SECURING CONTROL SYSTEMS AGAINST CYBER ATTACKS

S. Lüders, CERN, Geneva, Switzerland

Abstract

Modern accelerator control systems are increasingly based on commercial-off-the-shelf products (VME crates, Programmable Logic Controllers (PLCs), Supervisory Control And Data Acquisition (SCADA) systems, etc.), on Windows or Linux PCs, and on communication infrastructures using Ethernet and the TCP/IP protocol. Despite the benefits coming with this (r)evolution, these "modern" control systems and infrastructures usually completely lack adequate levels of robustness, resilience and security. Even worse, new threats are inherited, too: Worms and viruses spread within seconds via the Ethernet cable, and attackers are becoming interested in breaking into control systems. This paper will discuss the initial security risks, what precautions are needed to protect control systems against cyber threats and how to provide a secure environment without sacrificing operability.

INTRODUCTION

The enormous growth of the worldwide interconnectivity of computing devices (the "Internet") during the last decades offers computer users new means to share and distribute information and data. In industry, this results in an adoption of modern Information Technologies (IT) in their plants and, subsequently, in an increasing integration of the production facilities, i.e. their process control and automation systems, and the data warehouses. Thus, information from the factory floor is now directly available at the management level ("From Shop-Floor to Top-Floor") and can be manipulated from there.

Today's accelerator control systems do not differ significantly from the control systems* used in industry. Modern IT is increasingly used, hardware for accelerator control system is nowadays generally based on common-of-the-shelf devices (VME crates, PLCs, Ethernet connected power supplies and fan trays, Ethernet-to-RS232 gateways, oscilloscopes), standard operation systems (Microsoft's Windows or Linux derivatives, VxWorks or LynxOS), standard software applications used for SCADA systems (like PVSS, EPICS, WINCC, Wonderware, Labview to name but a few), standard programming languages and middle ware (e.g. JAVA, C++, AJAX, XML, CORBA, OPC) as well as standard Oracle or MySQL data bases.

However, different to industry due to the academic freedom in the High Energy Physics (HEP) community, accelerator control systems are produced by a wide, decentralized community. Consequently, this leads to many different, heterogeneous systems which often depend on each other and, thus, necessitate open interfaces for inter-communication. Furthermore, the decentralized development very often requires remote access for the corresponding experts for future expert interventions like applying bug fixes.

Due to this trend, the risk of suffering from a security breach also increases: With the thorough inter-connection of campus and controls networks, the adoption of modern IT standards, and the usage of standard IT components, accelerator control systems are also exposed to the inherent vulnerabilities of the corresponding hardware and software. This security risk can be expressed as in the following intuitive illustrative formula:

$$Risk = Threat \times Vulnerability \times Consequence$$

These different factors are explained in the following and examples are given.

Threats

The inter-connection of campus and controls networks exposes the attached control systems to the hostile external world unless there is an "air gap" between both – an assumption that is rarely true. The number of potential "threats" increases as worms and viruses can now easily propagate to control systems and attackers start to become interested in control systems, too. It is a fact that the corresponding attack procedures were already presented and discusses on the usual "black-hat"-conferences.

Just recently, the Wall Street Journal reported that "cyber-spies have penetrated the U.S. electrical grid and left behind software programs that could be used to disrupt the system" [1]. Thus, the fear of cyber-attacks by terrorist on control systems is driving governments world-wide to act ("America's failure to protect cyberspace is one of the most urgent national security problems facing the country" [2]).

There is no argument why attackers would ignore HEP systems just because it's HEP†. For example, on the day of the LHC start-up in September 2009, Greek activists have deliberately tried to break into a web server hosted at one of the LHC experiments and defaced one webpage.

Additional to these external threats by attackers are internal threats like operators or engineers, who unintentionally download configuration data to the wrong

* Throughout this paper, the term "control system" commonly denotes all controls-related systems like distributed process control systems, automation systems, SCADA systems, safety systems, etc. A "system expert" has the expertise in its configuration.

† Some more examples of incidents involving control systems can be found here: The Register, 2000, http://www.theregister.co.uk/2000/04/27russia_welcomes_hack_attacks; Computer Crime Research Centre, 2005, http://www.crime-research.org/analytics/1718; Security Focus, 2005, http://www.securityfocus.com/news/6767; eWeek.com, 2005, http://www.eweek.com/article2/0,1759,1849914,00.asp; Security Focus, 2006, http://www.securityfocus.com/news/11465; CSO online, 2007, http://www2.csoonline.com/exclusives/column.html?CID=32893; CNN online, 2007, http://edition.cnn.com/2007/US/09/26/power.at.risk/index.html

device, or broken controls devices that flood the controls network and, thus, bring it to a halt.

Nevertheless, the major part of the factor "threat" originates from outside and cannot be significantly reduced. Thus, protective measures have to be implemented to prevent external threats penetrating control systems. These protective measures should also prevent insiders from (deliberate or accidental) unauthorized access.

Vulnerabilities

The adoption of standard modern IT hardware and software in control systems also exposes their inherent vulnerabilities to the world. PLCs, oscilloscopes and other controls devices (even valves or temperature sensors) are nowadays directly connected to Ethernet, but often completely lack security protection [3]. Control PCs are based on Linux and Microsoft Windows operating systems, where in particular the latter is not designed for control systems but for office usage. Even worse, compared to office PCs, these control PCs cannot be patched that easily, as this has to be properly scheduled beforehand. In addition, controls applications may either not be compliant with a particular patch or software licenses to run controls applications may become invalid. The same argumentation might apply to the deployment of anti-virus software, virus signature file updates, and local firewalls. At CERN, a few oscilloscopes got compromised since nobody realized that these were running the Windows operating system and thus, nobody bothered to patch them or installed anti-virus software.

Finally, using emailing, web servers or web cameras has become normal on control systems today; even web cameras and laptops can now be part of them.

The "vulnerability" factor can either be minimized by guaranteeing a prompt fix of published or known vulnerabilities, and/or by adding pro-actively protective measures to mitigate the unknown, potential or non-fixable vulnerabilities. The mitigation strategies will be discussed in the next Chapter.

Consequences

Within the High-Energy Physics (HEP) community, control systems are used for the operation of eventually very large particle accelerators and beam lines, the attached physics experiments, as well as for the technical infrastructure (e.g. power & electricity, cooling & ventilation). All are running a wide variety of control systems, some of them complex, some of them dealing with personnel safety, some of them controlling or protecting very expensive or irreplaceable equipment.

The consequences from suffering a security incident are inherent to the design of e.g. accelerators or experiments. These assets and their proper operation are at stake. A security incident can lead to loss of control, loss of beam time, and, thus, reduced efficiency, or loss of physics data. Even worse, consequences might be the damage or destruction of unique and expensive equipment and hardware: collimators being moved maliciously at the wrong moment into the beam line, protective devices failing to trigger a beam dump when needed, or beam dumps being initiated out of synchronization with the beam abort gap – and this is just the beginning of a much more exhaustive list.

Furthermore, physical destruction is just one side of the medal. Although the technical impact of the aforementioned incident at the LHC was small, the (negative) publicity created by it was enormous. In addition, the costs for dealing with this particular incident were not negligible and accumulated to 16 working days spent by the security team, IT experts, and experts of that particular experiment. The affected web server was under quarantine for more than three weeks, and has subsequently been reinstalled from scratch.

DEFENSE-IN-DEPTH FOR SECURE CONTROL SYSTEMS

Developers, system experts, and users inside the HEP community have to face the fact that the world is changing. Control system technology for accelerator controls *is* already employing standard IT solutions and techniques. And with this adaptation of the "interesting" IT-technologies (like Ethernet, the TCP/IP protocol, wireless access points, web pages, emails, Microsoft Windows, USB sticks) to the level of control systems, also the corresponding security technologies have to be inherited in order to mitigate the risks.

The worldwide most commonly used approach for mitigation is that of a "Defense-in-depth", which ensures that vulnerabilities are protected by multiple different protective means.

Focussing on single aspects corresponding to an "M&M"-principle ("hard on the outside, soft in the inside") has to be avoided: measures such as "Network security that's all you need!", "Firewall protection is sufficient", "...we're deploying only Linux" or "Our control system and network protocols are proprietary" should be put in the realm of wishful thinking.

On the contrary, and in the end much more advantageous, the "Defense-in-Depth" approach pro-active security measures must be applied to every possible level:

- ... the security of the device itself;
- ... the firmware and operating system;
- ... the network connections & protocols;
- ... the software applications;
- ... third party software.

These multiple layers offer the flexibility of not necessarily needing to act immediately to every (new) security risk – which, however, does not mean ignoring them completely. For example, segregated and well-protected networks allow "buying time" in order to postpone patching of control PCs to a more convenient moment, e.g. the maintenance window.

Such a "Defense-in-depth" approach must jointly be implemented by operators, system experts, developers, users, manufacturers and system integrators. At CERN, a

dedicated working group, the Computing and Networking Infrastructure for Controls project (CNIC) [4], has been set up implementing the recommendations of the British Centre for the Protection of the National Infrastructure (CPNI) [5] and the ISA SP99 standard of the American Instrumentation, Systems, and Automation Society (ISA) [6]. Their basic principles are detailed in the next sections. Implementation details of several other HEP laboratories worldwide can be found in [7].

Network Segregation

The basic means for all communication is the TCP/IP-based controls network for accelerator operations. It has to be segmented into smaller and separately protected network domains serving dedicated functions and purposes (e.g. injector domain, main accelerator domain, infrastructure domain, safety system domain). Interfaces between different network domains must be restricted and properly defined. All incoming and outgoing traffic of such a domain must be filtered (e.g. using professionally configured firewalls) such that only authorized traffic can pass. A direct connection to the Internet – taking also wireless access points and (GPRS) modems into account – must be avoided by all means.

Furthermore, it is highly beneficial, too, if the complete network domain for accelerator operations and its whole infrastructure are well separated from those used for development and testing. This also implies a clear separation between accelerator controls domains and the control systems deployed for e.g. running physics experiments or beam-lines – including the proper definition of accelerator/experiment interfaces.

Sensitive devices like PLCs have to be protected separately (e.g. by dedicated firewalls or using Virtual Private Networks), or have to be replaced by security-tested and robust ones.

Additional intrusion detection systems on the controls domains might be advantageous to detect the usual viruses and worms, but are currently only capable of analysing a small fraction of control-specific network protocols. In particular because of this, intrusion prevention systems have to be deployed with care, since blocking misidentified network traffic can lead to a halt of operation.

Remote access from outside onto any network domain must be carefully controlled (e.g. using "Application Gateways") or suppressed completely. Ultimate control of remote access permissions must remain by the shift leader for operations in order to allow him also to supervise activities outside the control room.

Patching, Patching, Patching

Although it is a given fact, that interventions on control systems and in particular on control PCs need proper scheduling, this should not be misused as an argument never or rarely to patch control PCs. A worm like the recent "Conficker" worm can infect thousands of Windows PCs within seconds, making no distinction between control PCs and others. A patch for the corresponding vulnerability (MS08-067) has been published in October 2008 and should have been applied by now to all instances. Despite this, the global number of PCs infected by this worm is still increasing.

Thus, operating systems of control PCs[‡] should be patched regularly, regardless of whether the operating system is based on Microsoft's Windows or on Unix/Linux.

In order to ease patching, the control system itself should be designed such that a restart of an individual control PC does not affect the overall availability and functionality. This is also advisable with respect to the rather short life-cycles of PCs. Dedicated test procedures have to validate the compliancy of a patch with the existing controls applications before the patch is widely deployed. A dedicated test stand is also beneficial in such a case.

A corresponding implementation at CERN [8] has proven that prompt patching of control PCs is feasible. The CERN Computer Management Framework (CMF) allows system experts to assume responsibility for the security of their control systems and their Windows PCs. CMF informs these experts of upcoming patches and provides them with means to deploy test patches. Bulk installation of patches can easily be scheduled and subsequently applied according to the corresponding maintenance plan. Furthermore, CMF allows a tight management of all installed software applications, and provides local firewall configurations and anti-virus software packages with automatic signature file updates. With CMF, the MS08-067 patch has been applied to nearly all Windows-based control PCs within a few weeks in line with the individual maintenance schedules.

Even if the discussion currently focuses on patching Windows PCs, this does not imply that Linux PCs or Apple Macs are more secure. While this was valid in the past, today all platforms suffer under the insecurity of applications running on them. Web browsers and browser plug-ins (e.g. for Adobe Reader, Java, ActiveX, QuickTime) are platform independent and so are their vulnerabilities [9]. Unfortunately, since web browsers are still part of any control PC and with web-based threats on the rise [9], patching gets even more important.

CERN has produced a similar framework to CMF for the Linux platform: Linux For Controls allows fine grained installation and patch control for Linux-based control PCs.

Robustness

With the means described above, control PCs can be reasonably secured, but usually control systems consist of more than that: hardware devices for accelerator controls are connected to the network, too. These Lynx O/S driven VME crates, power supplies and fan trays, PLCs, oscilloscopes, etc. usually completely lack security

[‡] In the following, the term "controls PC" includes also all other devices with embedded Windows or Linux operating systems like standard oscilloscopes from e.g. LeCroy or Tektronix.

protections – a fact proven by dedicated CERN security tests [3].

Thus, such devices must be tested for existing vulnerabilities prior to their deployment e.g. using standard IT vulnerability scanners like NMAP and Nessus[§]. Such scans also verify their robustness when receiving non-conform network packets and under high load of network traffic. Devices failing such scans should either be replaced or put under additional security protection. On the other hand, successful passes will confirm a basic level of robustness of those devices.

In addition to this, the configuration of these devices should be reviewed with respect to security. Protocols and services not needed for the operation (e.g. email, SNMP, Telnet, web servers) should be disabled or removed. Ideally, this is done in collaboration with the device manufacturer or vendor. The "Cyber-Security Procurement Language for Control Systems" [10] provides copy & paste paragraphs for procurement and service contracts.

Authentication & Authorization

Access to control PCs, their operating system, controls applications, control devices, and network domains used for controls must be tightly controlled and monitored. Authentication and authorization of operators, developers and experts must be handled restrictively and with sufficiently separated rights. Only those persons should be authorized, who have a professional need for access. The access rights have to be adapted accordingly (e.g. operators must be allowed to make general settings; system experts need access to specialised settings; guest access must be restricted to "read-only" – if at all). All access attempts must be logged and monitored regularly.

Accounts shared by several or many people must be avoided and replaced by individual accounts. Modern IT-technologies like magnetic strip readers or RFIDs allow the inconvenient and multiple keying of passwords to be avoided. Passwords for any remaining shared accounts must be tightly restricted in circulation. In addition, hidden accounts (e.g. installed by third party software) must be identified and disabled.

Access rights for remote maintenance, especially that of third parties, must be handled even more restrictively, or the operations might suffer from hidden side-effects like changing parameters without the acknowledge of the shift leader for operations.

In parallel to these software-based access controls, sufficient physical access protection measures must be deployed.

(New) Development Life Cycle

In the long-term the development, testing, deployment, and operation of (more complex) controls applications might benefit from established IT technologies. Here methodologies for the whole software development life cycle have existed for years [11], and today's accelerator controls applications should be able easily to adapt to this, too.

Control system development must be separated from operations and conducted on separated network domains – at least until debugging requires the interaction with the actual accelerator hardware. Therefore proper interfaces and procedures must be established, e.g. using software versioning and deployment platforms like CVS, Git or Subversion[**]. For the sake of safe and stable running, online changes to controls software during accelerator operation must be tightly controlled and avoided wherever possible.

The development life cycle should also include configuration management and documentation. All configuration parameters (e.g. threshold settings, device settings like IP addresses), dependencies, and system documentation must be stored centrally. Access must be properly secured, such that manipulations are only performed by authorized parties, and all changes are logged.

The Human Factor

While standard IT security offers many technical solutions, it necessitates the engagement and collaboration of both sides, that of the IT experts and that of the controls experts, in order to obtain satisfactory results.

While the operators, systems experts, and developers have an in-depth knowledge of their control systems, they (might) lack insight into certain IT concepts, in particular that of security. Vice versa, it is not granted that IT experts are sufficiently experienced to handle control system-specific aspects. Therefore, dedicated seminars and training sessions on the cyber-security of control systems have to be given in order to raise their awareness and knowledge. A close collaboration of both sides provides synergy effects and increases mutual trust.

An agreed and approved security policy dedicated for control systems should detail and define what is permitted and what is not inside the controls networks. This will avoid future misunderstandings.

Guidelines & Standards

Sponsored by the U.S. Department of Homeland Security, many private and commercial associations in the U.S. and in Europe (e.g. CIDX, ISPE, NERC, CPNI) have begun to develop and publish a large number of guidelines and standards. However, it is questionable, whether this cacophony is justified, or whether a smaller set of detailed and complete in-depth documents would be sufficient.

The aforementioned set of good practice guidelines of the U.K. CPNI [5] cover the basic aspects of control system cyber-security and provide initial mitigation

[§] http://www.nmap.org and http://www.nessus.org. Companies like Wurldtech and Mu Dynamics provide more sophisticated vulnerability testing.

[**] http://www.nongnu.org/cvs/, http://git-scm.com/ and http://subversion.org/

strategies. The SP99 series of the U.S. Instrumentation, Systems, and Automation Society is more detailed and covers the whole life cycle starting with the project definition, implementation, commissioning and deployment, and its operation [6]. The U.S. National Institute of Standards and Technology (NIST) also produced wide-ranging guidelines (SP800-53, -53A, and -82) [12]. These are in direct competition with the "Critical Infrastructure Protection"-standards (NERC CIP-002 to CIP-009) of the US Federal Energy Regulatory Commission (FERC) and the North American Electric Reliability Corporation (NERC) [13]. These CIP standards are now binding by law for all bulk electrical systems in the U.S. Finally, also the International Standardization Organization (ISO) has begun to extend its ISO 2700 series on guidelines and recommendations on information security management, risks and controls by dedicated aspects for control systems [14].

SUMMARY

Modern accelerator control systems today use many established IT technologies, especially Ethernet and the TCP/IP protocol. While this new functionality is appreciated, negative side effects like inherent security risks are often ignored. With the inheritance of standard IT technologies, network-based (virus and worm) attacks and vulnerabilities of the Microsoft Windows and Linux operating systems enter the scene. However, contrary to office PCs, control systems usually lack the standard protective measures.

In order to address this problem, leading organizations in the field of control system cyber-security recommend inheriting the corresponding IT security technologies for the protection of control systems. A useful approach is based on a "Defense-in-depth" strategy, which shields potential vulnerabilities on several levels and by different means: on the network layer, on the layer of the operating system, on the layer of software applications, and in the area of access control, authentication and authorization. Ideally, this is a collaborative effort between operators, system experts, developers, users, and manufacturers – and necessitates a tight collaboration between controls and IT experts, too.

"Secure" operation must become an additional objective for control system owners, and not just an answer to the push of a few security experts: Security must become an inherent property of any control system. Nevertheless, "security" is a permanent and iterative process. Continually improving security and addressing new security issues as they arise is the goal; ultimate security will remain utopia.

ACKNOWLEDGMENTS

The author would like to thank his colleagues in the CERN Security Team, and those involved in the realization of the CNIC project, in particular A. Bland, P. Charrue, L. Cons, I. Deloose, M. Dobson, U. Epting, N. Høimyr, S. Lopienski, D. Myers, H. Nissen, S. Poulsen, and M. Schröder.

REFERENCES

[1] Wall Street Journal, "Electricity Grid in U.S. Penetrated by Spies", 2009, http://online.wsj.com/article/SB123914805204099085.html.

[2] J.D. Rockefeller, "S. 773: Cyber Security Act of 2009", 111th Congress of the U.S. Senate, 2009.

[3] S. Lüders, "Control Systems Under Attack?", ICALEPCS05, Geneva, 2005, FR2.4-6O.

[4] U. Epting et al., "Computing and Network Infrastructure for Controls CNIC", ICALEPCS05, Geneva, 2005, O2_009; S. Lüders et al., "Update on the CERN Computing and Network Infrastructure for Controls (CNIC)" ICALEPCS07, Knoxville, 2007, WPPB38.

[5] Centre for the Protection of the National Infrastructure (CPNI), "Good Practice Guidelines Parts 1-7", 2006.

[6] The Instrumentation, Systems, and Automation Society (ISA), „Scope, Concepts, Models and Terminology", 2007, ISA 99.00.01; ISA, „Establishing a Manufacturing and Control Systems Security Program", 2007, ISA 99.00.02; ISA, „Operating a Manufacturing and Control Systems Security Program", 2007, ISA 99.00.03; ISA, „Specific Security Requirements for Manufacturing and Control Systems", 2007, ISA 99.00.04.

[7] S. Lüders et al., "Summary of the Control System Cyber-Security (CS)2/HEP Workshop", ICALEPCS07, Knoxville, 2007, MOPA01.

[8] I. Deloose, "The Evolution of Managing Windows Computers at CERN", HEPix, Rome, 2006.

[9] Symantec Corp., "Internet Security Threat Report Volume XIV", April, 2009.

[10] Idaho National Labs, „Cyber Security Procurement Language for Control Systems", 2007, http://www.msisac.org/scada.

[11] International Organization for Standardization (ISO), "Systems and Software Engineering — Software Life Cycle Processes", 2008, ISO 12207.

[12] National Institute of Standards and Technology (NIST), „Recommended Security Controls for Federal Information Systems", 2007, NIST SP800-53; „Guide for Assessing the Security Controls in Federal Information Systems", 2007, NIST SP800-53A Draft; „Guide to Industrial Control Systems (ICS) Security", 2007, NIST SP800-82 Draft.

[13] The Federal Energy Regulatory Commission (FERC), The North American Electric Reliability Corporation (NERC), „Cyber Security Standards", 2008, CIP-002 to 009.

[14] International Organization for Standardization (ISO), „Information Technology — Security Techniques — Specification for an Information Security Management System", 2005, ISO 27001 and sqq.