PROTECTING ACCELERATOR CONTROL SYSTEMS IN THE FACE OF SOPHISTICATED CYBER ATTACKS[*]

S. M. Hartman[†], Spallation Neutron Source,
Oak Ridge National Laboratory, Oak Ridge, TN 37831, USA

Abstract

Cyber security for industrial control systems has received significant attention in the past two years. The news coverage of the Stuxnet attack, believed to be targeted at the control system for a uranium enrichment plant, brought the issue to the attention of news media and policy makers. This has led to increased scrutiny of control systems for critical infrastructure such as power generation and distribution, and industrial systems such as chemical plants and petroleum refineries. The past two years have also seen targeted network attacks aimed at corporate and government entities including US Department of Energy National Laboratories. Both of these developments have potential repercussions for the control systems of particle accelerators. The need to balance risks from potential attacks with the operational needs of an accelerator present a unique challenge for the system architecture and access model.

THE ERA OF STUXNET

In the summer of 2010, reports appeared in the media describing a new piece of computer malware infecting Windows computers. The malware was named Stuxnet. As computer security companies began analysis of this software, several interesting findings emerged. The software was large and much more complex than typically seen, and likely required significant financial support and technical expertise to develop. Stuxnet included four previously unknown exploits of the Windows operating system, while it is rare to see more than one used per malware package. It leveraged compromised digital certificates to install driver files, implying the developers or their cohorts had physical access to the certificate owner's facility. Stuxnet could spread or copy itself to infect other systems through multiple pathways, including local networks and removable storage media. It had an update mechanism and a command and control interface. And, most interestingly, Stuxnet was directed at industrial control systems, with the ability to reprogram programmable logic controllers (PLCs) and hide those modifications from the control system operator or engineers.

From a detailed analysis from Symantec [1], Stuxnet targeted Windows computers running Step 7, the programming software for a family of Siemens PLCs. Stuxnet installed itself into a Step 7 project (which also enables further infections if that project is copied to a different computer system) and modified the communication libraries used between the programming software and the remote PLC. After infecting a Step 7 installation, Stuxnet was designed to then look for a PLC with a specific configuration of field devices. If such a PLC were found, the runtime code on the PLC would be modified. This modification altered the set point of a frequency drive over a period of time. The modifications in the Step 7 communications libraries effectively hid the modifications to the PLC from the user interface layer. The set point appeared to be at the correct value for operators monitoring the system. Additionally, viewing the runtime PLC code displayed the unmodified code, hiding the compromise from system engineers.

The sophistication of the software used for this attack attracted considerable attention in the computer security arena. Additional details emerged, providing evidence that this attack was directed at a uranium enrichment facility in Iran. The frequency drive that Stuxnet targeted is believed to be the controller for centrifuges used in the enrichment process. The impact from the modification of the PLC code resulted in the destruction of an estimated 1000 centrifuges, slowing the enrichment program and altering geo-politics [2].

The implication that one or more nation states were potentially behind the creation and usage of Stuxnet brought further media coverage to Stuxnet [3]. Reports that Stuxnet marked the beginning of an era of “cyber-warfare” brought a level of awareness well beyond that typically received by the latest attack on Windows PCs.

While Stuxnet is often presented as the first use of malicious software for covert international political means, there have actually been several documented cases prior to Stuxnet. Perhaps the most dramatic, although without the fanfare of Stuxnet, was a 1982 attack orchestrated by the US Central Intelligence Agency against a Soviet gas pipeline. Details of the attack were not made public until the publication of a book and media reports in 2004 [4]. Reports describe a covert operation to plant a Trojan Horse into the computer control system for the automation of a trans-Siberian gas pipeline under construction. When the pipeline began operation, the embedded malware manipulated the pumps and valves to over-pressurize the system, resulting in what has been described as the greatest non-nuclear explosion and fire ever seen from space. The financial impact (the pipeline was expected to generate revenue of $8 billion per year) has been cited as contributor to the collapse of the Soviet Union [5].

[*] ORNL/SNS is managed by UT-Battelle, LLC, under contract DE-AC05-00OR22725 for the U. S. Department of Energy.

[†] hartmansm@ornl.gov

While these two particular attacks on control systems are among the most dramatic, there are likely many more attacks which are either not reported or don't receive the same level of attention. There have been reports of electrical blackouts from cyber attacks [6] and other attacks on critical infrastructure. But there have also been misattributions of system failures to cyber attacks. In November of 2011, widespread reports appeared describing the first foreign cyberattack on a public utility in the US after a failure at a water plant in Illinois. But further investigation concluded no evidence of malicious activity and no evidence of a computer intrusion [7].

ADVANCED PERSISTENT THREATS

Recent history has also brought increased awareness of cyber security to US National Laboratory operations. In April of 2011, Oak Ridge National Laboratory (ORNL) was the target of a so-called Advanced Persistent Threat [8]. An email message, purporting to be an ORNL benefits announcement, was sent to approximately 500 ORNL employees [9]. Of these employees, about 10 percent clicked on a link, downloading an executable file that could exploit an unpatched vulnerability in certain Windows computer systems. One employee had sufficient privileges on that and other ORNL computers to allow the malware to spread to other systems. Four days later, the intrusion was detected but the decision was made to monitor the systems to determine the nature of the attack. On April 15, just over a week after the initial email, ORNL elected to disconnect its network from the internet to protect against any exfiltration of data from ORNL systems [10]. ORNL remained isolated from the internet until May 1, with additional restrictions on email remaining in place for more than a month.

Over the following months, several other National Laboratories including Pacific Northwest National Laboratory [11] and Jefferson Laboratory [12] were also impacted by cyber events. In each case, the Laboratories elected to disconnect their networks from the internet to limit further damage or unintended release of data while cleaning and rebuilding systems in response to the attack.

Such attacks have been dubbed Advanced Persistent Threats (APT). An APT is a targeted attack rather then an opportunistic one, and is carried out by what is believed to be a sophisticated attacker. The approach is to gain access to a network, use this foothold to gather information on the network topology and security, and then leverage this information to gain additional access. The focus of such attacks is generally thought to be access to proprietary data or intellectual property.

Reports emerged in late summer of 2011 [13] detailing a widespread instance of an organized APT aimed at a broad range of targets including governments, technology corporations, defense contractors, and non-governmental organizations. The intrusions were dubbed "Operation Shady RAT," with RAT standing for "remote access tool." A white paper [14] by the computer security company McAfee documents the scope of the attacks based on logs McAfee attained from a command and control computer used in the attacks. Details in the report indicate at least one (unnamed) Department of Energy National Laboratory was a victim, with logs showing intrusions going back to at least July 2006 and lasting for at least three months. The report documents at least seventy other victims, with intrusion durations lasting more than two years in some cases.

Operation Shady RAT is just one example of such an organized, systematic attack used to gain access to data. These attacks mark a shift from the opportunistic attacks that have been the focus of much of cyber security in past years, or the attacks aimed at direct financial gain through individual credit card or banking information.

ACCELERATOR CONTROL SYSTEMS

The most common cyber risks for an accelerator control system to date have been those from more routine computer malware infections rather then something specific to an accelerator system. However, the visibility of large scale accelerator projects may invite unwanted attention. During commissioning of the Large Hadron Collider (LHC) at CERN for instance, publicity seekers defaced a web server associated with the Compact Muon Solenoid experiment [15]. The web server attacked was not connected to the accelerator control system and did not have any impact on LHC operations. However, this did result in a lot of unwanted attention for the laboratory.

A Stuxnet-type attack or an APT attack are unlikely to be directed at the computer control system for a particle accelerator. However, these control systems are not completely removed from these attacks and the fallout from such an attack can be of real concern.

Industrial Control Systems

Current accelerator control systems make use of many of the same technologies used by industrial Supervisory Control and Data Acquisition (SCADA) systems. The use of custom hardware has generally been superseded by commercial-off-the-shelf (COTS) products including industrial PLCs, network attached embedded devices and networked measurement equipment (oscilloscopes, spectrum analyzers, etc). Accelerator control system toolkits such as Experimental Physics and Industrial Control System (EPICS) are well integrated with these COTS devices. Control system workstations and servers generally run common operating systems such as Linux or Windows. Relational Databases provide the backend for system configuration and runtime logs. With the use of these common technologies, accelerator control systems inherit some of the same risks and threats faced by industrial SCADA systems and office computer systems [16].

The Siemens PLCs targeted by Stuxnet are used by a number of accelerator control systems. Although it is highly unlikely that the effort employed in the initial Stuxnet attack would ever be deployed against an accelerator control system, now that the code is available, there is the potential for it to be used by others for more generic attacks. Since Stuxnet, a number of demonstrations of potential attacks using the capabilities developed for Stuxnet have been shown. These attacks could target industrial control systems, critical infrastructure and even the locks in a prison [17].

For an accelerator control system, with its complexities and need for ongoing evolution, the ability to interface to a wide variety of COTS devices and industrial control devices is critical. At the Spallation Neutron Source (SNS) at ORNL, a significant amount of the field input/output devices are industrial PLCs from Allen Bradley. An open communication protocol [18] is used to interface to this commercial system from the EPICS-based accelerator control system. Many networked industrial control system devices are relatively unsecured from a cyber security perspective. While the protocols may be open standards or proprietary, the devices themselves generally implement limited security. Effectively, any client who has a network path to the device and an implementation of the protocol can interact with the device and, perhaps, control its outputs. This provides a great convenience in developing a complex, heterogenous control system like that used for an accelerator. Devices from a variety of vendors using various protocols can be integrated into a common control system framework. However, this accessibility does imply a level of risk. If a control system device driver can be used to manipulate the device over the control network, other applications on the network – or with remote access to that network – can do the same. Securing what is installed on the network (authorized or unauthorized) and access to the network (physical or remote) is critical in ensuring the integrity of the control system.

Laboratory Networks

With evidence of National Laboratory networks being the focus of directed attacks, control system networks may be vulnerable by association. Intellectual property of the type being sought through APT-type attacks is not likely to be found on an accelerator controls network. Information about machine design and operational parameters is generally openly shared through conference presentations and other means. Little or no proprietary information is kept on such networks. However, there is the potential for inadvertent fallout from an attack on the laboratory network interfering with the operations of an accelerator system.

In the case of the attack on ORNL in April 2011, steps were taken to physically isolate the accelerator control system network from the ORNL network shortly after the intrusion was first detected. The accelerator control system network is generally isolated from the ORNL network with a single point connection through a firewall. The firewall is restrictive enough to protect against most likely network-based threats and provides an easy means of complete isolation if needed. This option for complete isolation has been exercised on several occasions and can be used to protect the control system network from ORNL systems or ORNL systems from the control system network in the event of a problem on either side.

Segmentation of networks proved to be of value during the April incident at ORNL. Although central laboratory systems, including web services and mail, were greatly impacted, organizations within ORNL which were segmented from the main network experienced less disruption and were more able to continue business as usual, although without access to the outside world. While ORNL's network was still significantly impaired, the SNS, with its accelerator control network isolated, operated as normal, with a 1 MW beam on target and all accelerator, machine protection and safety systems operational.

Remote Access

The SNS functionality which was impaired during the April ORNL incident was that of remote access. With a complete separation between the accelerator control system network and the ORNL network, it was only possible to monitor accelerator systems from within the central control room. Tools to provide read-only access for monitoring the control system were no longer functional. The ability to remotely access the network for trouble shooting or repair by system experts was also lost. For the several weeks of operation in this configuration, this loss of access was mostly an inconvenience. However, for longer term operation of a user facility accelerator, some level of remote access is needed by system experts to ensure reliable operation and efficient recovery from downtime events.

The inconvenience and operational impact of isolating the control system network is proportional to the duration of the isolation. The complete network separation made it difficult to keep development systems synchronized with the production system. The use of offline tools for system analysis was also impacted. Having the network separated also increased the desire to use removable media as a mechanism for moving data in to or out of the production accelerator environment. Such use brings with it the increased likelihood of inadvertently introducing malware through infected data storage media.

Risks

In considering the overall cyber security risks to an accelerator control system, the most likely problems are those that have been inherited from desktop computer world. These include malware spread by interconnected networks, infected removable media, and errors by the system experts who have control over the network and networked computer equipment. Protection relies on the same fundamentals required for protecting business networks: defense in depth through network segregation, patching, access control restrictions, system design, and user training [19].

However, for an accelerator control system, operational schedules and availability needs typically result in software patches not being applied as quickly as they might be in an office setting. Additionally, digital oscilloscopes and other test and measurement equipment are frequently built using a desktop operating system. Patches for these systems can be delayed over concern for the patch impacting operation of the device. Moreover, these devices frequently fall out of normal patch cycles since they are not centrally managed and may come and go from the network as need arises. Effectively, this means that it is likely that there are vulnerable systems on any accelerator control system network. Therefore, it is important that system design and procedures for access to the network are appropriate to limit risk.

The greatest potential for problems is from remote access. While it is necessary to provide remote access for monitoring, troubleshooting and system expert access, such access must be designed in a way to keep risks to an acceptable level. Providing an easy means for read-only remote access, through a gateway for instance, can be an effective means of limiting need for directly accessing the network remotely. Direct login access to the network needs to be available for system experts, but should be inconvenient by design to limit unnecessary use. One-time passwords, two factor authentication, or multiple hops to access the network are effective means of limiting risk from remote access. Ensuring a single point of contact from the restricted network to external networks is key in avoiding backdoors. This provides a single point for monitoring, a mechanism for complete separation if needed, and a single configuration point to maintain access rules. Inadvertent backdoors to control system networks, bypassing system firewalls, have led to malware infections in electrical and nuclear power plant networks [20].

For a well designed network with appropriate perimeter controls, the next threat is likely malware infection through the inadvertent actions of accelerator staff. Actions such as using a laptop on an unsecured network and then on the controls network, or careless use of memory sticks or other removable media provide an easy means for introducing computer malware. Staff training and policy are the best protections since engineering controls to protect against this threat are limited.

IMPLICATIONS FOR ACCELERATOR PROJECTS

Stuxnet was designed for a very particular target, likely a uranium enrichment plant centrifuge controller. It also required substantial resources to develop and deploy. As such, it is a unique risk of a level unlikely to be directed at an accelerator control system. However, it did raise the awareness of potential risk to industrial control systems and provided tools that can potentially be used by less sophisticated attackers on other systems. This introduces another potential threat vector that may impact controls systems for accelerators. Directed APT-type attacks aimed at gathering intellectual property from National Laboratory networks are also unlikely to be directed at the control system network of an accelerator. However, the consequences from these attacks can indirectly interfere with accelerator systems and operations.

The broader impact from both of these developments has been an increased awareness of cyber risks. Project sponsors for operational and planned accelerators are going to have increased expectations for mitigations of cyber risks. Any intrusion into an operational accelerator control system will likely draw close scrutiny and can impact the sponsor's evaluation of the project. There would also be a direct impact to the operational schedule, with clean-up and reconfiguration potentially resulting in extended downtime after any such incident.

But while some of this risk has been inherited from the use of some of the same technologies used in business information technology (IT) systems, the management of these risks cannot have an identical solution to that used for enterprise IT systems. The operational and functional requirements of an accelerator control system are very different from that of an IT system. Issues of system stability and performance, the complexity of interactions across heterogeneous systems, availability requirements, and the long operational lifetime of a control system all differentiate a control system from an enterprise IT system.

It is critical that accelerator system experts are actively engaged in the system design and security configuration rather then leaving this solely to IT professionals. It is also important that the risk mitigation approaches are in line with the level of risk faced, while not unnecessarily impacting the scientific, engineering and operational needs of the accelerator.

REFERENCES

[1] N. Falliere, L. O. Murchu and E. Chien, "W32.Stuxnet Dossier," Symantec Security Response White Paper, Version 1.4, February 2011.

[2] E. MacAskill, "Stuxnet cyberworm heads off US strike on Iran," *The Guardian,* January 16, 2011.

[3] J. Markoff, "A Code for Chaos," *The New York Times,* October 2, 2010, page WK5.

[4] A. Russell, "CIA plot led to huge blast in Siberian gas pipeline," *The Telegraph (UK),* February 28, 2004.

[5] W. Safire, "Oh, that big 1982 Siberian Explosion?" *The New York Times,* February 4, 2004.

[6] T. Claburn, "CIA Admits Cyberattacks Blacked Out Cities," *InformationWeek,* January 18, 2008.

[7] E. Nakashima, "Water-pump failure in Illinois wasn't cyber-attack after all," *The Washington Post,* November 25, 2011.

[8] F. Munger, "Cyber attack forces ORNL to shut down Internet access; experts probing Advanced Persistent Threat," blogs.knoxnews.com/munger/2011/04/cyber-attack-forces-ornl-to-sh.html.

[9] K. Zetter, "Everyone Has Been Hacked. Now What?" *Wired.com,* May 4, 2012.

[10] J. Vijayan, "Oak Ridge National Lab shuts down Internet, email after cyberattack," *Computerworld,* April 19, 2011.

[11] F. Y. Rashid, “DOE Lab Shuts Down Email, Web Access After Sophisticated Cyber-Attack,” *eWeek.com,* July 7, 2011.

[12] W. Jackson, “Energy lab restoring website, investigating attack,” *Government Computing News,* July 7, 2011.

[13] E. Nakashima, “Report on ‘Operation Shady RAT’ identifies widespread cyber-spying,” *Washington Post,* August 2, 2011.

[14] D. Alperovitch, “Revealed: Operation Shady RAT,” McAfee White Paper, Version 1.1.

[15] L. Greenemeier, “Hackers attack Large Hadron Collider computers to prove they’re vulnerable,” *Scientific American,* September 12, 2008.

[16] S. Lüders, “Stuxnet and the Impact on Accelerator Control Systems,” Proceedings of ICALEPCS2011, Grenoble, France, October 2011, p. 1285-1288.

[17] K. Zetter, “Researchers Say Vulnerabilities Could Let hackers Spring Prisoners From Cells,” *Wired.com,* July 29, 2011.

[18] K. U. Kasemir, L. R. Dalesio, “Interfacing the ControlLogix PLC over Ethernet/IP,” Proceedings of ICALEPCS2001, San Jose, California, November 2001.

[19] S. Lüders, “Securing Control Systems Against Cyber Attacks,” Proceedings of PAC2009, Vancouver, BC, Canada, p. 1785-1789.

[20] K. Poulsen, “Slammer worm crashed Ohio nuke plant network,” *Security Focus,* August 19, 2003.