UPDATE ON THE CERN COMPUTING AND NETWORK INFRASTRUCTURE FOR CONTROLS (CNIC)

S. Lüders on behalf of the CNIC WG, CERN, Geneva, Switzerland

Abstract

Over the last few years modern accelerator and experiment control systems have increasingly been based on commercial-off-the-shelf products (VME crates, PLCs, SCADA systems, etc.), on Windows or Linux PCs, and on communication infrastructures using Ethernet and TCP/IP. Despite the benefits coming with this (r)evolution, new vulnerabilities are inherited too: Worms and viruses spread within seconds via the Ethernet cable, and attackers are becoming interested in control systems. Unfortunately, control PCs cannot be patched as fast as office PCs. Even worse, vulnerability scans at CERN using standard IT tools have shown that commercial automation systems lack fundamental security precautions: Some systems crashed during the scan, others could easily be stopped or their process data be altered [1]. During the two years following the presentation of the CNIC Security Policy at ICALEPCS2005 [2], a "Defense-in-Depth" approach has been applied to protect CERN's control systems. This presentation will give a review of its thorough implementation and its deployment. Particularly, measures to secure the controls network and tools for user-driven management of Windows and Linux control PCs will be discussed.

INTRODUCTION

The enormous growth of the worldwide interconnectivity of computing devices (the "Internet") during the last decade offers computer users new means to share and distribute information and data. In industry, this results in an adoption of modern Information Technologies (IT) to 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.

Unfortunately, the adoption of standard modern IT in distributed process control and automation systems* also exposes their inherent vulnerabilities to the world [1]. Furthermore, this world is by far more hostile than a local private controls network as the number and power of worms and viruses increase, and attackers start to become interested in control systems. Partial protection can be obtained through the usage of properly configured firewalls and through well-defined network architectures. However, some means of security incorporated into standard IT equipment cannot be directly applied to controls equipment since both differ in hardware but also in the concepts of availability and manageability.

At CERN, control systems are used for the operation of the whole accelerator complex, including the Large Hadron Collider (LHC), for the LHC experiments, as well as for the technical infrastructure (e.g. electricity, cooling & ventilation) and for fixed-target experiments.

In order to cope with the growing usage of standard IT technologies in control systems at CERN, the corresponding operation principles have been reviewed taking the aspect of security into account. This paper will give an update on the implementation of the Security Policy presented by the CERN Computing and Network Infrastructure for Controls (CNIC) working group two years ago [2].

CNIC SECURITY POLICY

The CNIC Security Policy gives directions on how to secure CERN control systems. It is necessary to reduce the overall risk to a minimum in order to obtain maximum security, where "risk" is defined as:

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

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

The "consequences" are inherent to the design of CERN's accelerators and the affiliated experiments. 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. Thus, CERN's assets and their proper operation are at stake. A security incident can lead to loss of beam time and physics data, or — even worse — damage to or destruction of unique equipment and hardware.

The "vulnerability" factor can be either minimized by guaranteeing a prompt fixing of published or known vulnerabilities, and/or by adding pro-active measures to secure the unknown, potential or not-fixable vulnerabilities.

In order to protect such vulnerabilities against being exploited (either because there is no patch available or a patch could not be applied), the Security Policy follows

* Commonly denoted in the following as "control systems", where a "system expert" has the expertise in its configuration.

the recommendations of the U.K. CPNI [3]. It is based on a "Defense-in-Depth" approach, where pro-active security measures must be applied to every possible level:

- …of the security of the device itself;
- …of the firmware and operating system;
- …of the network connections & protocols;
- …of the software applications;
- …of third party software;
- …together with users, developers, and operators.

"Defense-in-Depth" is, thus, based on four major pillars: "Network Security", "Central Installation Schemes", "Authorization & Authentication", and "User Training".

However, sufficient protection should not provide false security. Incidents will happen. Therefore, the Security Policy also defines rules to deal with "Incident Response & System Recovery", as well as with regular security audits. The next chapters will outline the major implementations in detail.

NETWORK SECURITY

In order to contain and control the network traffic, the CERN network has been separated into defined "Network Domains". For example, each of the LHC experiments is now using such a dedicated Domain. CERN's accelerator complex and the control systems which monitor and control the technical infrastructure (e.g. electricity distribution, cooling & ventilation, facility management as well as safety and access control systems) use another. "Domain Administrators" were assigned to take full responsibility for a Domain. In particular, they supervise the stringent rules for connecting devices to it. Additional network segregation allows a system expert further to protect vulnerable devices like Programmable Logic Controllers (PLCs).

Each Domain is interconnected with CERN's "General Purpose Network" used for office computing, and other Domains via dedicated network routers. The traffic crossing any two Domains is restricted to a minimum by the usage of routing tables, with only mandatory traffic passing such boundaries. A large fraction of this traffic is either dedicated data exchange with CERN's Computing Centre, or currently inevitable due to the ongoing commissioning of the LHC accelerator.

Windows Terminal Servers (WTS), and to a lesser extent Linux-based application gateways, have become the only means for remote user access.

CENTRAL INSTALLATION SCHEMES

From experience at CERN, only centrally managed and patched PCs have shown to be secure in the long run. However, since control PCs are very different in handling than office PCs, a more flexible installation and management scheme was needed. While the operating systems, antivirus software, and basic software applications should continue to be managed and maintained by the IT service providers, it is up to the system expert to take over full flexibility of the configuration of the PCs of his system — and full responsibility for securing it. Such a scheme will also help the expert to recover e.g. from a security incident.

Schemes for the central installation of CERN Scientific Linux and of Windows, respectively, have been created.

Linux For Controls (L4C)

L4C is based on a bottom-up approach to the configuration of the PCs in a hierarchical manner, and uses the same techniques having proven to be successful for managing Linux clusters in the CERN Computing Centre, i.e. using Quattor (http://quattor.web.cern.ch) for configuring PCs ("the nodes"). Settings that are specific to a node are defined in so-called "node templates". The individual nodes join sets of PCs with a common configuration. These sets in turn can join super-sets. L4C supports CERN Scientific Linux 3 and 4, including all regular security updates and enhancements. Basic templates are maintained by L4C support, all other templates by the system experts allowing him full flexibility. The templates can be hosted in central or user owned configuration databases (CDBs).

The Linux installation of a PC from scratch uses CERN's "Automated Installation Management System" (AIMS) service and is based on RedHat's "Kickstart" software. Booting a Linux PC via the network using PXE Boot (Preboot Execution Environment) pulls the "Kickstart" and configuration profile from the CDB. From this information, Quattor is able to perform a full automatic installation.

From here, the CDB informs a node about any new changes in the configuration profile. Triggered by CDB, a local daemon then downloads the modified profile and subsequently applies the changes. For each node in the CDB, a webpage shows the names, versions, hardware architecture and the repository from which all the packages for a node are to be installed. However, in order to verify that all packages are installed as foreseen the system expert has to log onto that node.

Computer Management Framework (CMF)

CMF [4], on the other hand, has implemented a top-down structure, focussing on sets of PCs with a defined configuration. The installation of PCs is handled through a top-down tree of so-called "Named Sets of Computers" (NSCs). Each NSC assigns a list of applications to its members, where these members can be individual PCs or other, nested NSCs. A PC that is member of different NSCs will receive the applications from any of them. CMF is taking care of clashes. Depending on the type of NSC, the administrator of the NSC, i.e. the system expert who maintains that NSC, has full autonomy of his configuration ("locally managed"), or CERN's IT department is still providing a basic configuration (i.e. that of an office PC), and takes care of patches.

A long list of basic applications has been provided as CMF "packages". Packages can be applications but also group policy settings, regular scheduled tasks, or Visual Basic scripts. NSC administrators can easily create

additional packages using the CMF web-interface and simple scripting. The installation of a package can be either forced ("applied"), such that it is installed automatically after a small notification period, "denied", such that it cannot be installed at all, or offered ("published") to the interactive user. In the latter case, the interactive user can use the CMF web-interface to select/deselect packages he wants to install/de-install.

The installation of these packages is controlled via a small local program (the "CMF Agent"), being installed on every CMF Windows PC. It handles all pending (de)installation tasks, interacts with the user, and performs regular inventory checks which are passed back to a central CMF configuration management database.

The initial installation from scratch is based on the Windows pre-installation environment and PXE Boot. The preferred operating system (Windows XP SP2 or Windows 2003 terminal server) can either be chosen from a list, or predefined for automatic installation on the CMF web-interface. After the operating system has been installed, the CMF Agent controls the subsequent installation of all packages being applied to that particular PC.

With PXE Boot and the proper configuration of his NSCs, the system expert has full liberty to install sets of PCs in parallel or to run pilot installations prior to mass deployment. CMF ensures that all members of a NSC are identically configured, and that corrections or modifications are propagated accordingly. The configuration database provides always an up-to-date status (e.g. PC hardware, operating system version, patch levels, installed applications and their versions) via the CMF web-page. Queries can be run to assess a particular situation (e.g. listing all PCs missing patch XYZ).

AUTHENTICATION & AUTHORIZATION

Several dedicated authentication & authorization schemes have been developed at CERN serving the accelerator complex [5] and LHC experiments [6]. These are based mainly on general standards like Role-Based Access Control [5] or on commercial solutions [6]. Details can be found in these proceedings [5][6].

Major challenges still to be overcome are the generalization to one common central scheme at CERN.

USER TRAINING

A series of awareness campaigns and training sessions have been held for users, operators, and system experts at CERN. Monthly CNIC meetings provide a forum for questions and discussions.

Furthermore, CERN has raised aspects of Control System Cyber Security at several conferences and workshops (e.g. at the CS2/HEP workshop [7]), interacted with major vendors of control systems, and is now leading the "EuroSCSIE", the European Information Exchange on SCADA (Supervisory Control And Data Acquisition) Security, with members from European governments, industry, and research institutions that are dependent upon and/or whose responsibility it is to improve the security of SCADA and Control Systems.

INCIDENT RESPONSE & SYSTEM RECOVERY

Even with a stringent Security Policy incidents can never be prevented completely. Therefore, handling incidents on a Domain have been and will be jointly performed by CERN's Computer Security Team and the corresponding Domain Administrator. The acting Computer Security Officer has the right to take appropriate actions in justified emergency cases.

After incident analysis, the central installation schemes CMF and L4C allow for a rapid system recovery by the corresponding system expert.

SUMMARY

Due to the continuing integration of common IT technology into control systems, the corresponding IT security vulnerabilities and cyber-attackers end up threatening control systems, and, thus, CERN's operation and assets. However, control systems demand a different approach to security than office systems do.

This paper has presented a thorough rule-set to secure CERN's control systems. Its implementation uses a "Defense-in-Depth" approach based on network segregation, central installation schemes, authentication & authorization, user training, incident response & system recovery, and security auditing.

ACKNOWLEGDMENTS

The author would like to thank all his colleagues at CERN involved in the realization of CNIC making it such a success; in particular A. Bland, P. Charrue, I. Deloose, M. Dobson, U. Epting, N. Høimyr, S. Poulsen, and M. Schröder.

REFERENCES

[1] S. Lüders, "Control Systems Under Attack ?", ICALEPCS, Geneva, October 2005.

[2] U. Epting et al., "Computing and Network Infrastructure for Controls", ICALEPCS, Geneva, October 2005.

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

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

[5] S. Gysin et al., "Role-Based Access Control for the Accelerator Control System at CERN"; K. Kostro et al., "Role-Based Authorization in Equipment Access at CERN"; and A. Petrov et al., "User Authentication for Role-Based Access Control", these proceedings.

[6] P. Chocula, "Cyber Security in ALICE", these proceedings.

[7] S. Lüders, "Summary of the Control System Cyber-Security CS^2/HEP Workshop", these proceedings.