Written Testimony of

Dr. Seymour Goodman

Chair, Committee on Improving Cybersecurity Research in the United States

Computer Science and Telecommunications Board
National Research Council

Before the
Subcommittee on Terrorism, Unconventional Threats and Capabilities
Armed Services Committee
U.S. House of Representatives

April 1, 2008

Mr. Chairman, distinguished members of the Subcommittee: Thank you for the opportunity to appear before you today to discuss the subject of holistic approaches to cybersecurity enabling network centric operations.

My name is Seymour Goodman, and I am professor of international affairs and of computing, at the Sam Nunn School of International Affairs and the College of Computing at the Georgia Institute of Technology. I recently served as chair of a committee of the National Research Council on cybersecurity research in the United States; this committee produced a report entitled "*Towards a Safer and More Secure Cyberspace*." The National Research Council is the operating arm of the National Academy of Sciences, National Academy of Engineering, and the Institute of Medicine of the National Academies, chartered by Congress in 1863 to advise the government on matters of science and technology.

According to the Joint Chiefs of Staff, net-centric operations are the operational concept under which U.S. military forces and mission partners have "rapid access to relevant, accurate, and timely information, and also the ability to create and share the knowledge required to make superior decisions in an assured environment amid unprecedented quantities of operational data."[1] It goes without saying that access to such information and the ability to create and share information are capabilities that will depend heavily on modern information technology. (A number of NRC reports address matters related to net-centric operations in a naval context, including *FORCEnet Implementation Strategy* (2005), *C4ISR for Future Naval Strike Groups* (2006), and *Network-Centric Naval Forces: A Transition Strategy for Enhancing Operational Capabilities* (2000).)

[1] CONNECTING THE WARFIGHTERS, Joint Net-Centric Operations (JNO) fact sheet, J-6, available at http://www.jcs.mil/j6/c4campaignplan/JNO_fact_sheet.pdf.

But in order to leverage these capabilities effectively, commanders must be able to count on their availability when they need them, must believe that they are providing trustworthy and uncompromised information, and must know that adversaries do not have advance knowledge of ensuing military activities. Moreover, all of these things must be true in the face of an adversary wanting to compromise these capabilities. Ensuring the availability, integrity, and confidentiality of information are the classical goals of cybersecurity, and high-confidence authentication is often added to this list.

My remarks will focus on the link between cybersecurity and net-centric operations.

Given the need for net-centric operations to be conducted in a secure environment, two thrusts are necessary. The first could be characterized as "do what you already know how to do." There is much that is known about cybersecurity technologies and practices today that is simply not put into practice, and even the widespread deployment of relatively unsophisticated cybersecurity measures can make it more difficult for an adversary to conduct a cyberattack.

The second could be characterized as "learn more about how to be secure." That is, even assuming that everything known today was immediately put into practice, the resulting cybersecurity posture—though it would be stronger and more resilient than it is now—would still be inadequate against today's threat, let alone tomorrow's. Reducing this gap—a gap of knowledge—will require both traditional and unorthodox approaches to research.

Traditional research is problem-specific, and there are many cybersecurity problems for which good solutions are not known. (A good solution to a cybersecurity problem is one that is effective, is robust against a variety of attack types, is inexpensive and easy to deploy, is easy to use, and does not significantly reduce or cripple other functionality in the system of which it is made a part.) Research will be needed to address these problems.

But problem-by-problem solutions, or even problem-class by problem-class solutions, are highly unlikely to be sufficient to close the gap by themselves. Unorthodox, clean-slate approaches will also be needed to deal with what might be called a structural problem in cybersecurity research now, and these approaches will entail the development of new ideas and new points of view that revisit the basic foundations and implicit assumptions of security research.

To motivate my description of necessary cybersecurity research, consider the story of the U.S.S. Yorktown, an Aegis cruiser that was the Navy testbed for "smart ship technology" in the late 1990's. As you know, the Aegis system has been an important element of the Navy's concept for network-centric operations. A widely used commercial network operating system—Windows NT—was installed on the Yorktown to control a variety of important ship-board applications, including navigation and

propulsion. In September 1997, a crewman mistakenly entered an invalid number into a database. He thereby caused a "divide-by-zero" error that crashed the network—and the ship was left dead in the water for several hours.

What are some lessons for cybersecurity research that might be drawn from this episode?

- Net-centric operations may have a very intimate connection to commercial information technology. Indeed, the day has long since passed when the DOD can rely on custom-built information technology—and its reliance on commercial IT for all kinds of functions means that insecurities in the commercial IT base may have a potentially devastating effect on vital military functions.

- Humans are part of any IT system. One might argue, as the Navy did at the time, that it was therefore "human error" that crashed the network rather than a problem with the network itself. But because we assume that cyber-adversaries are smart and highly motivated, inducing human error is a strategy that an adversary might well employ.

- A decision could have been made to provide a back up means of controlling ship propulsion, so that a crashed network would not leave the ship dead in the water. A decision to do so would not have depended on a detailed knowledge of cybersecurity as cybersecurity is traditionally construed, but rather on a philosophy of system design that anticipates failures and provides for ways of mitigating and containing their impact.

- The Yorktown was a testbed for new technologies, and thus one might argue that failures should be expected. True enough, but the argument is incomplete. Testbeds often have a way of turning into a legacy base—that is, even though we built testbeds and experimental applications thinking that we can throw them away when we "get serious" about an application that will be deployed for real, in practice the design concepts from these testbeds and experimental applications often remain embedded in the new generation. This reality suggests that understanding how to provide security for legacy systems is a vital dimension of cybersecurity research.

These comments are not intended to denigrate the conceptualization of cybersecurity as a technological problem, because in many ways, it is a technological problem. One of the six categories of needed research outlined in our report is *blocking and limiting the impact of compromise.* This category is relatively traditional, including the design and development of secure information systems and networks that resist technical compromise. Somewhat unusual in the topics for inclusion in this category was the need for research to understand how to contain the damage from a penetration, how to lock down a system under attack, and how to recover quickly from a successful attack. Because absolute security of an information system never can be guaranteed, that

research is needed so that recovery from a successful attack can be accomplished as expeditiously as possible.

But it would be a bad mistake to conceptualize cybersecurity as only a technological problem. Indeed, we found in our work that areas ranging from anthropology, sociology, design, economics, law, psychology, human factors, and organizational theory were relevant to cybersecurity.

Consider, for example, a proposition that very few cybersecurity experts would deny—the most effective security measures or technologies provide very little benefit if they are not deployed in operational systems, and even if they are deployed, they provide very little benefit if they are not used, or even worse, misused or bypassed by users because they are not well understood or they interfere with getting work done. Today, a great deal of security functionality is often turned off, disabled, bypassed, and not deployed because it is too complex for individuals and enterprise organizations to manage effectively or to use conveniently.

It is easy to believe that in military organizations, a senior commander can simply order his subordinates to comply with all necessary security measures—and to some extent, this is true. Nevertheless, under the pressure of combat operations, it is often the case that faithful execution of security procedures gives way to the expediency of circumventing those procedures if they are cumbersome. Indeed, you might want to inquire whether the use of secure STU-III telephones increases or decreases at the onset of combat operations.

Such reasons suggest that cybersecurity construed in purely technological terms may well be ineffective in an operational context. Thus, our view of necessary cybersecurity research includes a category focused on *promoting deployment and effective use* of cybersecurity technologies. This category includes research on technologies that facilitate ease of use by both end users and system implementers, incentives that promote the use of security technologies in the relevant contexts, and the removal of barriers that impede such use. Measures to provide incentives and to remove barriers to the use of security technologies and procedures may have legal, economic, psychological, social, and organizational dimensions.

The NRC report also covered four other categories of necessary research:

- *Enabling accountability*. This category includes matters such as remote authentication, access control and policy management, auditing and traceability, maintenance of provenance, secure associations between system components, intrusion detection, and so on. In general, the objective is to hold anyone or anything that has access to a system component—a computing device, a sensor, an actuator, a network—accountable for the results of such access. An example of research in this category is attribution. Anonymous attackers cannot be held responsible for their actions and do not suffer any consequences for the harmful actions that they may initiate. But many computer operations are inherently

anonymous, which means that associating actors with actions must be done explicitly. Attribution technology enables such associations to be easily ascertained, captured, and preserved. At the same time, attribution mechanisms do not solve the important problem of the unwittingly compromised or duped user, although these mechanisms may be necessary in conducting forensic investigations that lead to such a user.

- *Deterring would-be attackers*. This category includes legal and policy measures that could be employed to penalize or impose consequences on cyberattackers, and technologies that support such measures. In principle, this category could also include technical measures to retaliate against a cyberattacker. One illustrative example of research in this category would facilitate the prosecution of cybercriminals across international borders. Many cybercrime perpetrators are outside of U.S. jurisdiction, and the applicable laws may not criminalize the particulars of the crime perpetrated. Even if they do, logistical difficulties in identifying a perpetrator across national boundaries may render him or her practically immune to prosecution. Research is needed to further harmonize laws across many national boundaries to enable international prosecutions and to reduce the logistical difficulties involved in such activities. Other illustrations are provided in the main text of the report.

- *Crosscutting problem-focused research*. This category focuses elements of research in the above categories onto specific important problems in cybersecurity. These include security for legacy systems, the role of secrecy in cyberdefense, coping with the insider threat, and security for new computing environments and in application domains.

- *Speculative research*. This category focuses on admittedly speculative approaches to cybersecurity that are unorthodox, "out-of-the-box," and also that arguably have some potential for revolutionary and nonincremental gains in cybersecurity.

The committee also examined the lack of substantive progress in closing the gap between the nation's cybersecurity posture and the cyberthreat. Indeed, it observed that after more than 15 years of cybersecurity reports pointing to an ominous threat, and more than 15 years in which the threat has objectively grown, there is not a national sense of urgency about cybersecurity.

The committee concluded that the lack of adequate action in the cybersecurity space could be largely explained by three factors:

- Past reports have not provided the sufficiently compelling information needed to make the case for dramatic and urgent action. If so, perhaps it is possible to paint a sufficiently ominous picture of the threat in terms that would inspire decision makers to take action. Detailed and specific information is usually more convincing than information couched in very general terms, but

unfortunately, detailed and specific information in the open literature about the scope and nature of the cyberthreat is lacking. Many corporate victims of cyberattack, for example, are reluctant to identify themselves as being victims for fear of being cast in a bad light relative to their competitors.

- Even with the relevant information in hand, decision makers discount future possibilities so much that they do not see the need for present-day action. If that is the case, then nothing short of a highly visible and perhaps ongoing cyber-disaster will motivate actions. Decision makers weigh the immediate costs of putting into place adequate cybersecurity measures, both technical and procedural, against the potential future benefits (actually, avoided costs) of preventing cyber-disaster in the future—and systematically discount the latter as uncertain and vague.
- The costs of inaction are not borne by the relevant decision makers. The bulk of the nation's critical infrastructure is owned and operated by private-sector companies. To the extent that these companies respond to security issues, they generally do so as one of the risks of doing business. But they do much less to respond to the threat of low-probability, high-impact (i.e., catastrophic) threats, although all of society at large has a large stake in their actions.

Although these observations were made regarding information technology outside the military sphere, I believe that they—and especially the last two factors—are highly relevant to DOD cybersecurity issues as well.

One might also consider the fact that net-centric operations, broadly writ, depend on dramatically increased access and functionality afforded by modern information technology. But increased access also multiplies the routes through which an adversary can attack us, and increased functionality has required ever more complex systems that are inevitably riddled with vulnerabilities. From a security standpoint, the consequence has been that our increasing dependence on these technologies provides formerly weak adversaries with unprecedented ways of attacking us.

To address these vulnerabilities, the report suggests that we need to reduce the likelihood that an adversary will succeed in penetrating our cyber-defenses and to increase the ease of recovering from successful penetrations of those defenses. But a third logical possibility, also addressed in the report, is to design systems so that critical activities can take advantage of advanced information technology when appropriate and possible but do not require such technology in order to function. In some cases, this may mean providing adequate means for backup in case the necessary IT is unavailable or under attack; in other cases, it may mean foregoing some of the advantages afforded by network-centric operations because the risk is just too large to manage even with backups in place.

Finally, I was asked to comment on coordination within the Federal government of cybersecurity research, which our report addressed. It was our impression that the scope and nature of cybersecurity research across the federal government were not well understood, including by government decision makers, and that no entity within the

federal government had a reasonably complete picture, including classified and unclassified, of the cybersecurity research efforts that the government supports from year to year. To illustrate the issue, in 2004, the President's Information Technology Advisory Committee, backed by the National Coordination Office for Networking and Information Technology Research and Development, was able to determine the DARPA investment in cybersecurity research and development (R&D) for FY 2004 only within a factor of about four (that is, PITAC determined that figure to be between $40 million and $150 million).

Our report argues that an effort to develop a complete picture should distinguish clearly between research and development, including both classified and unclassified R&D; disaggregate (and publish) government-wide budget figures associated with different areas of research focus; and track budget figures from year to year. Further, the report argues for a sustained, coherent, and comprehensive approach to cybersecurity research, and the lack of a mechanism for drawing this complete picture suggests that the U.S. government is not well-organized for supporting such an approach.

Thank you. I will try to answer any questions you might have.