Input to the Commission on Enhancing National Cybersecurity

Submitted by: Roger R. Schell
Organization: Aesec Corporation

Topic Addressed: Critical Infrastructure Cybersecurity

Executive Summary of Comments

Industrial Control Systems (ICS) which control our **critical infrastructures**, including capabilities like bulk power generation, are unusually vulnerable to cyber attacks. The National Research Council (NRC) has reported that such an attack on the bulk power system "could deny large regions of the country access to bulk system power for weeks or even months. An event of this magnitude and duration could lead to turmoil, widespread public fear, and an image of helplessness". While this is the subject of many ongoing current and proposed future efforts, very few are addressing the core underlying challenge of ICS security: namely, that the ICS is fundamentally insecure and in its present form cannot be made secure. A primary cause of this is that they are built on top of operating systems that cannot be made secure.

A promising and innovative architectural approach to address this challenge uses a layered security strategy to instantiate ICS applications, bulk power generation, on top of an existing very high assurance general purpose operating system designed to meet verifiable protection as codified by the National Security Agency (NSA) in TCSEC/TNI "Division A". In essence, this alternative proposes to "teach ICS" to the underlying fundamentally secure cyber systems. This includes rigorous application of information hiding, layering of modules, centralized protection mechanism, formal specification of security policies, trusted distribution, strong configuration management, and strongly-typed implementation language. This approach integrates an existing security kernel as the operating system and associated hardware into a reusable Trusted Device, reminiscent of a reusable commercial computer motherboard. On top of this kernel, applications would execute the physical actions required by ICS application, e.g., the power system. Since these actions always produce physical results, the application can either be proven correct or exhaustively tested for all physical possibilities, and thus guarantee cyber physical safety.

We recommend managing the inherent scale and complexity of critical infrastructure control systems using a coordinated multi-disciplinary approach to create a holistic and reusable security architecture that integrates varied, disciplines. Extant ICS technology and practice are systematically decomposed to reflect its implicitly diverse security policies. High assurance trusted systems technology is adapted to more effectively host ICS functions and enforce associated policies. Architectural longevity is provided by proactively mitigating hardware level vulnerabilities including supply chain issues. The resulting carefully designed security architecture and its embodied Trusted Device will result in a powerful new secure ICS paradigm. The government then has a unique opportunity to change the cyber security game. They should aggressively engage ICS manufacturers by sponsoring prototypes and providing a market using proven commercial security kernel technology. This recommended approach can within a couple of years make our critical infrastructure dramatically more trustworthy.

Topic Area Challenges and Approaches

The following comments provide information on current and future challenges, promising and innovative approaches to address those challenges, recommendations, and references to inform the work of the Commission.

Current and future trends and challenges in critical infrastructure cybersecurity

An Industrial Control System (ICS) built on commercial operating system with public internet connections cannot be made secure. The large amount of code, the many pathways through it, and the continual changes means that it will never be secure and will always have exploitable holes. The CEO of RSA said at the April 2015 RSA Conference, "**computer security has failed**". The case to be made that he has accurately portrayed the current and future trends is pretty simple:

1. The current cyber security problem is even worse than people perceive because of the devastating damage that can be done with **software subversion** by a witted adversary.
2. What we are currently doing, mostly **monitoring and surveillance, is not working**, and we know from science never will, regardless of how much more resources we throw at it.
3. The dominant future trends are not much better, since we are not using scientific principle of the "reference monitor", whose high assurance implementation is called a "**security kernel**". Verifiable and mature security kernels have been highly effective when used, but for a number of years this has been largely ignored by the government, industry and educational communities, perpetuating the problem.

There are many cyber attacks on critical infrastructure that can cause short-term disruptions, but the research proposed here is concerned with attacks that can cause long-term disruptions. Our focus is on the attacks that can result in physical damage to critical infrastructure that can be extremely expensive and time-consuming to repair. Moreover, automated ICS devices in critical infrastructure number in the millions, are often installed in remote, hard to reach locations, and are typically left in place for decades before replacement. The "penetrate-andpatch" model that is the current "best practice" in IT security does not scale in the critical infrastructure environment. Creating a secure critical infrastructure calls for a radically different model.

This is illustrated by the Bulk Power Generation system, which has been identified as one of the most critical infrastructures that could suffer lasting physical damage from coordinated cyber attacks [1]. These attacks are known to be able to cause physical damage to turbine generators and to Extremely High Voltage Transformers, both of which are components with extremely large replacement cost ($2M to $8M or more each), transport cost for heavy units (250,000 to 850,000 lbs), and time to replace (months to years). While the power grid is engineered with capacity to cope with loss of one or even two major generation plants, a coordinated cyber attack could result in destruction of generating capacity far beyond the power grid's capability to compensate, resulting in multiple-month-long extended blackouts over large regions of the continent.

The current and future trends and challenges in critical infrastructure cybersecurity are still much as summarized by then Secretary of Defense Leon Panetta in 2012 [2]. He warned that the United States was facing the possibility of a "cyber-Pearl Harbor" and was increasingly vulnerable to foreign computer hackers who could dismantle the nation's power grid, transportation system, financial networks and government. He described the collective result as a "cyber-Pearl Harbor that would cause physical destruction and the loss of life, an attack that would paralyze and shock the nation and create a profound new sense of vulnerability."

Progress being made to address the challenges

There are two distinct answers to the question of the progress being made to address the challenges. On the one hand the current "best practices" approach had made little progress, and little can be expected by continuing to pursue this approach. On the other hand, the security kernel technology (based on the reference monitor abstraction) has made tremendous strides in demonstrating it ability to provide security for of large, complex, distributed systems – although it almost totally ignored by the current trends.

In the first case, over at least the past few decades considerable effort has been invested in various "patches" trying to secure existing critical infrastructure control cyber systems. The results are piecemeal, marginal, ineffective, un-scalable, and generally woefully inadequate – "an arms race we cannot win" [3].

With all the time and money we have invested in trying to "teach security" to existing ICS power generation systems we have accomplished little to nothing in protection against (even moderately) determined adversaries. As Kaspersky has put it, "Most automated control systems were not created with security in mind… ideally all Industrial Control System (ICS) software would need to be rewritten incorporating all the security technologies available and taking into account the new realities of cyber attacks…the alternative would be a secure operating system, one onto which ICS can be installed and which can be built into the existing infrastructure."

In the second case tremendous progress is described in some detail in a recent publication [4] by Prof. Mark Heckman of USD based on the Reference Monitor abstraction. He notes, "*An associated systematic security engineering and evaluation methodology was codified as an engineering standard in the Trusted Computer System Evaluation Criteria (TCSEC). This paper explains how the TCSEC and its Trusted Network Interpretation (TNI) constitute a set of security patterns for large, complex and distributed systems and how those patterns have been repeatedly and successfully used to create and evaluate some of the most secure government and commercial systems ever developed.*"

This paper points out that the potential positive impact of the **security kernel is transformative**:

1. At least half a dozen security kernel-based operating systems that ran for years (even decades) in the face of nation-state adversaries **without a single reported security patch** – not ever! That is truly a paradigm shift. What alternative approach has come close to that?
2. Although it can take 10-15 years and 10s of millions of dollars to build and evaluate a security kernel, once completed the systematic engineering process used greatly

facilitates long-term maintenance of security assurance through technology refresh. This has repeatedly delivered **affordable secure systems in a couple of years**.

3. The major investment for a general purpose **security kernel is highly reusable**. This is demonstrated by OEM deployments of highly secure systems and products, ranging from enterprise "cloud technology", to general purpose data base management systems (DBMS), to secure authenticated Internet communications, by applying commercially available security kernel technology. What other way do we have to accomplish that?

The most promising approaches to addressing the challenges

To be sure there are vocal (mostly uninformed, poorly advised, or threatened) detractors. However, the transformative potential, when contrasted with the grave cyber security dangers, **makes it only prudent to aggressive apply** the security kernel in at least reference implementations for key elements of the critical infrastructure. Such an aggressive application needs cyber security engineers well-educated in the scientific principles of the Reference Monitor and their application to systems. But to make that sustainable an even more pressing need is to produce future university faculty with a mastery of the Reference Monitor, and its security kernel implementation. This can be approached as three strong "legs of the stool" that are primary considerations in the systematic engineering of secure systems:

1. Significantly mitigating software **subversion** – Software subversion is the likely tool of choice for a determined adversary to breach system security. As the ACM has noted, "addressing security **requires a different mindset** from traditional engineering." The primary means for software subversion are Trojan horses and trap doors.
2. Mandatory access control **(MAC) policy** – Only MAC policy can with high assurance enforce secure information flow. Only MAC can confine Trojan horse attacks.
3. **Verifiability** – The reference monitor implementation, known as a security kernel, is the only known technology for systematically achieving verifiable protection.

Heckman summarizes several real-world examples where, "This is demonstrated by OEM deployments of highly secure systems and products, ranging from enterprise 'cloud technology' to general purpose data base management systems (DBMS) to secure authenticated Internet communications, by applying commercially available security kernel technology."[4] Heckman additionally describes completed research prototypes in the past few years for things like source-code compatible secure Linux and a standards-compliant highly secure Network File Service (NFS). Because the Commission may not be familiar with the commercially available OEM security kernel technology, and example product description for the Gemini Multiprocessing Secure Operating System (GEMSOS) Real Time Operating System (RTOS) is attached as a separate document.

What can or should be done now or within the next 1-2 years to better address the challenges

As noted earlier several government leaders have expressed concern that we face an existential cyber security threat to industrial control systems (ICS) in the critical infrastructure, such as the power grid. Use of a security kernel can within a couple of years make our critical infrastructure

dramatically more trustworthy. The government has a unique opportunity to change the cyber security game. To do so we recommend they aggressively engage ICS manufacturers by sponsoring prototypes and providing a market using proven commercial security kernel OEM technology.

Many ICS implementations, including bulk power generation, use as their fundamental control component what is called a PLC (Programmable Logic Controller). The most basic action is for the government to sponsor a leading PLC manufacturer to incorporate a mature security kernel operating system to produce a reference implementation of a verifiably secure PLC. In parallel with that, to provide a viable market for such secure PLC, the government needs to take a leadership role with a specific segment of the critical infrastructure, e.g., bulk power generation, to not only provided a complete functional prototype demonstration with the secure PLC, but also to formulate regulatory requirements that future PLCs must be verifiably secure.

What should be done over the next decade to better address the challenges

Over the next decade the various individual segments of the critical infrastructure (starting with that use for the above prototype) with government leadership, needs to systematically create, validate and promulgate, technical standards and compliance requirements to insure critical components, and their composition, is verifiably secure.

Future challenges that may arise and recommended actions that individuals, organizations, and governments can take to best position themselves today to meet those challenges

Successfully pursuing the recommended approach will create and validate a new and innovative security architecture with high impact and short term payoff for the nation's critical infrastructure, e.g., bulk power generation. These improvements will also be directly applicable to control systems for other critical infrastructures, including oil and gas production, water, transportation, communications, banking, etc.

Yet, futures challenges will arise to other components of our social fabric that are increasingly reliant on information technology. The focus on high assurance security architecture will point the way to systematically extend the results to these new challenges. Introducing the reusable Trusted Device at the core of the architecture makes it significantly agnostic to the specific application context, while retaining its properties for substantial reduction in development time and resources when applied in other contexts. This will yield demonstrable mitigation of coordinated cyber attacks, including subversion of the sort widely reported in StuxNet.

REFERENCES

1. NERC-DOE HILF Report: High Impact, Low- Frequency (HILF) Event Risk to the North American Bulk Power System, http://www.nerc.com/pa/CI/Resources/Documents/HILF_Report.pdf.

2. Bumiller, Elizabeth and Shanker, Thom. (2012, October 12). Panetta Warns of Dire Threat of Cyberattack on U.S. *The New York Times*. Retrieved from http://www.nytimes.com/2012/10/12/world/panetta-warns-of-dire-threat-of-cyberattack.html?
3. Kelly Jackson Higgins, *IBM: The Security Business 'Has No Future'*, Information Week Dark Reading, 4/10/2008; http://www.darkreading.com/ibm-the-security-business-has-no-future/d/d-id/1129423
4. Heckman MR, Schell RR. *Using Proven Reference Monitor Patterns for Security Evaluation*. Information. 2016 Apr 26;7(2):23; http://dx.doi.org/10.3390/info7020023

ATTACHMENT 1 (Separate Document)

GEMSOS™ Security Kernel RTOS

Product description

Aesec Corporation
725 Cowper Street #46
Palo Alto, CA 94301-2650