



Calhoun: The NPS Institutional Archive
DSpace Repository

Faculty and Researchers

Faculty and Researchers' Publications

2011

The Vulnerability of Nuclear Facilities to Cyber Attack; Strategic Insights: Spring 2010

Kesler, Brent

Monterey, California. Naval Postgraduate School

<https://hdl.handle.net/10945/25465>

This publication is a work of the U.S. Government as defined in Title 17, United States Code, Section 101. Copyright protection is not available for this work in the United States.

Downloaded from NPS Archive: Calhoun



<http://www.nps.edu/library>

Calhoun is the Naval Postgraduate School's public access digital repository for research materials and institutional publications created by the NPS community. Calhoun is named for Professor of Mathematics Guy K. Calhoun, NPS's first appointed -- and published -- scholarly author.

Dudley Knox Library / Naval Postgraduate School
411 Dyer Road / 1 University Circle
Monterey, California USA 93943

The Vulnerability of Nuclear Facilities to Cyber Attack

Brent Kesler

Introduction

In June 2010, U.S. Senators Susan Collins, Joseph Lieberman, and Tom Carper introduced the Protecting Cyberspace as a National Asset Act. One of its many aims is to protect critical infrastructures in the United States from cyber attack. In January 2011, Brandon Milhorn, staff director of the Senate Homeland Security and Governmental Affairs Committee, defended the bill, saying that it would prevent a hacker from opening the floodgates of the Hoover Dam. Peter Soeth, a spokesman for the US Bureau of Reclamation, the agency which manages the Hoover Dam, objected to that example, arguing that “These types of facilities are protected by multiple layers of security, including physical separation from the internet, that are in place because of multiple security mandates and good business practices.”¹

This dispute over the Hoover Dam demonstrates the classic pattern of debate over critical infrastructures and their vulnerability to cyber attacks. Most of the process control systems designed to manage critical infrastructures, such as electric grids, oil pipelines, and water utilities, use specialized hardware and proprietary protocols. However, since the 1990s, the managers of these infrastructures have been integrating their control systems with computer networks built from commercial off-the-shelf operating systems, such as Windows and Unix.² This has simplified the task of managing facilities remotely, but it has also made process control systems vulnerable to attack over the internet. Alarmists point to these connections as vulnerabilities that pose almost epic threats; skeptics immediately dismiss such fears, claiming that the necessary measures to prevent a catastrophic cyber attack have already been implemented. History suggests the truth lies somewhere in between.

As a relatively young field, national cyber security policy has been open to speculation about potential threats. However, in 2011, network operators have accumulated enough experience and data from real world attacks to draw a more realistic picture of the threats facing critical infrastructures. This paper will examine the history of cyber security incidents at nuclear facilities to assess the extent to which recorded vulnerabilities pose an “epic” threat. Specifically, it will examine three cyber incidents that occurred at U.S. nuclear facilities between 2003 and 2008. It will then turn to details of the 2010 Stuxnet attack against the Iranian nuclear program to outline similarities with the three U.S. incidents. The lessons from these four incidents suggest that situational awareness and other security measures are too weak in their current state to guarantee that a catastrophic attack will never happen. However, it will also argue that launching a catastrophic attack is not simple and requires a sophisticated adversary. The article will then turn to gaps in nuclear regulation that policy makers should consider when formulating cyber security policies, not only for nuclear facilities, but for other critical infrastructures.

¹ David Kravets, “No, Hackers Can’t Open Hoover Dam Floodgates” *Threat Level*, (*Wired* blog), February 3, 2011. <http://www.wired.com/threatlevel/2011/02/hoover/>

² Martin Stoddard et al, *Process Control System Security Metrics – State of Practice*, Institute for Information Infrastructure Protection, August 2005.

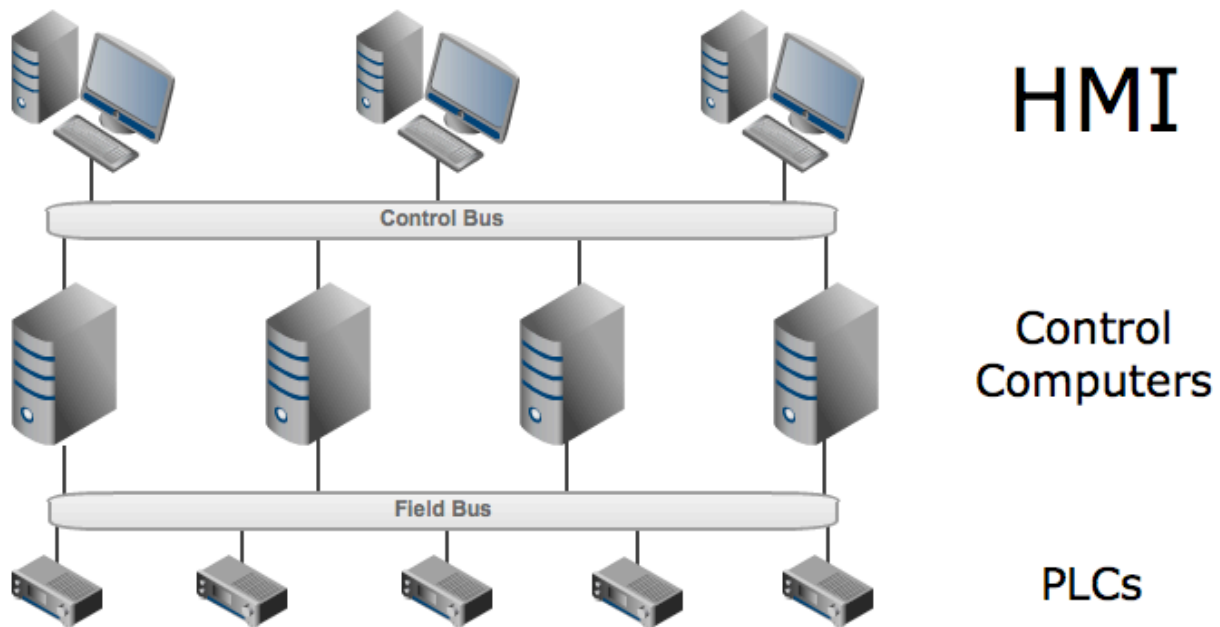


Figure 1: Highly simplified representation of a process control network

Process control systems

Historically, critical infrastructures have used two kinds of control systems: supervisory control and data acquisition (SCADA) systems that quickly gather remote field data, and distributed control systems (DCS) that manage automated manufacturing processes. Over time, these systems began to share many of the same technologies and features, making them less distinct from each other. However, given their separate histories, much of their distinct terminology remains. Other terms, such as integrated control systems (ICS) or instrumentation and control (I&C) are also used, depending on the traditional practice of the facilities using these systems. This paper will collectively refer to these technologies as process control systems (PCS).³

Process control systems come in any number of complex architectures, but a general pattern holds for most facilities. The *control network* is the collection of computer systems which directly monitor and control plant operations. At the top are the *human-machine interfaces* (HMI) that display data from plant equipment and allow technicians to adjust their operations. These are often Windows or Unix based computers. HMI communicate over a *control bus* with other computers that monitor and control operations using software that is less user-friendly. These computers communicate over a *field bus* with *programmable logic controllers* (PLC), hardware that directly adjusts the various motors, sensors, actuators, and other physical components at the heart of a plant's operations.⁴ This is a highly simplified description of a control network; structure and terminology will vary.

Power plants also have office networks for business purposes. The office networks often collect data from control networks and have connections with a wider corporate network over the internet.

³ Stoddard et al, *PCS Security Metrics*.

⁴ K. Korash et al. *Emerging Technologies in Instrumentation and Controls: An Update*. (Oak Ridge: Oak Ridge National Laboratory, 2006), 25-28.

Connecting control networks with business offices and the larger corporate network makes it easier for managers to match plant operations with business goals and improve efficiency. However, it also opens a path that malicious hackers on the wider internet could follow to the plant's process control systems.

Vulnerability of process control systems

Operators of process control systems used to believe they were invulnerable to cyber attack for two main reasons. The first reason is the assumption that PCS are isolated from the internet; the second is that PCS generally use proprietary protocols and specialized hardware not compatible with ordinary computers and common network protocols like Ethernet and TCP/IP. These assumptions have led some PCS operators to see the threat of a cyber attack as alarmist. For example, a 2002 article published in *CIO Magazine* outlines the numerous security precautions taken by the Massachusetts Water Resource Authority (MWRA) and concludes that a cyber attack against its PCS would have no effect:

[M]ost public utilities rely on a highly customized Scada system. No two are the same, so hacking them requires specific knowledge -- in this case, knowledge of the MWRA's design and access to that customized software. ... Scada is not networked, except in two places.⁵

He added:

[PLCs] follow the lowest level, most basic instructions (such as turn on and turn off), and report them to Scada ... If something is wrong, the PLC says, "Help me" in the form of an alarm. The alarm sounds at the water site and at the Scada operations centers. The alarm also flashes on the computers, and it can't be shut off until a formal acknowledgement of the alarm is made and physically logged by a human being⁶.

However, many operators have been moving towards open protocols and off-the-shelf hardware to manage their process control systems, even connecting them to the internet—sometimes inadvertently.⁷ These trends have made PCS vulnerable to hackers, often with dangerous results. This fact had been demonstrated even before the MWRA article and has been repeatedly confirmed by penetration testers hired to assess cyber security at critical infrastructures. At the 2006 Black Hat Conference, presenters from IBM Internet Security Systems' X-Force team outlined a penetration test at an unnamed power plant. While meeting with plant management in a conference room, the testing team found a unprotected wireless access point, used it to access the plant's business network, and from there accessed the plant's control network using a ten-year old exploit. In X-Force's experience, only knowledge of common internet protocols was necessary to interfere with PCS systems, but any hacker who wanted to take the extra step to learn about PCS protocols could

5 Scott Berinato. "Debunking the Threat to Water Utilities", *CIO Magazine* (March 15, 2002).
http://www.cio.com/article/30935/Debunking_the_Threat_to_Water_Uilities

6 Ibid.

7 A common cause of an inadvertent connection is a "rogue access point". Employees sometimes set up a wireless network in their office without telling systems administrators. If the access point is not well protected, a hacker can use it to bypass the firewalls and intrusion detection systems that administrators have set up to protect office computers from the wider internet.

find technical specifications online.⁸

Past PCS attacks have even caused physical damage to critical infrastructures. For example, in 2000 a former contractor hacked into the Maroochy Water District's PCS system in Queensland, Australia, and released 80,000 liters of raw sewage into parks, rivers, and even the Hyatt Regency Hotel; the smell drove away local residents, river water turned black, and marine life died as a result.⁹ In March 2007, Idaho National Laboratory conducted a test of the so-called "Aurora vulnerability". This vulnerability would allow an attacker at a remote high voltage circuit breaker to physically destroy a generator by quickly opening and closing the breaker. Details of this vulnerability have been designated "For Official Use Only" by the Department of Homeland Security.¹⁰

Cyber attacks against PCS, whether intentional or unintentional, are likely underreported. No regulation exists requiring power plants to report problems with or attacks against their control systems. In the case of the Aurora vulnerability, ES-ISAC (Electric Sector Information Sharing and Analysis Center) and the Nuclear Energy Institute issued advisories that required no action.¹¹ In April 2009, the North American Electric Reliability Corporation (NERC) issued a letter stating that many power companies were choosing not to identify critical assets in order to avoid complying with cyber security standards, leaving them exposed to such vulnerabilities as Aurora.¹² NERC explains this behavior as a misconception of cyber threats; most operators do not see their own systems as critical to the Bulk Electric System, so they fail to realize that a cyber attack could affect multiple systems at once, and through them the power grid as a whole. In another case, an unnamed power plant suffered a targeted attack and lost process control systems for two weeks. However, since the attack did not disrupt power generation, the attack was not reported to government agencies.¹³

Process control systems at nuclear power plants

The United States has 104 nuclear power plants generating 98,000 megawatts of electricity, roughly 20% of the electricity generated within the US. These plants generally have process control systems, often designed by the same companies that provide these systems to non-nuclear power plants.¹⁴ However, the operators of non-nuclear plants usually have better hardware and cyber security experience than their colleagues at nuclear facilities. Since installation and upgrades of PCS are

8 David Maynor and Robert Graham. "SCADA Security and Terrorism: We're not crying wolf", (paper presented at the Black Hat conference, Las Vegas, Nevada, July 29-August 3, 2006).

<http://www.blackhat.com/presentations/bh-federal-06/BH-Fed-06-Maynor-Graham-up.pdf>

9 Marshall Abrams and Joe Weiss. "Malicious Control System Cyber Security Attack Case Study - Maroochy Water Services, Australia" National Institute of Standards and Technology, Computer Security Resource Center (August 2008).

http://csrc.nist.gov/groups/SMA/fisma/ics/documents/Maroochy-Water-Services-Case-Study_report.pdf

10 Joe Weiss. "One reason why we need regulation", *ControlGlobal.com Unfettered Blog* (December 18, 2008).

<http://community.controlglobal.com/content/one-reason-why-we-need-regulation>

11 Ibid.

12 Michael Assante. "Critical Cyber Asset Identification" (Letter to Industry Stakeholders from the North American Electric Reliability Corporation, April 7, 2009).

<http://www.nerc.com/fileUploads/File/News/CIP-002-Identification-Letter-040709.pdf>

13 Joe Weiss. "Control system cyber events, 60 Minutes, disclosure, and FUD", *ControlGlobal.com Unfettered Blog* (November 13, 2009).

<http://community.controlglobal.com/content/control-system-cyber-events-60-minutes-disclosure-and-fud>

14 Ken Barnes, Briam Johnson, and Reva Nickelson. *Review of Supervisory Control and Data Acquisition (SCADA) Systems*, Idaho National Engineering and Environmental Laboratory, January 2004, page 9.

costly and time-consuming, most non-nuclear PCS operate for eight to fifteen years, the expected lifespan of the hardware used. However, nuclear plants face even higher costs and more stringent safety requirements for their PCS, so they often choose to continue using their original control systems rather than upgrade. A nuclear PCS can be in service for twenty to thirty years, well past the life expectancy of the hardware. Many plants are still using systems based on analog electronics rather than digital.¹⁵ This is confirmed by the experience of nuclear engineer Joe Weiss, now a managing partner of Applied Control Solutions, a consultancy specializing in control system cyber security. Mr. Weiss worked for five years managing a nuclear instrumentation program for the Electric Power Research Institute (EPRI). However, nuclear plants prefer to use tested technologies so Mr. Weiss did not get to do "bleeding edge" research until he managed EPRI's research program for fossil fuel plant instrumentation. This meant that nuclear plants had often adopted modern information technology for their process control systems, but had less experience implementing cyber security on those systems than their colleagues at other electric power plants. This experience gap often led nuclear operators to assume they were less exposed to cyber threats than non-nuclear power plants.¹⁶

In the past five years, US government-funded research into the cyber security of process control systems has focused mainly on oil and gas utilities and the electric grid. While nuclear power plants face many of the same issues in protecting their infrastructure, the key difference is the nuclear reactor. Non-nuclear generators can be completely shutdown, but nuclear reactors run for one to two years once the fuel is installed. Even when the reactor is "shutdown", the fuel still produces decay heat and must be cooled, or the reactor core may melt. The partial meltdown of Three-Mile Island Unit 2 occurred during a reactor shutdown due to operator errors and equipment malfunctions.¹⁷ If such errors and malfunctions can be replicated by a cyber attack, then a reactor meltdown is possible. To determine the danger of this threat, it is necessary to examine cyber incidents that have occurred at nuclear power plants.

Davis-Besse worm infection

On January 25, 2003, at 12:30 AM Eastern Standard Time, the Slammer worm began exploiting a vulnerability in Microsoft SQL Server. Within ten minutes, it had infected 75,000 servers worldwide—90% of vulnerable hosts. The design of Slammer was simple; it did not write itself to the hard drive, delete files, or obtain system control for its author. Instead, it settled in system memory and searched for other hosts to infect. Removing the worm was as simple as rebooting the server after closing network port 1434, Slammer's point of entry. Installing a patch Microsoft had released six months earlier would eliminate the vulnerability Slammer exploited and prevent another infection.

Although Slammer carried no malicious payload, it still caused considerable disruption. It searched for new hosts by scanning random IP addresses. This generated a huge volume of spurious traffic, consuming bandwidth and clogging networks. Slammer's random IP scans disabled data-entry terminals at a 911 call center in Bellevue, Washington (population 680,000), shutdown 13,000 Bank of America ATMs, and forced Continental Airlines to cancel several flights when their online

¹⁵ Ibid, page 23.

¹⁶ Joe Weiss. "Nuclear plant cyber security has a ways to go", *ControlGlobal.com Unfettered Blog*, March 25, 2008. <http://community.controlglobal.com/content/nuclear-plant-cyber-security-has-ways-go>

¹⁷ Ronald L. Krutz. *Securing SCADA Systems*. (Indianapolis: Wiley Publishing, 2006), 29.

ticketing system and kiosks could not process orders.¹⁸ South Korea suffered a nationwide internet outage lasting half a day.¹⁹

The Slammer worm also infected computer systems at the Davis-Besse nuclear power plant near Oak Harbor, Ohio. The worm traveled from a consultant's network, to the corporate network of First Energy Nuclear, the licensee for Davis-Besse, then to the process control network for the plant. The traffic generated by the worm clogged the corporate and control networks. For four hours and fifty minutes, plant personnel could not access the Safety Parameter Display System (SPDS), which shows sensitive data about the reactor core collected from coolant systems, temperature sensors, and radiation detectors—these components would be the first to indicate meltdown conditions. Power plants are required to notify the NRC if an SPDS outage lasts longer than eight hours.

The reactor at Davis-Besse had been offline for nearly a year before its Slammer infection due to the discovery of a hole in the reactor head.²⁰ Although Slammer's scanning traffic did block sensors from providing digital readouts to control systems, it did not affect analog readouts on the equipment itself; plant technicians could still get reliable data from sensors by physically walking up to them and looking at them, though this process is slower than retrieving data over a network.

Davis-Besse had a firewall protecting its corporate network from the wider internet, and its configuration would have prevented a Slammer infection. However, a consultant had created a connection behind the firewall to the consultancy's office network. This allowed Slammer to bypass the firewall and infect First Energy's corporate network. From there, it faced no obstacle on its way to the plant control network. In response, First Energy set up a firewall between the corporate network and the plant control network.

The Davis-Besse incident highlighted the fact that most nuclear power plants, by retrofitting their SCADA systems for remote monitoring from their corporate network, had unknowingly connected their control networks to the internet. At the time, the NRC did not permit remote operation of plant functions.²¹ That policy would change by 2008.

Browns Ferry shutdown

The August 19, 2006, shutdown of Unit 3 at the Browns Ferry nuclear plant near Athens, Alabama, demonstrates that not just computers, but even critical reactor components, could be disrupted and disabled by a cyber attack. Unit 3 was manually shutdown after the failure of both reactor recirculation pumps and the condensate demineralizer controller.²² Without the recirculation pumps, the power plant could not cool the reactor, making a shutdown necessary to avoid melting the reactor core.

18 Robert O. Harrow, Jr. "Internet Worm Uncovers New Holes", *SecurityFocus* (January 29, 2003), <http://www.securityfocus.com/news/2186>

19 Stacy Cowley and Martyn Williams. "Slammer Worm Slaps Net Down, But Not Out" *PCWorld* (January 25, 2003), http://www.pcworld.com/article/108988/slammer_worm_slaps_net_down_but_not_out.html

20 Kevin Poulsen. "Slammer worm crashed Ohio nuclear plant network", *SecurityFocus* (August 19, 2003), <http://www.securityfocus.com/news/6767>

21 Ibid.

22 US Nuclear Regulatory Commission. "Effects of Ethernet-based, non-safety related controls on the safe and continued operation of nuclear power stations" *NRC Information Notice* (April 17, 2007).

The condensate demineralizer is a kind of programmable logic controller (PLC); the recirculation pumps depend on variable frequency drives (VFD) to modulate motor speed. Both kinds of devices have embedded microprocessors that can communicate data over Ethernet, a popular standard for local access networks (LAN). However, both devices are prone to failure in high traffic environments. A device using Ethernet broadcasts data packets to every other device connected to the network. Receiving devices must examine each packet to determine which ones are addressed to them and to ignore those that are not. It appears the Browns Ferry control network produced more traffic than the PLC and VFD controllers could handle; it is also possible that the PLC malfunctioned and flooded the Ethernet with spurious traffic, disabling the VFD controllers; tests conducted after the incident were inconclusive.

The failure of these controllers was not the result of a cyber attack. However, it demonstrates the effect that one component can have on an entire PCS network and every device on that network. Combined with the Davis-Besse worm infection, the Browns Ferry shutdown presents a possible attack scenario. If a worm like Slammer had infected the control network of an active plant and attempted to spread not only through UDP, but also through Ethernet, it could have disabled the recirculation pumps as well as the sensors that would alert plant personnel to the problem.

Hatch automatic shutdown

Due to the growing network connections between control systems and office computers, even seemingly simple actions can have unexpected results. On March 7, 2008, Unit 2 of the Hatch nuclear power plant near Baxley, Georgia, automatically shutdown after an engineer applied a software update to a single computer on the plant's business network. The computer was used to collect diagnostic data from the process control network; the update was designed to synchronize data on both networks. When the engineer rebooted the computer, the synchronization program reset the data on the control network. The control systems interpreted the reset as a sudden drop in the reactor's water reservoirs and initiated an automatic shutdown.²³

This innocent mistake demonstrates how malicious hackers could make simple changes to a business network that end up affecting a nuclear reactor—even if they have no intent to interfere with critical systems. This incident is probably the least critical of those examined so far, since it *activated* safety systems rather than disrupting them. However, it also demonstrates that plant operators do not fully understand the dependencies between network devices. This would make it difficult to identify and protect all the vulnerabilities in a process control system.

Stuxnet: a proof of concept

The Stuxnet attack against the Iranian nuclear program demonstrates the impact that a sophisticated adversary with a detailed knowledge of process control systems can have on critical infrastructures. Stuxnet is believed to have destroyed 984 centrifuges at Iran's uranium enrichment facility in Natanz.²⁴ An analysis of the event by the Institute for Science and International Security (ISIS),

23 Brain Krebs, "Cyber Incident Blamed for Nuclear Power Plant Shutdown" *Washington Post*, June 5, 2008.

<http://www.washingtonpost.com/wp-dyn/content/article/2008/06/05/AR2008060501958.html>

24 William J. Broad, John Markoff, and David E. Sanger. "Israeli Test on Worm Called Crucial in Iranian Nuclear Delay". *New York Times*, January 15, 2011.

based on open source technical data about the Stuxnet computer worm and the Iranian nuclear program, found that Stuxnet may have been designed specifically for that purpose. However, Stuxnet also demonstrates the limitations that even such a sophisticated adversary would face in launching an attack against process control systems. The ISIS report finds that the Stuxnet attack, though it successfully disrupted the Iranian centrifuge program, did not slow down Iran's accumulation of low-enriched uranium.²⁵ The attack is remarkable for its sophistication, but it did not pose an epic threat to Iran.

However, that sophistication must be considered when assessing the vulnerability of nuclear facilities to cyber attack. The Stuxnet worm targeted specific PCS components used in the Iranian centrifuge cascades: a frequency converter manufactured by Iranian firm Fararo Paya, another frequency converter manufactured by Finland's Vacon,²⁶ and the S7-315 and S7-417 programmable logic controllers made by Siemens.²⁷ The PLCs controlled the frequency converters to modulate the speed at which the centrifuges spun. Stuxnet commanded the PLCs to speed up and slow down the spinning centrifuges, destroying some of them, while sending false data to plant operators to make it appear the centrifuges were behaving normally. The *New York Times* report suggests that Stuxnet's authors may have learned about vulnerabilities in the Siemens controllers thanks to a partnership between Siemens and the Idaho National Laboratory aimed at assessing vulnerabilities in such components. These products are general PCS components not unique to the Iranian nuclear program; Siemens reports that at least 24 of its customers were infected by Stuxnet, though they suffered no damage.²⁸

The reason Stuxnet did not disrupt every vulnerable PCS it infected is that it was programmed to disrupt only systems that had the same configuration as the centrifuge cascade used at Natanz.²⁹ Antivirus company Symantec began detecting Stuxnet traffic in June 2009, mostly in Iran, but also in neighboring countries. However, since it did not spread aggressively and did not damage the systems it had infected, it raised little alarm.³⁰ Only at the Natanz enrichment facility did it have a major effect. Experts cited by the *New York Times* report suggest that Israeli intelligence provided the specific technical details necessary for Stuxnet to limit its damage to the Iranian nuclear program.

While the *New York Times* article only presents a possible scenario, that scenario and the evidence reflect the challenges of executing a catastrophic cyber attack against a nuclear facility. Programming is a cyclical process of trial and error. For an amateur hacker working only with a computer, the costs of testing software are trivial. Testing software designed for process control systems, however, requires access to the system in question, which is usually expensive. Malicious hackers could run tests on a remote PCS they had compromised, but an unsuccessful test could raise alarms or damage the system before the hackers were ready for the next stage of an attack. The Stuxnet authors would need a dedicated testbed to refine their code. Stuxnet also incorporated technical information specific to the Iranian facility. These resources are out of the reach of amateurs and would require

25 David Albright, Paul Brannan, and Christina Walrond. "Stuxnet Malware and Natanz: Update of ISIS December 22, 2010 Report". ISIS Report, February 15, 2011, pg 2.

26 David Albright, Paul Brannan, and Christina Walrond. "Did Stuxnet Take Out 1,000 Centrifuges at the Natanz Enrichment Plant? Preliminary Assessment". ISIS Report, December 22, 2010, pg 3.

27 Albright, Brannan, and Walrond. "Stuxnet Malware", pg 1.

28 "SIMTAC WinCC / SIMTAC PCS-7: Information about Malware / Viruses / Trojan horses". Siemens, accessed April 14, 2011.
<http://support.automation.siemens.com/WW/llisapi.dll?query=stuxnet&func=cslib.cssearch&content=adsearch%2Fadsearch.aspx&lang=en&siteid=cseus&objaction=cssearch&searchinprim=0&nodeid0=10805583>

29 Albright, Brannan, and Walrond. "Stuxnet Malware", pg 1.

30 Broad, Markoff, and Sanger. "Israeli Test".

the kind of funding and actionable intelligence that comes from state sponsorship.

The Stuxnet attack also incorporates elements of the other three incidents examined in this paper. First, it disrupted the systems that monitored physical components, like the Davis-Besse worm infection. Second, it interfered with programmable logic controllers, like the Browns Ferry data storm. Third, it relied on there being some path from ordinary office computer to process control systems, as in the Hatch automatic shutdown. At the same time, the Stuxnet authors innovated on these features: Stuxnet did not simply disrupt sensor output, it faked it; it did not simply interfere with PLCs, it gave them specific instructions; finally, it did not rely on an internet connection to Natanz—it also traveled between computers on worker’s thumb drives³¹ and infected components destined for Natanz at their source in the Iranian chain of supply.³²

Skeptics and alarmists can both use the Stuxnet attack to justify their positions. Alarmists can point to the vulnerability of PCS and its direct effect on Iranian national interests. However, skeptics can argue that the Stuxnet attack required specific knowledge of a particular facility and cannot be generalized to other systems, the same argument used by the Massachusetts Water Resource Authority. Further, the impact could hardly be described as catastrophic. However, it is important to look at the Stuxnet attack in the context of history. Cyber attacks have evolved from the work of amateurs and professional criminals into a serious endeavor for states engaged in international disputes. States have begun to use cyber attacks not just to gather intelligence or control information networks, but to damage physical infrastructures. While the damage is nowhere near a “digital Pearl Harbor”, the trend is clear: states are actively pursuing cyber attacks as an instrument of foreign policy while advancing the technical know-how such attacks require.

Lessons

These four incidents hold important lessons for the cyber security of nuclear facilities and critical infrastructures in general. First, skeptics claim that PCS are immune from attack since they are not connected to the internet. However, the Davis-Besse incident shows that this is a misconception; even operators who try to monitor and protect every connection cannot be sure they know about all of them. Stuxnet even traveled on portable thumb drives to infect computers that were not connected to the internet. Second, skeptics argue that PCS are immune from attack since they are different from ordinary computers. However, all four incidents demonstrate that PCS have become interoperable with ordinary computers, making them vulnerable. Third, vulnerabilities are more complicated than both skeptics and alarmists realize. Alarmists often invoke the danger of hackers taking control of a power plant, but these incidents show how unintelligent computer viruses and even malfunctions in small devices can have big unexpected effects. This suggests that even though nuclear facilities are vulnerable to attack, a malicious hacker would have difficulty making sure an attack works precisely as planned. Even so, states are working make cyber attacks more precise, supplementing their methods with intelligence from other sources.

Cyber security and nuclear safety regulations

As states take a greater interest in launching cyber attacks against nuclear facilities, they should also

31 Albright, Brannan, and Walrond. “Did Stuxnet Take Out 1,000 Centrifuges?” pg 7.

32 Albright, Brannan, and Walrond. “Stuxnet Malware”, pg 2.

take a greater interest in protecting their own facilities against attack. This means translating the lessons of previous incidents into workable guidance and regulation for plant operators. So far, this has been lacking, both from the United States government and the International Atomic Energy Agency (IAEA). The nuclear industry does not have the expertise to handle such threats on its own, as evidenced not only by the incidents covered here, but also by the lack of compliance with NERC critical asset identification standards.³³

However, the agencies charged with providing the necessary guidance may not have that expertise themselves. The U.S. Nuclear Regulatory Commission (NRC) did not issue an Information Notice after the Hatch shutdown as it had for the Davis-Besse and Browns Ferry incidents. The NRC is aware of its expertise gap and is actively addressing it. In January 2008, the NRC's newly established Computer Security Office launched a working group to develop an Information Security Strategic Plan (ISSP) for 2010 to 2015. The working group found that cyber security issues at nuclear plants were handled in an "ad hoc" manner, since the NRC's staff with cyber experience were both limited and widely dispersed about the country. The NRC set up an Information Security Steering Committee to coordinate the activities of these dispersed experts under the ISSP, including the development of new rules and regulatory guidance for cyber security at nuclear facilities. Part of that process will be implementing a 2008 recommendation from the Office of the Inspector General to develop a program of cyber security inspections at nuclear power plants.³⁴ The ISSP outlines plans to use the NRC's licensing and inspection authority to enforce cyber security standards at nuclear facilities,³⁵ however, it is too early to judge the effectiveness of these efforts.

While the IAEA lacks the enforcement powers of the NRC, it still has an important role to play as inspector and advisor to the nuclear programs of other nations. However, it seems to be a bit slower than the NRC in developing its cyber security expertise. Its most recent technical guidance on the matter seems to be "Security of Information and Instrumentation & Control Systems at Nuclear Facilities" released in 2007. However, this guidance fails to account for documented PCS incidents in both nuclear and non-nuclear facilities and the reported experience of penetration testers. For example, the guidance states that cyber security at nuclear facilities can be achieved using the same methods and tools developed for IT security.³⁶ However, the Browns Ferry data storm was created by either a failed PCS component or normal network operations; IT security would not have predicted the resulting failure of the reactor pump VFDs. Since then, Stuxnet has further demonstrated the inadequacy of basic IT security, since it infected PCS components in the Iranian supply chain rather than looking for a direct network connection to Natanz. The guidance also recommends developing a network diagram documenting all external connections, however, the assumption that all external connections were known and controlled was the basis for the supposed invulnerability of PCS. Even in the IT world, penetration testers have found that network diagrams are often grossly inaccurate and only create a false sense of security. While the IAEA guidance does give some sound advice for basic cyber security, it does not begin to address the unique challenges presented by PCS. The IAEA is continuing to develop its expertise in this area, especially since the Stuxnet attack, however, the current state of official guidance and regulation suggests that those responsible for protecting nuclear facilities from cyber attack are less prepared than their potential

33 Assante. "Critical Cyber Asset Identification".

34 Stephen D. Dingbaum. "NRC's Planned Cyber security Program (OIG-08-A-06)". Memorandum Report from the Office of the Inspector General (March 18, 2008).

35 Nuclear Regulatory Commission, "Information Security Strategic Plan". May 18, 2009.
<http://www.nrc.gov/reading-rm/doc-collections/commission/secys/2009/secy2009-0077/enclosure.pdf>

36 International Atomic Energy Agency. "Security of Information and Instrumentation & Control Systems at Nuclear Facilities", *IAEA Nuclear Security Series No. XX Technical Guidance*. 2007, page 13.

aggressors.

Conclusion: A mixed bag

While some cyber security incidents have occurred at nuclear power plants, crossing the imaginary boundary between IT and PCS and shutting down reactors, so far the potential for damaging a nuclear reactor appears theoretical. Scott Lunsford, a penetration tester for IBM, says government mandated safeguards would prevent a hacker from triggering a meltdown. So far, no catastrophic damage has resulted from a cyber attack against a nuclear facility. The same cannot be said for other sectors, as in the case of the Maroochy water incident, and the Stuxnet attack has demonstrated that states are likely pushing the development of new tactics and capabilities in cyberspace.

Although the experience of the nuclear sector lags behind that of non-nuclear facilities in cyber security and PCS, nuclear plants must also comply with stronger safety regulations and inspections. Although the NRC's cyber regulations are still being developed, its existing regulations have put several incidents on the public record that would have gone unreported by non-nuclear power plants. This parallels the trend of cyber security in e-commerce. In the early 2000's, banks and online merchants commonly suffered cyber attacks that potentially revealed their customers private data to hackers. To protect their reputations, they hired consultants to quietly fix their systems under a non-disclosure agreement. Eventually, California passed SB1386, requiring any company that did business in the state of California to notify their customers if a hacker could have potentially accessed their private data. After the law went into effect in July 2003, the extent of the hacking became public knowledge and companies began to invest in cyber security to reassure their customers before they suffered an attack. Oil, gas, and electric companies have been active in protecting their PCS from cyber attack, however, they still have little incentive to report the attacks they suffer. No regulation requires it, and companies fear their information could be made public under the Freedom of Information Act if they do. The years from 2010 to 2015 could prove decisive in the field of PCS security. If the NRC can implement the same sort of rigorous inspection and reporting requirements for cyber security as they have for physical security and safety, it may open the field up to greater public scrutiny and spur the investment needed to better protect critical infrastructures.

About the Author

Brent Kesler graduated from Dartmouth College in 2003 with a degree in computer science. Until 2006 he wrote *Security in the News*, a daily report of developments related to malicious hacking, malware, cyber security best practices, and homeland security. He later worked as research coordinator for the Institute for Information Infrastructure Protection (I3P), helping manage federal research projects related to critical infrastructure protection. In 2010, Mr. Kesler graduated from the Monterey Institute of International Studies, building on his cyber security background with a master's degree in international policy and a focus on terrorism.