fourth technology provider said that it had increased its overall security efforts by establishing three new groups. One group has 10 to 20 people located in two laboratories, one being off shore. Their work is devoted to security analysis and internal technology "security consulting" activities.
- This same company also remarked it knew of only five companies (large and small) with a grand total of 30 people working on anti-virus research.

Funding. Some technology providers indicated efforts by funding levels not only of their own companies but of outside examples:

- The example of the Singapore government starting its own \$80 million research initiative³ because it couldn't import strong U.S. encryption technology.
- Though hesitant to invest much effort in security, one company believed that there would be a \$2 billion industry potential once export controls are lifted.

² Estimating that this company spent about \$10 million in R&D (based on a percentage of annual revenues), this would mean that the company expended between \$500,000 and \$1 million on contract research and internally funded research. Of that amount, only \$50,000 to \$100,000 was company-funded information assurance R&D (1/2% to 1% of total R&D).

³ This funding is assumed to be over a five-year period.

- One niche technology provider has about 38% of its annual revenues, which are between \$25 million and \$30 million, coming from government contract funding for R&D and consulting.
- Another niche provider said that about 6% to 8% of its annual revenues, which are between \$3 million and \$5 million, is invested in research, mostly for advanced product development.
- One large technology provider stated that it would be spending over \$2 billion in 1997 for R&D, with an unspecified amount that included efforts on availability, integrity, and assurance (reliability).
- One large technology provider, with overall R&D expenditures in excess of \$500 million, sponsored 122 collaborative research projects in 1996 at various universities, representing \$2 million in research grants and another \$2 million in equipment donations. The amount that might be related to information assurance could not be specified, though it was thought to be modest.

Calculating a Gross Estimate of Commercial Information Assurance R&D Investment

The authors of this report have calculated, using publicly available data, that the gross commercial information assurance R&D funding ranges between \$120 million to \$355 million per year.

Estimate Approach

An estimate of the potential IA research funding can be attempted using aggregate industry data, individual corporation financial reports (10K's), and factors generated from the anecdotal data provided in the interviews. The anecdotal information indicates that there may be somewhere between 1% and 3% (maximum) of overall R&D funding devoted to IA research. This is a gross estimate and must be recognized as such. There is currently no generally established basis for determining the validity of this estimate, so it should be used with caution in future government funding decisions. This estimate is provided simply to attempt a grasp at the order of magnitude of potential commercial information assurance R&D investment per year that might be available or have been recently experienced.

Two Bases for Estimating

Two bases are available to approach the development of and estimated range of investment.

- Basis 1: Estimate based on total market.
- Basis 2. Estimate based on overall U.S. industry R&D data.

Basis 1:
Total
market

The telecommunications and computing market is significant and growing in size. The world-wide market was estimated at \$610 billion in 1996 and the U.S. portion of this market was estimated to be \$425 billion.⁴ The market growth is forecasted at 12% annually and is expected to reach \$1 trillion by 2001. This market is now composed of hardware (52.4%), software (17.3%), and services (30.3%). Thus, some 69.7% of the market (\$425 billion world-wide; \$171 billion U.S.) is hardware and software revenues that can be related to R&D investments. If, on average, 10% of this market was devoted to R&D, we would have about \$17 billion in U.S. hardware and software R&D. If between 1% and 3% of this total R&D had been devoted to information assurance R&D, we could have had a range of \$170 million to \$510 million in information assurance R&D expenditures by the U.S. private sector in 1996.

Basis 2:
Overall U.S.
industry R&D

A second basis for estimating information assurance R&D funding levels can be through the use of overall industry R&D information. As previously indicated in the discussion of Subfinding 4.6 (page 17), the total share of U.S. private-sector R&D carried out by companies in the electronics and information fields jumped from 32% in 1981 to 41.9% in 1988, and then moved up to 43.6% in 1995.

Alternative Office of Technology Policy scenarios for 1995 U.S. industry R&D spending show a range of between \$57 billion (with 1981 R&D intensity and portfolio) and \$107 billion (current 1995 OTP estimate).⁵ Using the current OTP estimate, the U.S. 1995 R&D spending in electronics (43.6%) would have been about \$47 billion. If approximately 66% of this figure is related to telecommunications,

⁴ International Data Corporation, *Black Book*, 1996.

⁵ "Telecom Paces U.S. Electronics Sales Gains," *New Technology Week*, 12 May 1997, p. 9.

computers, and electronic components,⁶ then the R&D share we are concerned with would be about \$31 billion. If the 1% to 3% range estimate for commercial information assurance R&D is applied to this figure, we could have had between \$310 million and \$930 million in U.S. commercial information assurance R&D spending in 1995.

Gross Estimate

The two bases yield a wide range: a minimum of between \$170 million and \$310 million) and a maximum of between \$510 million and \$910 million. Assuming that the average of these ranges will get us closer to the truth, we arrive at an average range of \$240 million to \$710 million.

Using the average range and further assuming that we may have overestimated by as much as 50% (i.e., half of the companies have no investments in information assurance R&D), we get a more conservative lower bound of \$120 million and a more conservative upper bound of \$355 million as a possible range for our estimate of commercial information assurance R&D investment per year. Thus, we can very cautiously assert⁷ that possibly somewhere between \$100 million and \$350 million is a gross, somewhat conservative, estimate of commercial funding for information assurance R&D that is currently being expended per year or that has been recently expended.⁸

It can be seen that the estimated information assurance R&D funding of this range estimate comes close to using the same approach and publicly available data found in the financial statements of the sampled major technology providers (see Table B-1 on page B-7).

⁶ Ibid. Electronics sector consists of electronic components, consumer electronics, telecommunications, defense communications, computers and peripherals, electromedical equipment, industrial electronics, and other related items. First quarter 1997 sales data (\$107 billion), compiled by the Electronic Industries Association's Marketing Services Department for the U.S. Department of Commerce, included electronic components (31%), telecommunications (14%), and computers and peripherals (21%). This growth was 8% greater when compared to the same quarter in 1996, with telecommunications growing 13%, computers and peripherals growing 7%, and electronic components growing 5%.

⁷ While we know of no other such estimate, and the estimate provided herein may seem plausible and may be the best available, we urge caution in its application.

⁸ It is further conjectured, based on anecdotal and publicly available data, that the niche companies combined *internal* funding accounts for no more than 1% of this amount (i.e., \$1 million to \$3.5 million).

Applying the anecdotal percentages (1% to 3%) to the providers' overall R&D total (\$18.6 billion), a potential range of commercial IA research funding for these major technology providers is between \$186 million and \$558 million. Adopting our previous conservative approach, half of the above amounts would yield a range of \$93 million to \$229 million as the likely amount of information assurance R&D investment being expended by these companies.

Table B-1. Revenues, Overall R&D Expenditures, and Estimated IA Research Funding^a

Company	No. of Employees (Total/R&D-Eng.)	Annual Revenues (\$M)	Overall R&D Expenditure (\$M)	% of Revenues	Estimated IA Research (1% -3% of R&D) (\$M)
Oracle	23,113/3,125	4,200	389	9	3.9 - 11.7
Sybase	5,484/?	1,000	164	16	1.6 - 4.9
IBM	240,615/?	76,000	4,654	6	46.5 - 139.6
HP	112,000/?	38,400	2,718	7	27.2 - 81.6
Sun	17,400/?	7,100	657	9	6.6 - 19.7
Microsoft	20,561/6,861	8,700	1,432	16	14.3 - 43.0
Novell	5,818/1,806	1,400	276	20	2.8 - 8.3
3COM	5,190/1,227	2,300	233	10	2.3 - 7.0
CISCO	8,782/2,420	4,100	399	10	4.0 - 12.0
Lucent	124,000/?	23,300	2,703 ^b	11	27.0 - 81.1
Intel	48,500/9,100	20,800	1,808	9	18.0 - 54.2
Motorola	139,000/15,800	28,000	3,152	11	32.0 - 94.6
TOTAL	N/A	215,300	18,585	N/A	186.2 - 557.7

a. Source: Company 10K's and IDA estimates.

b. The Lucent number here (and which appears as a different number in Table 2) was taken from two different sections in Lucent's 10K. We were unable to reconcile them, given the information provided in the 10K's.

Reference List

- “Business World Global 1000,” and “Top 100 Global Companies,” *Business Week*, July 7, 1997, pp. 52-97.
- Department of Defense, *Department of Defense Trusted Computer System Evaluation Criteria* (TCSEC), DoD 5200.28-STD, 1985.
- Fox, Christopher and William Frakes, “The Quality Approach: Is It Delivering?,” *Communications of the ACM*, Vol. 40, No. 6, June 1997, pp. 24-29.
- Frye, Collen, “Software 500,” *Software Magazine*, July 1997, pp. 41-67.
- International Data Corporation, *Black Book*, 1996.
- Jacobson, Ken, “Industry R&D Spending Patterns Shifting,” *New Technology Week*, May 12, 1997, pp. 6-7.
- National Research Council, *Computers at Risk: Safe Computing in the Information Age*, National Academy Press, 1991.
- National Research Council, *Cryptography’s Role in Securing the Information Society*, National Academy Press, 1996.
- National Research Council, *For The Record: Protecting Electronic Health Records*, National Academy Press, 1997.
- National Research Council, *Realizing the Information Future: The Internet and Beyond*, National Academy Press, 1994.
- Office of Technology Policy, *Globalizing Industrial Research and Development*, 1997. Available through the National Technical Information Service, PB96-119201NB.
- “Telecom Paces U.S. Electronics Sales Gains,” *New Technology Week*, 12 May 1997, p. 9.
- Workshop Report: The Role of Optical Systems and Devices in Security & Anticounterfeiting*, edited by Bahram Javidi. Sponsored by National Science Foundation, Defense Advanced Research Projects Agency, and Air Force Office of Scientific Research. Held at the Institute for Defense Analyses, Alexandria, VA, February 26-28, 1996.

Acronym List

2D	two dimensional
3D	three dimensional
API	applications programming interface
ASIC	Application Specific Integrated Circuit
ATM	Asynchronous Transfer Mode
BIOS	Basic Input/Output System
CASE	computer-assisted software engineering
CDSA	Common Data Security Architecture
CIO	Chief Information Officer
CORBA	Common Object Request Broker Architecture
CoS	Class of Service
COTS	commercial off-the-shelf
CRADA	Cooperative Research and Development Agreement
CY	calendar year
DARPA	Defense Advanced Research Projects Agency
DBMS	database management system
DCE	Distributed Computing Environment
DCOM	Distributed Common Object Model
DGSA	Department of Defense (DoD) Goal Security Architecture
DoD	Department of Defense
DVD	Digital Versatile Disks
http	HyperText Transfer Protocol
IA	Information Assurance
IDA	Institute for Defense Analyses
IDS	Intrusion Detection Systems
IETF	Internet Engineering Task Force
INFOSEC	Information Security
I/O	input/output
IP	Internet Protocol
IPv6	Internetworking Protocol (newest version)
IRC	Information Security (INFOSEC) Research Council
LAN	local area network

MEI	minimal essential infrastructure
MIT	Massachusetts Institute of Technology
MLS	multilevel security
NGI	Next Generation Internet
NIARA	National Information Assurance Research Agenda
NLM	Net Layer Module
NSA	National Security Agency
ORB	object request broker
OTP	Office of Technology Policy (U.S. Department of Commerce)
PC	personal computer
PCCIP	President's Commission on Critical Infrastructure Protection
PGP	Pretty Good Privacy
PIN	personal identification number
PKI	Public Key Infrastructure
PVS	Prototype Verification System
QoS	Quality of Service
R&D	Research and Development
RFC	Request for Comments
RPC	remote procedure call
RSVP	ReSerVation Protocol
RTP	Real-Time Transport Protocol
S/MIME	Secure Multipurpose Internet Mail Extensions
SET	Secure Electronic Transaction
SMIB	Security Management Information Bases
SPKI	Simplified Public Key Infrastructure
TCB	Trusted Computing Base
TCP/IP	Transmission Control Protocol/Internet Protocol
TCSEC	Trusted Computer System Evaluation Criteria
URL	Uniform Resource Locator
U.S.	United States