TESTIMONY OF MICHAEL J. ASSANTE

PRESIDENT AND CHIEF EXECUTIVE OFFICER
NATIONAL BOARD OF INFORMATION SECURITY EXAMINERS OF THE UNITED STATES INC.

BEFORE THE
SENATE COMMITTEE ON HOMELAND SECURITY AND GOVERNMENTAL AFFAIRS

U.S. SENATE

Hearing on

SECURING CRITICAL INFRASTRUCTURE IN THE AGE OF STUXNET

November 17, 2010

Testimony of Michael J. Assante
National Board of Information Security Examiners
November 17, 2010

Good morning, Chairman Lieberman, Senator Collins, and members of the Committee. I am pleased to appear here this morning to testify on securing critical infrastructure in the age of Stuxnet.

My name is Michael Assante and I am the Chief Executive Officer of the National Board of Information Security Examiners ("NBISE"). NBISE is a newly-created, not-for-profit, certification body comprised of dedicated staff and a board of experts in information security practice and policy. NBISE is developing assessments, examinations, and certifications designed to uphold the highest standards of professionalism and practice in essential information security disciplines. I am here in this capacity and as someone who has worked in the field of critical infrastructure protection with a focus on industrial control systems security. I have served in the U.S. Navy, been responsible for both physical and cyber security of one of the largest electric utilities in the United States and worked on control system security research at the Idaho National Laboratory. I also recently held the position of the Chief Security Officer at the North American Electric Reliability Corporation ("NERC"), which serves as the Electric Reliability Organization ("ERO") in the United States and much of Canada.

I am pleased that this hearing has been convened to explore the implications of advanced cyber threats on the security of our nation and its critical infrastructure, as exemplified recently by the Stuxnet worm. The Stuxnet code is a worthy centerpiece for this discussion, but I believe this is neither the first nor the last attempt to compromise and use operational systems to effect physical outcomes. Stuxnet is, at the very least an important wake up call for digitally-enhanced and reliant countries; and at its worst, a blueprint for future attackers.

There are many lessons that we must learn from this particularly sophisticated piece of malicious code. Because it will set the course for cyber strategy and policy, our response to this demonstration of the new cyber reality is critical. Developing and implementing effective indicators, defenses, and countermeasures to cyber threats like Stuxnet demands that we look not just to the security community but also to the system designers, planners, engineers, and operators of our essential technology and physical infrastructures. We must take a prudent and proactive approach that enhances our ability to learn and apply knowledge fast enough to manage the dangerous consequences that come with these types of attacks. We can no longer ignore known system weaknesses and simply accept current system limitations. We must admit that our current security strategies are too disjointed and are often, in unintended ways, working against our efforts address the highly-advanced security challenges facing our cyber-dependent critical infrastructures.

My statement will paint a very difficult challenge, but it is important to note that I remain an optimist about our ability to close the gap. This nation, as it has done in countless past contests, should turn to its men and women, in and out of uniform, to muster an effective defense. Our obligation is to effectively organize, train, and equip them to be successful in this important task.

THE DISRUPTIVE INNOVATION THAT IS STUXNET

Simply put, Pandora's Box was opened years ago as the United States became reliant upon digital technology to help operators control complex processes. Stuxnet is an important harbinger of things that may come if we do not use this opportunity to learn about this threat (and others like it) to our infrastructure control systems. No one should be shocked that cyber exploits can be engineered to successfully compromise and impact control systems. Study after study has identified common vulnerabilities found across control system products and implementations. The exploitation of a hard-coded password design in one vendor's implementation will not be an uncommon or isolated occurrence.

Stuxnet is a good example of a cyber threat that was thought to be hypothetically possible, but not considered probable by many. Its features, capabilities, and intended technology target/purpose should disturb security professionals, engineers, businessmen, and government leaders alike. There are three specific reasons why I make this important statement:

First, it appears there is a group of well-resourced people possessing the necessary motivation, who have successfully acquired the knowledge, skills, and capabilities to systematicall develop and launch a highly-sophisticated attack targeting control systems to effect a desired physical outcome. The public occurrence of such a cyber attack is important as it dispels conventional thinking that it is just "too hard" for an attacker to assemble the necessary information, gain familiarity with the technology, acquire the knowledge of specific implementations, configurations and accesses to devise an attack that could disrupt or damage the physical components of an industrial process. It requires more resources and skill, by a cyber attacker, to attack control systems with sufficient confidence to achieve a specific and intended outcome. The authors of Stuxnet have certainly established themselves in this category. I am, however, concerned that an attacker with less means might still be capable of creating havoc or unintentionally causing a more isolated accident.

What is shocking to control system security experts is not that it was done, but that it was done in such a manner as to rely upon pre-programmed code that had the ability to autonomously analyze the system that had been compromised and identify the specific conditions desired for the delivery of its "digital warhead." Most had anticipated more manual attacks capable of achieving negative consequences to physical process, through a more stepped process of compromise, discovery, learning, and action.

The authors of Stuxnet were able to characterize the environments that it would compromise and develop enough capability and logic to defeat anticipated security measures, deal with different configurations, practices and architectures, and act in a very restrained manner until it found its ultimate target. The lesson that we must not gloss over is that highly resourced actors can assemble people capable of planning on how to deal with system variances, security controls, obscure and proprietary technology, and complex industrial processes.

This complex and sophisticated code had been propagating for months before it was identified and recognized as something different. Security researchers inspecting the code soon discovered a code base that was professionally developed and tested. This threat was wrapped in layers of obfuscation like many pieces of custom malicious code, but a closer inspection revealed advanced techniques for circumventing a number of anticipated security controls to include signature-based security tools, behavior-based detection engines and the requirement for files signed with certificates. The authors took advantage of combining known vulnerabilities with a never before seen use of four "zero-day" (or previously unknown) exploits to compromise systems. This allowed the code to escalate privileges, embed itself, propagate further, receive updates, and seek out its intended target. It is not beneficial to speculate on the worm's ultimate target, but I was relieved to hear that it was programmed to inject code only when specific configurations and processes were identified. It is difficult to imagine what would have happened if the authors were interested in creating more general havoc and had programmed the worm to be less constrained in injecting changes to victim controllers.

Second, we must understand that the attacks we should be most concerned with are not designed to disable their digital targets, but to manipulate them in an unintended way to achieve a desired physical outcome. Many professionals have limited their thinking to dealing with the loss of individual elements or capabilities of their control systems and have failed to fully embrace the implications of calculated misuse. I attempted to shed light on this important topic in my April 7, 2009 letter to the electricity industry (also submitted for the record).

The ability of an attacker to access controllers, safety systems, and protection devices and inject valid or malicious code that can change set points, command action, change the expected logic, or suppress a measurement and/or action is alarming. I have participated in research that demonstrated this capability in a controlled environment to understand how it could be done and explore the potential consequences should such a weapon materialize. I believe that the analysis to date has indicated that Stuxnet may be such a weapon. To complicate matters, we have not sufficiently studied nor considered the potential for these types of attacks on large interconnected systems such as the electric grids or in highly controlled and potentially dangerous industrial processes.

We must not put our faith in the possibility that the necessary knowledge of these types of attacks will remain only in the hands of highly-educated and trained process control engineers. The first, generally accepted, industrial control system ("ICS") vulnerability was disclosed to the public in 2005. As of October 1, 2010, the national vulnerability database has 51 vulnerabilities currently listed, and organizations like Critical Intelligence are tracking 119 disclosed vulnerabilities. A dedicated attacker, even without first-hand, detailed knowledge of how to program and apply control system technology can still cause serious damage. I have worked with general cyber security researchers that have developed capabilities to compromise and affect industrial control systems. One can develop such capabilities by accessing production systems, acquiring and experimenting with components, and gathering technical information available on the web. The researchers demonstrated that gaining access to control system data streams, systems, and specific devices can be enough to attack the process being controlled or safeguarded. For example, in modern control systems most of the process safety depends on logic found in the controllers. By analyzing this code one can not only determine what the engineer wants to happen but also what the engineer wants to avoid. One can identify how to cause negative effects by studying the programmed logic, like master-stop conditions, to identify ways to achieve unsafe conditions or override safety shutdown logic.

It is critical that we reverse many of the trends in the control and safety system market that make an attacker's task of causing damage in the physical world easier. One such trend is the convergence of control systems and safety systems at the network-level. An all-too-familiar tale has us achieving greater efficiencies by doing away with silos and leveraging shared network resources. Legacy industrial systems relied on physically separate and functionally independent control and safety systems. The safety system could independently sense and act to ensure the safety or reliability of the system/process. Today, we see an alarming trend towards sharing network resources in space and cost-constrained industrial applications so that the safety system is now only functionally and logically separated from the control system. This dangerous trend provides far too great an opportunity to an attacker in high-consequence environments, enabling access to both safety and control systems from a single point of entry. This is significant as an attacker can explore and manipulate the safety system; removing planned safeguards, before misusing the control system to create a dangerous condition.

Finally, our current defense and protection models are not sufficient against highly-structured and resourced cyber adversaries capable of employing new and high-consequence attacks. Our defensive thinking has been shaped by the more frequent and more survivable threats of the past. As a result of this paradigm, our security architectures and security market solutions are mostly reactive by nature. This is a logical behavior because the market and security organizations primarily respond to attacks that have been observed in sufficient numbers to warrant the development of countermeasures and new practices. This behavior is less risky in applications that adhere to the restrictions of scale, time and geographic space. An example of this is the realm of physical security, where the capabilities of attackers typically evolve at a slower pace and an attacker is often constrained when using a new capability in a fashion that is limited by both time and their ability to be physically present in a single location. Computers and networks provide for less constrained circumstances and the ability to develop new attack techniques and tactics evolve very rapidly and change with the very technology it attacks. One of the more significant elements of a cyber threat, contributing to the uniqueness of cyber risk, is the cross-cutting and horizontal nature of networked technology that provides the means for an intelligent cyber attacker to impact multiple assets at once, and from anywhere, and without fear of attribution.

This means that while current cyber defense tactics, security architectures and tools are necessary and can be responsive to the most likely of threats, they are not sufficient to deal with new advanced threats. The optimist always points to a new type of security tool or practice as the solution to current protection inadequacies. I have watched the industrial system and security community rush towards the technology of white-listing as an answer to the Stuxnet worm. This technique can be effective and is an important option to deploy, but we should not believe that if it had been necessary to assure their success, the authors of the Stuxnet worm would not have simply developed a way to counter this measure. In support of this assertion, the experts at Symantec studied the capability of this worm and concluded that the authors had the resources necessary to counter anticipated market-produced security solutions.

I appreciate the strong desire to create reactive solutions but I believe that desire leads to tunnel vision limiting our approach to prevention and detection strategies. It is important that we do not place our faith solely in conventional security tactics and tools. The shortfalls of our current tools are becoming apparent to many specialists desperately working to clean up the Stuxnet infections. They have reported great difficulty in removing the problem. Many of the security research programs funded by the government are working to research and implement yesterday's general information technology (IT) security measures into today's operational ICS and Supervisory Control and Data Acquisition ("SCADA") systems. These efforts were proven ineffective in general IT systems against more advanced threats and don't represent a wise use of our resources. They will not significantly improve the safety and reliability of our physical infrastructures against well-resourced attackers. In short, we are preparing for the next battle by using pre-2003 strategies and weapons.

THE SUSCEPTIBILITY OF OUR CRITICAL INFRASTRUCTURE

I would like to provide some additional perspective by focusing on the electricity infrastructure because I am experienced with its operations and protection challenges. I would like to begin by prefacing my comments with an observation about the electric utility industry. I have seen this industry express a genuine desire to put system reliability first and can appreciate the unique pride that comes with serving communities with such an essential service. Cyber and physical security are two of many reliability risks faced by power system planners and operators. It is important to note that there are also various High Impact Low Frequency threats to consider. While the industry deals with some physical security events like copper theft on a regular basis, other more complex physical and technical threats or hazards, such as terrorism, electromagnetic pulse and space weather, are concerns as well and will require careful consideration to develop appropriate and effective mitigation.

Cyber-related threats pose a special set of concerns because they can arise virtually anytime, anywhere and change without warning. Unlike other operational and reliability concerns, such as extreme weather or the probabilistic failure of mechanical equipment, cyber-related events and threats could occur through accidents or by actors who intentionally manipulate or disrupt normal operations as part of a premeditated design to cause damage.

The susceptibility of our modern interconnected and digitally reliant infrastructures is well established. We must now find answers to research and engineering challenges associated with protecting those infrastructures for which significant damage, and the lack of availability for extended periods of time, would have catastrophic impacts on society. Efforts to modernize our nation's electric power infrastructure through the overlay of two-way digital communications and highly-automated digital control (to create the "smart grid") are based on the desirable promise of greater energy efficiency and system performance. Of course, more technology typically adds more complexity and interconnectedness. We should continue to seek progress, but also recognize the need to close the gaps in the software and system engineering foundations necessary to ensure that new smart grid functionality will be secure, safe, survivable, reliable, and resilient.

The most fundamental of these research and engineering challenges is how to design, configure, and operate the smart grid's systems and components in a manner that prevents an adverse cyber-physical event (whether accidental or malicious in origin) from having a catastrophic impact on the grid and on society at large. For examples of the kinds of adverse events, see the "Coordinated Attack Risk" chapter of the recent joint report by NERC and the U.S. Department of Energy ("DOE"), entitled *High-Impact, Low-Frequency Event Risk to the North American Bulk Power System*[1].

REGULATING CYBER SECURITY & NERC'S MANDATORY RELIABILITY STANDARDS

NERC-developed critical infrastructure protection ("CIP") Reliability Standards represent an early attempt to manage cyber security risks through mandatory standards with significant penalties for non-compliance. It is clear to me that the standards as written and implemented are not materially contributing to the management of risk posed by advanced cyber threats, such as the Stuxnet worm.

The standards are comprised of forty-three specific requirements designed to provide a minimum set of sound security practices that, if properly implemented, can serve as a simple foundation to be built upon. The perimeter protection model taken by the standards is more aligned with cyber threats of yesterday and is most effective against less structured types of cyber attacks. The standards also include significant gaps and exclusions, but their greatest weakness is in how they have been implemented. These standards have polarized the industry and have imposed requirements on a highly-dynamic and not fully understood area of system risk. The result has been a conscious and inevitable retreat to a compliance/checklist-focused approach to the security of the bulk power system. I have observed security programs that have suffered from resources being channeled into compliance activities and a hesitance, or even outright refusal to try ideas, measures, and security practices that exceed what is called for in the standards. Unfortunately the NERC CIP Standards have become a glass ceiling for many utility security programs, which prevents the emergence of the type of security programs we need to deal with Stuxnet-like attacks.

I believe the level of expertise needed to create standards that achieve security objectives and ensure safety and reliability must be found not in one quarter, such as within industry or a specific government agency, but within the community at large. There must be a process that will maximize the contributions from security, industry, and technology experts, while establishing a mechanism to identify the best performance measures to manage the risk. At this point it seems clear that saying "industry knows best" about what is important enough to deserve enhanced protections and how best accomplish mitigation of advanced cyber risks is not completely accurate. We are all amateurs in this quest. I have first-hand experience working with mandatory standards. It is clear to me that standards are a good tool to manage risk when it is either well-bounded and understood or when the standard simply codifies well-honed industry practices that are proven to be successful. Mandatory cyber standards fail both of these conditions, mainly because advanced cyber threats are not probabilistic in nature, but represent a

[1] Available at: http://www.nerc.com/files/HILF.pdf

co-adaptive risk. The best decisions will only come from an active engagement of experts in and out of government, industry, and the larger community under strong leadership.

Regulation, although necessary, should be re-evaluated and designed to emphasize learning, enable the development of greater technical capabilities through more qualified staff, and discourage the creation of a predictable and static defense. For example, both asset owners and technology vendors must be prepared to recognize and be required to report significant and unique security incidents to key stakeholders, including peer organizations. The requirement to report cyber incidents should be accompanied by a fair and limited safe harbor to promote this essential requirement (consider the model established by FAA and NASA regarding commercial pilot event reporting). Informed regulatory oversight will be necessary to shepherd the process so it is capable of producing timely results without harming the very systems we are trying to protect.

A joint rule-making arrangement between an office responsible for coordinating regulation across the critical infrastructure sectors and a more directly aligned government agency or independent commission is best suited for this leadership task. I believe that more clearly defined federal authority and funding is needed to address specific and imminent cyber security threats to critical infrastructure. If we were discussing the electricity sector the coordination authority would need to work jointly with FERC, as they better understand the reliability issues associated with the technology and operations of the sector. Again in the case of the power system, this effort will need to be closely coordinated with Canada as oir physical infrastructure is critically tied to that of our northern neighbor. Ultimately we will have to require security concerns to be factored into design choices and architectures, not just addressed by technical system security standards.

IMPLICATIONS

Cyber threats will continue to evolve and the extent of their potential to negatively impact our control systems is not yet fully appreciated. The potential for an intelligent attacker to exploit a common vulnerability that impacts many assets at once, and from a distance, is one of the most concerning aspects of this challenge. This is not unique to the electric sector. Addressing it, however, will require asset owners to apply additional, new thinking on top of sound operating and planning analysis when considering appropriate protections against these types of threats. It is imperative as a nation that we seek to broaden the understanding of cyber risk concerns facing the interconnected networks and critical infrastructure. We must develop and implement protection strategies that accept the unfortunate, though probable, reality that many of our networks are already contested territory. Accepting this important assumption will help stimulate industry and community efforts to develop new and prudent approaches to address the most material risks.

In the realm of cyber, we must recognize the potential for simultaneous loss of assets and common modal failure in identifying what needs to be protected or engineered to be more survivable. This requires a shift of our priorities from a prevention-heavy approach to reduce the likelihood of such an event from occurring to a greater focus on minimizing the possible consequences of such an event. We must tackle these types of threats by investing in the development of our security professionals, engineers, and operators, and establish countermeasures and mitigations in a far more comprehensive manner. This requires us to consider not only security, but also how we can design and engineer survivability into our complex systems, achieve resilience in our organizations, technology and process, and better prepare to respond and recover.

A significant cause for concern is that much of the information about cyber-security-related threats remains classified in the homeland security, defense and intelligence communities, with restricted opportunity to share information with security researchers, technology providers and affected private-sector asset owners. Our nation's critical infrastructure is placed at significant risk as a result of limited progress to support learning and the application of newly gained knowledge to protect or even respond to and recover from advanced cyber threats.

A mechanism is needed to quickly validate the existence of advanced threats and to ensure information is appropriately conveyed to and understood by asset owners and operators in order to mitigate or avert cyber vulnerabilities. A complex cyber threat cannot be easily contained and has the potential to undermine the integrity of systems owned by governments and private sector organizations alike. We must develop a better framework for tapping into the best and brightest, whether they are specialists holding clearances in the federal government, professionals conducting cutting-edge research into security problems (for example, Symantec engineers and ICS security specialists), those on the ground managing ICS, or others developing and providing technology, or managing complex control system environments. Efforts should be taken to develop standing and situation-based pools of expertise to quickly analyze specific threats and develop guidance to respond, mitigate and if possible, protect critical systems. Critical infrastructure organizations need to develop the ability to identify information relevant to the risks they face and work with the broader national security community to better understand adversary intent, capability (tactics, techniques and procedures both demonstrated and assessed) and opportunity to effectively prioritize and structure countermeasures and mitigations. Government agencies, asset owners, technology providers and researchers have clung to our different identities for too long. Having a common mission is not enough we need to develop the policies, practices, and tools to operate in a unified manner, much like coalition military operations. We need to raise both our individual and collective community capabilities to address these sophisticated and dangerous threats.

I would like to specifically emphasize one of the necessary investments to combat advanced cyber threats like Stuxnet. Though the size, configuration, and function of information networks can vary widely, there is a single feature common to each of them: behind every firewall, system architecture, and vulnerability assessment stands an information security professional. Through the years, working as a Chief Security Officer at a major utility, or supporting researchers, coordinating protection efforts while at NERC, I have gained an appreciation for the importance and the difference made by skilled and well-developed people. I have never understood why we have not embraced better training and development methods for our frontline security and operations staff. We train pilots using advanced simulators to deal with difficult conditions and mechanical failures. Why do we not use simulators to allow security and operational staff to experience low frequency but high consequence attacks against the systems they defend and operate? Why do we not use performance-based examinations to qualify our most important resources? We have allowed chance to be our school house, where targeted organizations simply suffer in silence, not willing to pass along the tough lessons they have learned to others.

Effective security and response against highly advanced cyber threats requires a current understanding of what adversaries are capable of, an opportunity to experience directed attacks to become familiar with observables and experiment with response actions, and the use of a team training framework to optimize defender tactics, techniques and procedures. We must embrace virtual gaming technology and look for ways to stimulate defensive technology with simulated attacks so what has traditionally been only a hypothetical or has been overlooked can made real for the purposes of learning and preparing. It is time that we more formally prepare these individuals and ensure they are competent, prepared, and capable of making the right decisions day-to-day and during emergencies.

CONCLUSION

I commend this Committee for its exploration of the implications that advanced threats like Stuxnet pose to our critical infrastructure and nation. I look forward to supporting your efforts in any way possible. Stuxnet should cause all of us to re-think how we are prepared to protect, respond, and survive future cyber challenges to operational technology like control, safety and protection systems. We must be better prepared to learn about our weaknesses; identify and understand new threats; and make better design, deployment, and operations decisions. We must waste no more time in debating our susceptibility. We must accept that well-resourced adversaries are currently able to achieve primacy by developing unique and creative tools to compromise and affect control system technology. These adversaries may also be capable of causing damage to industrial processes in difficult to anticipate ways. It is time to turn our attention to addressing known weaknesses, researching consequences, designing our security, training and preparing our operations staffs and finding ways to make our systems more resilient. The following steps are necessary:

- Remove and remediate architectural weaknesses, known vulnerabilities, and poor security designs in industrial control systems.

- Promote greater progress designing and integrating security/forensic tools into control environments. But, put your faith and focus in people not the tools of the day.

- Prioritize our efforts by jointly studying the potential consequences that may result from directed and well resourced attacks of control, safety and protection systems in high risk segments of the critical infrastructure. In the cases where the consequences are unacceptable we must assume the attacker can successfully defeat our security and therefore direct our efforts to engineering away that risk with more survivable designs and practices.

- Organize a well-funded, multi-year research & development program to design toward a more resilient infrastructure. The research should include safety system design; explore dedicated networks; architecting complex systems to severe system optimization and preserve core system functions, when needed.
- Establish new regulation in the form of risk-based performance requirements that value learning, promote innovation, and better equip/prepare control environments and the teams that protect, operate, and maintain them. The current regulatory structure will not, in my view, be capable of achieving this end. Legislation should include the need for more sharply defined federal authority to address specific and imminent cyber security threats to critical infrastructures in the form of emergency measures.
- Require critical infrastructure asset owners and control system vendors to report industrial control system specific security incidents and the U.S. government must provide up-to-date information to asset owners and operators on observed adversary tactics and techniques, especially when investigations reveal attacker capabilities to side-step or exploit relied upon security technologies.
- Invest in the workforce that defends and operates our infrastructure systems. Scalable immersive training environments and local simulator/stimulator training technology should be used to optimize the development of defender skills. The same workforce should then be qualified through periodic rigorous performance-based assessments and, where appropriate, examinations.

My greatest fear is that we are running out of time to learn our lessons. Stuxnet, although difficult to hijack or modify by others, may very well serve as a blueprint for similar but new attacks on control system technology. We know that ordinary high-risk practices, such as the use of USB sticks by plant personnel and contractors, must be modified. We know that well-known security weaknesses in ubiquitous technologies need to be re-evaluated and protected. We know that addressing security at the network and general IT layer only addresses one of many attack paths and we must start addressing the exploitable weaknesses of field control devices (such as Remote Terminal Units, Programmable Logic Controllers, and other Intelligent Electronic Devices). Ultimately, we know that our conventional approach to more common security threats will be necessary but woefully insufficient to protect these systems from the next Stuxnet-like cyber threat. We must act now to develop our greatest resource in this contest; the professionals that defend, operate, and protect our critical systems and infrastructure.

Respectfully submitted,

/s/ Michael Assante

Michael Assante
President and Chief Executive Officer

National Board of Information Security
Examiners of the United States, Inc.
2184 Channing Way, #304
Idaho Falls, ID 83404
(208) 557-8026
(973) 860-0921 – facsimile
michael.assante@nbise.org

NERC Chief Security Officer Letter to Industry, dated April 7, 2009: