Statement of Dr. William L. Scherlis
Professor, School of Computer Science
Carnegie Mellon University

Before the U.S. House of Representatives
Committee on Government Reform
Subcommittee on Technology, Information Policy, Intergovernmental Relations and the Census

Federal Information Technology Research and Development
July 7, 2004

Mr. Chairman and members of the Committee. My name is William Scherlis, and I am a Professor in the School of Computer Science at Carnegie Mellon University, where I direct a newly inaugurated PhD program in Software Engineering and lead a research project involving CMU and five other universities related to software dependability. Earlier in my career, I spent more than six years at DARPA managing research programs in areas including trustworthy systems, high performance algorithms, and software technology. I departed from that position to return to research in 1993. Our present project work is in collaboration with NASA, and has the goal of helping improve the safety and dependability for future generations of software-intensive mission systems. Software dependability is particularly important for NASA, and as you know it has broad significance for the security of our nation and its critical infrastructure.

I appreciate the opportunity to appear before you today to discuss the need for research and development for information technology. In this testimony, I address two areas of information technology (IT) that present particular challenges to federal agency CIOs and mission managers—cybersecurity and software dependability. These areas are among a number of IT challenges facing federal managers that are considered in an NSF-sponsored National Research Council study I led on IT for e-government (http://www.nap.edu/catalog/10355.html). I highlight these two areas because of their fundamental strategic and economic significance, and because of the importance of far-sighted strategic R&D to future agency systems.

Most IT leaders and computer scientists believe that IT is still evolving rapidly, and nowhere close to a plateau, either with respect to capability or quality. Mission managers will face a very different environment in the future—even in five years. We have a huge national stake in the definition of that environment.

This examination leads to three conclusions: (1) Strategic long-term federal IT R&D is more important now than before. (2) We must retain our national advantage in innovation leadership. (3) We need pro-active federal R&D leadership for critical IT challenges. This statement addresses primarily the “who,” “what,” and “why” questions posed by the Subcommittee.

1. Mission IT Challenges: Cybersecurity

Stop-gaps. Let us consider a mission manager who must provide an immediate cybersecurity solution for an organization. Following today's best practice, the manager will apply a range of available security interventions such as managed networks, firewalls, virus detection, intrusion detection systems, patch management, configuration management, and spam filters. Unfortunately, most of these interventions are stop-gaps that only partially address the weaknesses intrinsic in today's engineering practices and network and system architectures. With these interventions, the manager will be slightly less exposed in the current war of attrition (for example, between virus writers and virus detection tool creators). But statistics from the CERT and other sources suggest quite vividly that we will not win this war with the current set of interventions (http://www.cert.org and http://uscert.gov). There is a broad consensus that the stop-gaps are failing and more fundamental kind of progress is needed, despite the significant improvements in quality we have experienced over the years. The CERT security "exploit" and vulnerability curves continue to trend upwards geometrically (see chart below, from the CERT/CC).

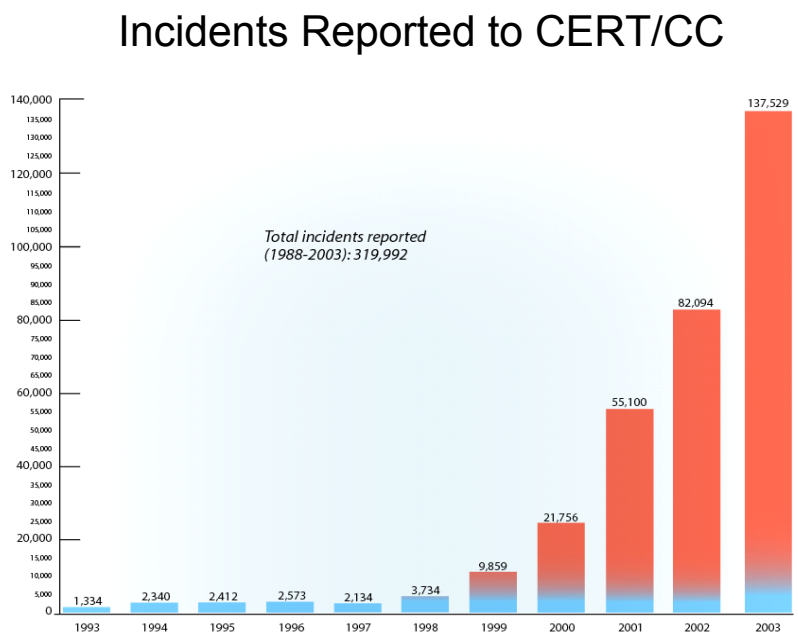

Product evaluation. Cybersecurity evaluation is a kind of product acceptance testing for security attributes of IT components and systems. The difficulty of evaluation is evident in established processes such as ISO 15408—NIST's Common Criteria (http://csrc.nist.gov/cc/). Commercial vendors may spend a year or more undertaking a system evaluation, which leads to useful and important assurances regarding specification and design. But the evaluation does not—and indeed cannot at this stage of development—assure an absence of malicious code or other vulnerabilities. Nor can it easily extend from a particular system instance to a family of configurations. These

difficulties are consequences of the fact that general-purpose technical means are not yet available to make such evaluations and, from an engineering standpoint, to make positive promises about software and IT components. Indeed, many of the fears regarding outsourcing and open source derive from the difficulty of evaluating the safety, dependability, and security of software, even when the code is fully available for inspection. The technical advances that are most promising in addressing this problem rely on deep mathematically-based techniques to analyze software code directly.

Engineering practices. We may conclude from the foregoing that we are unable to build secure systems—or evaluate those that we ourselves create—and that, more significantly, there are intrinsic characteristics of software and IT that place us in this quandary. This is not the case. Rather, it is simply the immature state of current technical understanding and engineering practice. Technical progress in the NITRD research portfolio suggests that strategic R&D effort can lead in the long run to fundamental changes in our ability to deliver higher levels of security (http://www.nitrd.gov). There are several recent reports from the Computer Science and Telecommunications Board of the National Research Council that summarize recent results and offer recommendations (http://www.cstb.org).

In the meantime, for the most highly critical systems, the present practice is to accept severe constraints on system capability and architecture in order to achieve the possibility of acceptable levels of assurance. In other words, our lack of ability to do evaluation forces us to limit the capability of the critical systems we build. There are emerging research results that demonstrate how key architectural and design commitments, coupled with analysis tools, may offer steps away from this state of limited capability. These results include, for example, self-healing architectures, framework designs for composable components, component evaluation tools, techniques for safe concurrency, etc.

Future systems. Many government agencies face cybersecurity challenges that lead the market in significant respects. (Additional examples of IT areas where government is a demand leader are offered in the e-government study cited above.) Managers who are planning next-generation systems can benefit by collaborating with the multi-agency R&D community in order to address these needs in a more strategic manner. This pattern—of investing in R&D in the supply chain in order to ensure future needs can be met—is well established in other industries, for example in supply chains for automobile and airplane design and manufacture.

In government, there is a record of success in major mission agency IT consumers applying this model, particularly DoD. The idea is to follow the established market model of supply chain management, which involves working with all levels of suppliers, helping them anticipate critical needs, and investing in R&D that has broad benefits. This market-based approach, in which major technology consumers, but particularly government mission agencies, collaborate through the supply chain to accelerate response to leading edge requirements, has stimulated much of what we take for granted in modern computing.

Many important cybersecurity improvements are already in the early stages of the R&D pipeline, including, for example, better security-focused engineering processes, evaluation and analysis tools, component designs, more secure network protocols and services, more securable system architectures, improved identity and authorization management schemes, architecture for self-healing and resilient systems, and others.

2. Mission IT Challenges: Software Dependability

A second area of IT challenge facing CIOs and mission agency managers is software dependability. Software dependability is critical for systems safety (such as for medical devices, plant control systems, and weapon systems), for management of sensitive records (such as for law enforcement, government financial and tax records, and health care records), and for critical infrastructure (such as for SCADA—supervisory control and data acquisition systems).

Software dependability is also fundamental to cybersecurity. More than 90% of the thousands of exploits reported annually to the CERT build on software flaws. Some observers have stated that these flaws exist because vendors feel they can "get away with" shoddy quality. This is not true. Rather, these flaws are the result of the fact that both the industry and the R&D community are still struggling with how to achieve high levels of quality in software engineering practice.

Process and quality. Today, software program managers generally employ traditional engineering management approaches that are taken from the statistical quality community. These involve process and metrics, with the metrics used as feedback to improve both product and process attributes. This is the essence of the ISO 9000 family, Six-Sigma, SEI's CMMI, and other process-based approaches. These are excellent, workable, and widely adopted approaches that have provided enormous benefit to projects of all sizes.

But the statistical approaches have a stop-gap character—they do not address all the challenges of producing dependable and secure software. In particular, software bugs are flaws of design and implementation—software components are not like physical systems that wear out over time. In cybersecurity, for example, when the threat model is based on frequencies of spam-zombie viruses, port scans, and spyware intrusions, the resulting system design may not be appropriate for a multi-point attack by a determined adversary. In addition, all software faults are not created equal, either with respect to the kinds of system failures they may trigger or the difficulty of repairing them. Some faults may be intrinsic in the system architecture, while others are mere coding oversights. The Common Criteria, noted above, assists in making this latter distinction with respect to security.

Software quality. A Defense Science Board task force noted that, "improvements to process, training, incentives, and procurement are critical, and yet improvements to the process without improvements to the technology cannot address the staggering complexity necessary for achieving a national competitive advantage." The President's

Information Technology Advisory Committee (PITAC), in its inaugural report, noted that "Our ability to construct the enormously complex software systems that lie at the core of our economy (is) painfully inadequate. Therefore, the increases in research on software should be given a high priority" (http://www.nitrd.gov/pitac/).

Industry leaders also recognize the seriousness of the technical challenge. Two months ago, the Business Roundtable, representing 150 U.S. CEOs, issued a report that notes, "Most of the significant cyber incidents that have harmed American business and consumers over the past several years have had at their root cause defective and readily exploitable software code." It adds, "Most software development processes used today do not incorporate effective tests, checks, or other safeguards, to detect those defects that result in product vulnerabilities" (http://www.businessroundtable.org). The Washington Post titled its report on this study, "Old Economy Fed Up With Cyber-Security."

Evaluation and assurance. As noted above, software assurances in present practice are usually based on indirect measurements—evaluation of process or organizational attributes as substitutes for direct evaluation of a product. Because existing measures are weak and overly approximate, it is difficult to build an ROI model—the "R" cannot easily be measured and so the "I" is not readily forthcoming. Improving our ability to measure the "R"—at every level from code-level fault identification to organization-level failure impacts—is part of the process of maturing the discipline. More concretely, research attention must be focused in areas such as: (1) Improved technical means to evaluate system components directly for critical security and dependability attributes, (2) Better techniques to engineer software with higher levels of security and dependability "out of the box," and (3) Principles of architecture, design, and coding that can reduce the overall impact of internal engineering faults (such as buffer overflow) on the kinds of failures that can result (such as compromise of data or loss of service).

There is some basis for optimism—including significant emerging research results (for example, relating to modeling of critical attributes, software analysis, and model checking techniques), and good indicators in industry practice that a business case is starting to emerge (an example from Microsoft is cited below).

Progress on this challenge has high leverage because it exists at every producer/consumer interface in supply chains for IT systems. Producers benefit by being able to provide concrete evidence of security, dependability, or other quality attributes. Consumers benefit when they (or third-parties acting on their behalf) can undertake effective acceptance evaluations.

3. The Essential Role of Government R&D

These challenges—cybersecurity and software dependability—have three important common characteristics. The first is that there is a compelling case, based on both need and opportunity, to develop R&D approaches that are strategic. These are fundamentally different in character from the current stop-gap practices. Second, in both areas there are government mission agencies whose requirements anticipate rather than follow market

demand. Cybersecurity and software dependability are important for everyone, but these needs are frankly more urgent for our national and infrastructural systems. Third, critical and mainstream solutions are converging rather than diverging. Many critical infrastructural systems, national defense systems, and enterprise systems are constructed from diverse components from diverse sources, and such systems may include extensive use of many familiar pervasive IT applications and components. In other words, what is pervasive is also becoming critical. And, engineering practices for critical systems are not that far beyond engineering practices for the pervasive systems. It is dangerous, therefore, to contemplate solutions where government and other critical consumers separate from the mainstream market.

Will industry do it? It is tempting to think that, while these problems—cybersecurity and software dependability—are challenging, the IT industry will eventually address them as a matter of course—that "industry will do it." In many other industries this is a legitimate conclusion, and best policy for meeting government requirements may be to follow the market. But this is generally not true for IT. And in fact it defies the historical truth of the past 50 years of IT innovation.

The leadership in IT innovation that we currently enjoy in the US is the legacy of several decades of effective and steady Federal R&D leadership, much of which is in response to the leading-edge requirements of government mission agencies. A recent National Research Council Report called "Innovation in Information Technology" notes that, regarding modern electronic commerce, "nearly every key technological component has been shaped by [federal] investment," including the Internet, web browsers, public-key cryptography, back-end database and transaction processing, and search engines. The report illustrates this by tracing the research origins of more than a dozen multi-billion dollar IT markets (http://books.nap.edu/catalog/10795.html).

The unique government role. Why is the government R&D role so essential for IT? The answer has three parts:

a. Non-appropriability. Many of the most significant research results in IT are non-appropriable. That is, there are foundational results that cannot successfully be confined to a single sponsoring organization—their impact diffuses broadly into the technical community and the market. It is difficult for firms to build an ROI model for investing in this kind of research, whose results may diffuse directly to competitors. Many of the most important concepts of information technology are in this category. The revolutionary concept of networked personal computing developed principally at Xerox PARC under a combination of federal and private sponsorship is an example. In the end, it didn't do much for Xerox, but it triggered the creation of an entire industry.

In the area of software dependability, there is recent evidence in the R&D community of progress in addressing some of the longer-term needs. Companies such as Microsoft and IBM are developing a new generation of engineering tools to improve software quality through direct analysis of software code and other artifacts. The Microsoft tools are partly responsible for the significant reduction in the frequency of "blue screens" in the past

couple of years. Of the tools I am aware of, most rely on fundamental technical concepts that were developed earlier in university labs, generally sponsored by NSF, DARPA, and other NITRD agencies. These concepts include, for example, finite-state model checking, binary decision diagrams, rule-based inference, and many program analysis techniques.

b. Commonalities. The early definition of standards is a critical element of the pattern of innovation in the IT industry. Government has had a long-standing role in facilitating the so-called pre-normative work that leads to the commonalities critical to the creation of markets for new IT capabilities. The early IETF had a pivotal role and provides an important model for consensus management in a community of innovators. For example, the first versions of the fundamental standards of the modern web-based Internet (IP, TCP, HTTP, HTML) were developed by a handful of researchers, all working under government sponsorship, collaborating through the IETF.

c. Education and universities. More obviously, the research programs in most universities and labs are closely coupled with the education enterprise. The direct engagement in the most aggressive research, including traditional exploratory basic research and far-sighted mission research, creates the next generation of inventors and innovators, who become tomorrow's industry leaders.

Industry senior managers recognize the need for R&D. Two months ago, Craig Barrett, the CEO of Intel, said to USA Today, "We have to invest more in R&D. If you have a worse education, a worse infrastructure, and you spend less of your gross domestic product on R&D, what makes you think you should be in a pre-eminent position?" Perhaps most importantly, he notes the importance to US employment of "research and development investment that is government funded," noting that the fraction of output that has gone to R&D has declined over the past two decades. "R&D creates the ideas for future products and services."

4. Moving Forward

The difficulties we are facing nationally with cybersecurity and software dependability are consequences of the limitations of our present engineering capability. Software is an unusual building material, almost unlimited in its potential for capability and scale—but we are still learning how to work with it successfully, and particularly how to create systems that are genuinely dependable and secure. This is ironic, because IT has become pervasive and, as noted recently by Alan Greenspan, is an important contributor to national productivity. But this immaturity need not persist—many in industry and universities believe that a combination of public-private partnerships and aggressive federal R&D can lead to fundamental change.

I conclude my statement with three observations:

1. Strategic long-term federal IT R&D is more important now than before.

The NITRD agencies identified in the President's Budget, along with the Department of Homeland Security, share leading-edge requirements for advanced IT capability. They recognize the necessity of stimulating innovation in order to ensure that their future mission requirements can be effectively met. By collaborating with each other they can share new technologies, spread risks, and build more effectively on the basic science portfolio principally sponsored primarily by NSF.

The NITRD process, which began with the High Performance Computing Act of 1991 and is now led by Dr. David Nelson and Dr. Peter Freeman, is an effective mechanism to support this coordination among mission agencies. NITRD has several coordinating teams focusing on specific issues and supporting strategy development. Cybersecurity is addressed by several of the teams, most notably the High Confidence Software and Systems (HCSS) coordinating group. Software dependability is also addressed by several teams, principally the Software Design and Productivity (SDP) coordinating group and the HCSS coordinating group (http://www.nitrd.gov/pubs/blue04/). Challenges related to achieving significant new levels of systems capability are addressed at several agencies, particularly at the NSF in the new Science of Design research program and at DARPA/IPTO (http://www.darpa.mil/ipto).

In the strategic planning process, it is essential to understand that we are not at a plateau in any aspect of IT capability. The capability, performance, and interconnection of IT systems are advancing at a rapid rate. Many in the industry recognize that Moore's Law will continue to hold for another decade or more, and that we are likely to experience several more decades of the kind of rapid innovation we experienced in the past ten years, which brought us the public Internet, World Wide Web, e-commerce, e-government, grid computing, and many other fundamental changes.

The critical challenges of IT are increasingly focused on quality and assurance, and this is why in this statement I have highlighted the areas of cybersecurity and software dependability. There are other areas of critical need. The NITRD agencies should be given the charge—and the resources—to address these challenges in a strategic long-term fashion.

2. Addressing these quality issues is fundamental to retaining our national advantage in IT innovation.

The NITRD innovation investment, including both mission-focused investment and the basic science component primarily at NSF, is essential to our national success in an international market that is becoming increasingly competitive with respect to IT innovation leadership. The patterns of IT innovation and the value of IT innovation leadership are now increasingly understood throughout the world—and many countries are making national-level commitments on this basis. Many believe that loss of IT innovation leadership will, for national security, have more severe consequences than the losses the U.S. has experienced, say, in steel or much of consumer electronics. (Note,

however, that this does not imply that offshore outsourcing and open source are necessarily bad for the U.S.—the economic story in these cases is nuanced, with both positives and negatives, and few general answers.) The point is that we need to maintain our strength in IT innovation. We can do this only by leading in education and in well-managed R&D.

3. Establish pro-active R&D leadership for the critical IT challenges.

Needs and opportunities have been well articulated in many studies and workshops in the past few years. It is now time to build on this progress by taking the next step, and defining some elements of a national strategy that is far-sighted in terms of impact, but actionable and concrete in terms of R&D activity. There are important new concepts and technical opportunities emerging in many labs across the country. The purpose of a national strategy is to link these efforts together to accomplish the next critical set of changes in the constantly-changing landscape of IT.

As noted in the previous section, development and execution of the strategy must be a collaborative process involving research leaders and far-sighted users in industry, academia, and government. An example of a collaboration focused on strategy development is the Accelerating Trustworthy Internetworking (ATI) initiative (http://www.ati2004.org). The most recent ATI workshop included active participation from industry, government, and academia. The purpose of the strategy is to provide a focus for government R&D in selected critical areas.

Long-term mission-motivated federal R&D is how the U.S. established its present IT leadership position, and it is what we must do both to retain this position and to address the new challenges that we now face. This will require collaboration of government, industry, and academia.

Thank you for giving me the opportunity to testify today. I would be pleased to answer any questions.