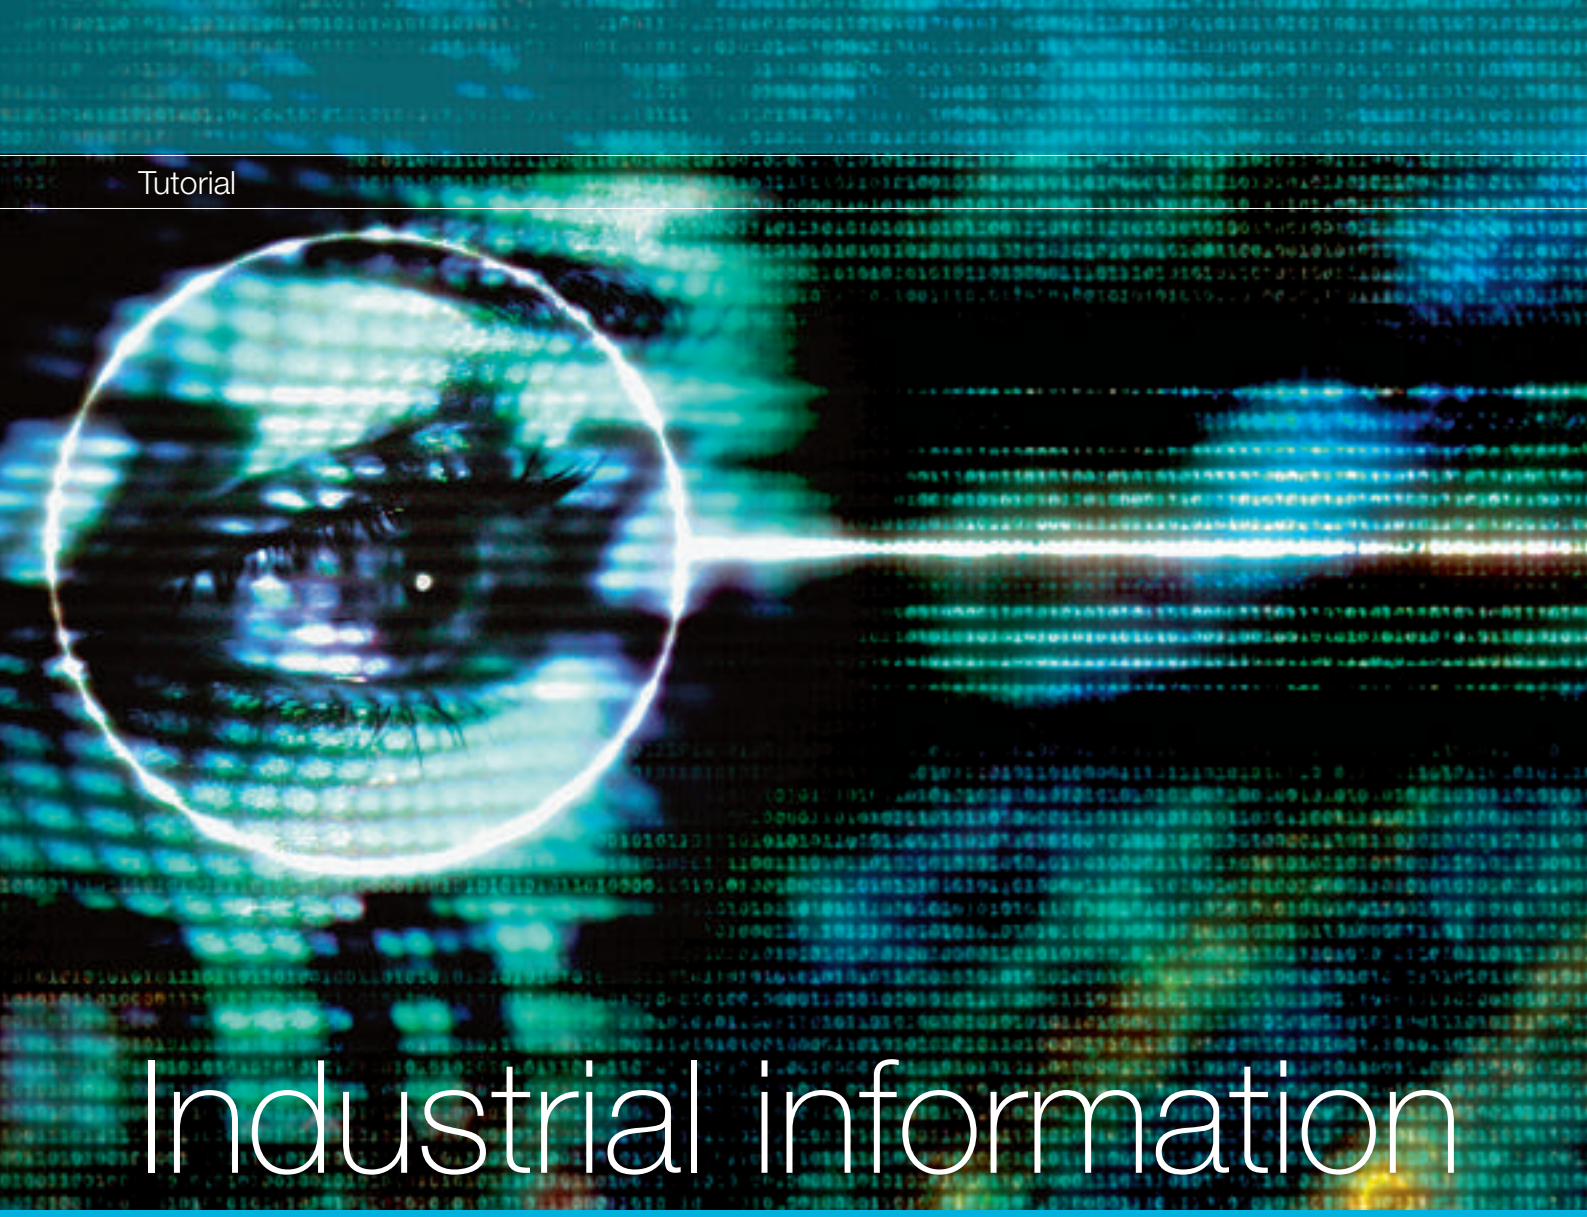

Industrial information system security Part 1

IT security in industrial plants - an introduction

Martin Naedele, Dacfey Dzung

The topic of IT security is becoming increasingly relevant in modern automated industrial plants. Modern automation systems provide high levels of inter-connectivity. Implementations are based on commercial IT platforms, many of which are known to be vulnerable to electronic attacks. This article is part 1 of a three part tutorial on IT security for industrial systems.

Security objectives are explained, attack types are illustrated, and available and suggested countermeasures and general good practices are briefly described. Part 2 of this series will address best practices in protecting against certain types of attacks, and Part 3 will survey emerging standards for automation system security.

In the past, industrial automation systems were not linked to each other and were not connected to public networks like the Internet. Today, the situation is somewhat different: because the market puts pressure on companies to make fast and cost efficient decisions, accurate and up-to-date information about the plant and the process status must be available not only on the plant floor, but also at the management level and even for supply chain partners [1]. This results in greater inter-connectivity between different automation systems and between automation and office systems. Modern industrial automation systems are, to a large extent, based on commercial operating systems, protocol implementations, and communication applications originally developed for the office IT environment. Many of these systems and implementations are known to be vulnerable to attacks, and with open and standardized Internet-technologies, expertise and knowledge of such vulnerabilities is easily available to potential attackers. By connecting industrial plants to the Internet or to other public networks these vulnerabilities are exposed. Thus, IT security issues must also be addressed in industrial automation systems.

What is IT security?

For many people, IT security is considered a synonym for encryption and for others, the foremost IT security issue concerns protection against computer viruses. In reality however, IT security has a much wider scope. The following eight *security objectives* are a suitable framework for structuring security requirements and properties of a system:

Confidentiality: The confidentiality objective deals with preventing the disclosure of information to unauthorized persons or systems. For automation systems, this is relevant both with respect to process specific information, such as product recipes or plant performance and planning data, and to the secrets specific to the security mechanisms themselves, such as passwords and encryption keys.

Integrity: The integrity objective deals with ensuring that modifications made by unauthorized persons or systems to specific information are detected. For automation systems, this applies to information such as product recipes, sensor values or control commands. Violation of integrity may cause safety issues, ie, equipment, the environment, or even people may be harmed.

Modern industrial automation systems are, to a large extent, based on commercial operating systems, protocol implementations, and communication applications originally developed for the office IT environment.

Availability: Availability means ensuring unauthorized persons or systems cannot deny access/use to authorized users. For automation systems, this refers to all elements of the plant like: control systems; safety systems; operator workstations; engineering workstations; manufacturing execution systems; and the communication systems between these elements and to the outside world. Violation of availability, also known as denial-of-service (DoS), may not only cause economic damages but also safety issues as operators may lose the ability to monitor and control the process.

Authentication: Authentication is concerned with determining the true identity of a system user and mapping this identity to a system-internal principal (eg, valid user account) under which this user is known to the system. Most other security objectives, most notably authorization, distinguish between legitimate and illegitimate users based on authentication.

Authorization: The authorization objective – also known as access control – is concerned with preventing people (or systems) who do not have permission from accessing the system. In the wider sense, authorization refers to the mechanism that distinguishes between legitimate and illegitimate users for all other security objectives, eg, confidentiality, integrity, etc. In the narrower sense of access control, it refers to restricting the ability to issue different types of commands to the plant control system. Violation of authorization may create safety issues.

Auditability: Auditability is concerned with being able to reconstruct the complete system behavior history from records of all (relevant) actions executed on it. This security objective is mostly concerned with discovering and finding reasons for malfunctions in the system and to establish the scope of the malfunction or the consequences a security incident. Note that auditability without authentication may serve diagnostic purposes, but does not provide accountability.

Non-repudiability: The non-repudiability objective means being able to provide irrefutable proof to a third-party of who initiated a certain action in the system, even if this actor is not cooperating. This security objective is relevant in establishing accountability and liability. In the context of automation systems, this is most important with regard to regulatory requirements, eg, US Food and Drug Administration (FDA) approval. Violation of this security objective may have legal

1 Security mechanisms used between control system and external networks as reported in [3]. Note that 5 % of systems are not secured in any war.

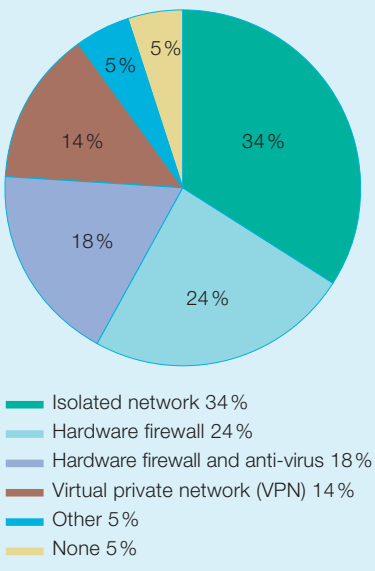

and commercial consequences, but no safety implications.

Third party protection: A successfully attacked and subverted automation system could be used for various attacks on the IT systems, data or users of external third parties using, for example, distributed denial-of-service (DDoS) or worm attacks. The third party protection objective deals with preventing this type of damage from occurring.

The importance of each security objective depends on the system, specifically its purpose and its assets. In automation systems, for example, confidentiality is important for production and performance data, while integrity and authorization is most relevant for operator commands, parameters, and control functions. For each system and installation, a *security policy*, stating the security objectives and specific system constraints must be in place before the security architecture for any system can be designed.

At this point it may be worth pointing out the difference between the concepts of *security* and *safety* as applied to an automation system or plant. Although no undisputed definitions for these terms exist, they tend to be used in the following way: security is concerned with the prevention of intentional malicious attacks whereas safety is concerned with the prevention of damage caused by a predominantly unintentional or random loss of integrity and availability of plant components, or by user error.

What types of attacks are there?

An *attack* is a violation of one or more security objectives and can be initiated either inside or outside the plant. Attacks may target a specific system or type of system, or they may, for example, in the form of viruses and worms simply victimize any vunerable system they encounter.

For each system and installation, a security policy, stating the security objectives and specific system constraints must be in place before the security architecture for any system can be designed.

A computer may be attacked for example to: obtain specific data stored on the computer (eg, production data); to abuse the processing or storage resources of this computer (eg, to store and distribute pirated software); to use applications installed on the computer to manipulate data or other systems (eg, a production plant controlled by the computer); or to prevent usage of this computer for its intended purpose.

A data transmission link is another possible target and may be attacked: to eavesdrop on transmitted information; to falsify the information sent to the recipient of the transmission; or to prevent the legitimate use of the transmission link by, for example, flooding it with messages.

Do such attacks really happen?

January 1998: External attackers took over the central control center for the Gazprom pipeline system. For an unknown period of time they were able to control the flow in the whole Gazprom pipeline network[1)].

March 2000: A disgruntled former contractor gained access to the control system of a sewage treatment plant in Maroochy Shire in Queensland/Australia. He flooded the surrounding environment with millions of litres of untreated sewage[2)].

December 2000: Attackers compromised the computer network of an unnamed power utility in the US via an unsecured data exchange protocol. They used the compromised hosts to play networked computer games. Their usage of computing resources and network bandwidth severely impedeed the utility's electricity trading[3)].

January 2003: The safety monitoring system of the Davis-Besse nuclear power plant in the US was infected with the "Slammer" worm. The worm bypassed the plant's firewalls via a contractor's laptop - which was connected to the power plant network at the same time - and via a modem to the infected enterprise network of the contractor company[4)].

August 2003: At CSX Transportation, a US railway company, a worm infected the communication network used for signalling, bringing all trains to a halt for half a day[5)].

Table 1: Which security mechanism for which security objective?

Security objective	Security mechanisms
Confidentiality	Encryption, Virtual Private Network (VPN), Secure Socket Layer (SSL)
Integrity	Cryptographic checksums, malware scanners
Availability	Redundancy, diversity, malware scanners
Authentication	Pass phrases, certificates, tokens/smartcards, biometrics, challenge-response protocols
Authorization	Hardened operating systems (no insecure or unused services, user accounts; tightly defines access control lists (ACLs) on resources, etc.), firewalls, personal firewalls, application level message filters, Virtual LAN (VLAN)
Auditability	Intrusion Detection System (IDS), logs
Non-repudiation	Digital signature
3rd party protection	Firewall (egress filtering), malware scanner (for outgoing data)

Footnotes:

1) http://www.gtiservices.org/security/riskassess/gazprom_attack_04261999.doc
2) http://www.theregister.co.uk/content/4/22579.html
3) http://zdnet.com.com/2100-11-526431.html?legacy=zdnn
4) http://www.theregister.co.uk/content/56/32425.html
5) http://www.csx.com/?fuseaction=company.news_detail&i=45722&news_year=-1

May 2004: The "Sasser" worm infected the signalling and control system of the Australian railway company, RailCorp. 300,000 commuters in and around Sydney had no transportation on this day[6].

These examples show that electronic attacks on industrial control systems really do happen. In addition, it can be safely assumed that a large number of attacks have not been reported in the press. Canadian researchers maintaining a confidential database of IT security incidents in industrial installations have observed an increase in incidents and a shift from internal to external attacks [2]. It is worth mentioning, however, as far as details about the above mentioned incidents are known, they were all possible only because suggested practices were disregarded. A survey conducted by ARC in 2004 indicates that a significant number of control systems connected to external networks have no security mechanisms 1.

Which security mechanisms should be used?

A risk of an attack exists if there is an exposed *vulnerability and a threat.* The *vulnerability* of an information system may be caused by a logical design flaw (eg, a wrongly designed protocol), an implementation mistake (eg, allowing a buffer overflow), or a fundamental weakness (eg, passwords and cryptographic keys that are guessable by trying out all possible permutations).

Threats to a system are the potential goals of attackers, for example, to disrupt production for a certain period of time. Threats may also be realized by an incidental, non-intentional exploitation of a vulnerability.

The *risk* of a given attack is determined by the *likelihood* of a successful attack and the *severity* of the damage it may cause. For a given system, a *threat analysis* must be performed during which risks are evaluated and ranked in importance. This analysis forms the basis for the security policy, where the relevant security objectives are specified. This finally determines the *security mechanisms* to be deployed. Security mechanisms reduce the risk to a system by making the exploitation of vulnerabilities less likely or by limiting the damage. 2 illustrates the interrelations between the IT security terms used in threat and vulnerability analysis.

Table 1 shows which security mechanisms are commonly used to reduce the risk of certain types of security objective violations:

Suggested best practices

Securing a system is difficult as it is necessary to spread effort and budget to efficiently and effectively prevent a wide variety of attacks. Based on experiences that have been made in this field over the last few decades, the following best practices should be taken into account when deploying technical security measures or implementing procedural controls:

Avoid a weak link: The effort spent on protecting the various interdependent security objectives required for a system has to be distributed so that all mechanisms facing an attacker are of comparable strength. Otherwise, an attacker could bypass a strong mechanism by breaking a weak one. In security systems humans are often the weakest link, the effect of which places greater importance on procedures and training.

There is no "security by obscurity": It used to be argued that automation systems were secure because very few people had sufficient detailed knowledge about their operations and protocols to attack them. Unfortunately this is no longer the case. Many automation experts exist in all parts of the world, and automation systems are largely based on well-documented open standards.

Least privilege: By only giving users the minimum permission necessary to

Footnotes:
6) http://news.com.au/common/story_page/0,4057,9455677%255E15306,00.html

2 IT security threat analysis terms and their interrelations.

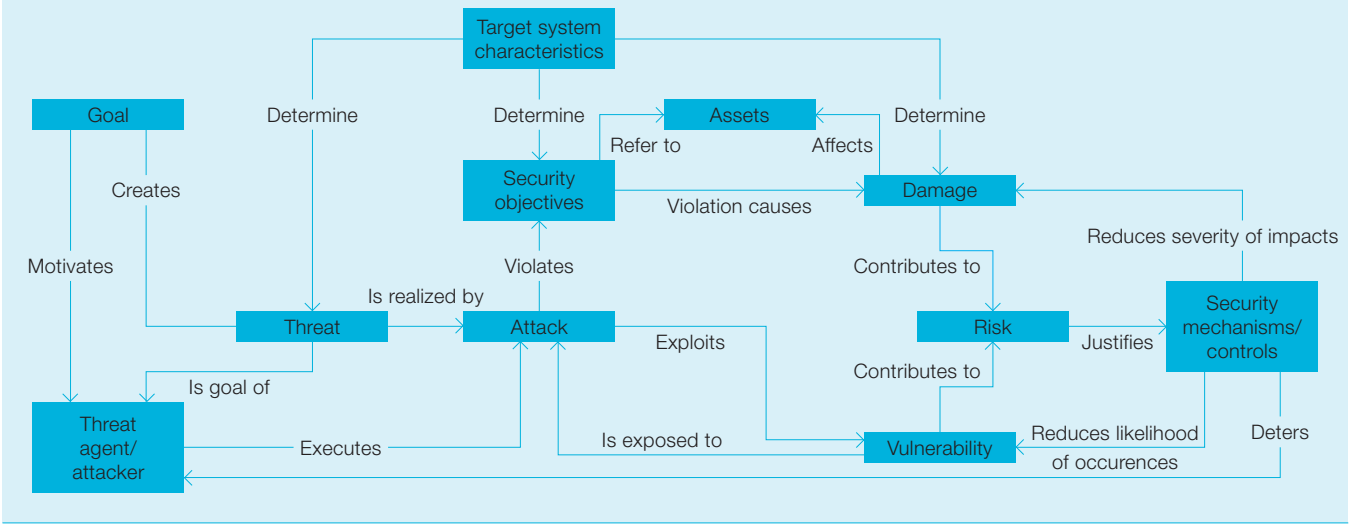

do their job, the risk of insider attacks or the abuse of compromised user identities is reduced.

Important secure system design principles

Two design principles should be considered when architecting a secure system:

Defense-in-depth: There are two commonly used basic approaches for securing physical and information systems: *hard perimeter* and *defense-in-depth*. The idea behind the hard perimeter approach is to put a single impenetrable wall around the system and to disregard all security issues inside. In the defense-in-depth approach several zones are placed around the object to be protected. Different types of mechanisms for detecting and delaying an attacker are used concurrently inside and around each zone. The outer zones contain less valuable targets. Properly implemented defense-in-depth security architectures are more resistant to attacks than hard perimeter architectures.

Security is a process, not a product: Due to changes in the operating environment and the availability of new attacks, no security system even if it is implemented flawlessly, will be able to fulfill its purpose forever without maintenance. Maintenance includes reviewing access control rules and updating installed software. These reviews compare actual state with the defined state of the system. In addition, they assess whether the defined state is still appropriate in the face of a changing business and risk environment. The need for maintenance means that continuous financial and staffing resources are needed to keep a system secure.

Where do we stand?

Even though reports about spreading worms and a general increase in network based attacks seem to dominate the news in IT related publications, this is not really a reason to forego the enormous benefits that full vertical and horizontal integration based on network interconnections can bring to an enterprise. As long as well-known good practices are respected in implementing and operating network connectivity between the control system and other networks, an adequate level of security can be achieved for any application. That security level then represents the residual risk that is regarded as acceptable after a thorough threat and risk analysis for the particular installation.

Part 2 of this tutorial will explain how an automation system may be protected against damages from worms and viruses. Part 3 will show what industry standards are emerging in the field and how they can be used to reduce the effort to secure a plant against targeted and untargeted attacks.

"Security" is not a fixed target. Continuous efforts are necessary to keep any system secure. In the case of control systems, both the plant operating enterprise and the automation vendor have to be involved in these efforts.

Glossary

Cryptographic Checksum	Check-bits calculated from the content of a document or message and a secret key, such that any unauthorized changes in the document can be detected.
Challenge-Response	A sends a challenge to B: a question whose answer can only be given if a common secret (password) is known. If B responds correctly, it has proven its identity to A, without having sent a password. (Sending a password is vulnerable to eavesdropping.)
Public Key Cryptography	Encryption method using a pair of keys: Encryption with the public key results in data which can only be decrypted by the receiver using the matching private key. The private key must be kept secret.
Digital Signature	Check-bits appended to a document or message, calculated from the document and a secret known only by the sender. Used to prove that the document originates from the claimed sender.
Digital Certificate	Digital "passport", attests that a public key belongs to the claimed owner. Issued (signed) by a trusted certification authority.
Virtual Private Network (VPN)	Private network running over encrypted tunnels through the public Internet.
Firewall	Device or program which inspects all incoming (ingress) or outgoing (egress) messages. Messages are filtered (blocked or forwarded) based on source/destination addresses or application level contents.
Intrusion Detection System (IDS)	Devices observing traffic in a network, in order to detect electronic intrusions and raise alarms. Detection is based on attack signatures or anomalies in the traffic patterns.
Secure Socket Layer (SSL)	Security protocol widely used to authenticate Web servers and to establish encrypted communication between Web browsers and servers. Uses digital certificates.

Dr. Martin Naedele
Dr. Dacfey Dzung
ABB Switzerland, Corporate Research
martin.naedele@ch.abb.com
dacfey.dzung@ch.abb.com

References:

[1] **Leffler, N. , Terwiesch, P.:** Aspects of Productivity, ABB Review 2/2004.

[2] **Byres, E. , Lowe, J.:** The Myths and Facts behind Cyber Security Risks for Industrial Control Systems, VDE Kongress 2004.

[3] **Forbes, H.:** Plant Floor Network Practices in Today's Factories and Plants, ARC insight 2004-53EMHLP, December 2004.

Further reading:

M. Naedele: IT Security for Automation Systems, in: **R. Zurawski** [Editor]: Industrial Information Technology Handbook. CRC Press, January 2005, ISBN 0-8493-1985-4.

M. Naedele: IT Security for Automation Systems - Motivations and Mechanisms, atp international, Vol 1 (1), 11/2003, and atp, Vol 45 (5), 5/2003.

R. Anderson: Security Engineering, Wiley, 2001.

Industrial information system security Part 2

Malware protection for industrial automation systems

Martin Naedele, Rolf Vahldieck

The previous part of this three-part tutorial on information system security in industrial networks introduced the basic terminology and explained the need for security measures. This need was, in part, motivated by a number of reported incidents where worms infected automation systems. There is no doubt that worm infections pose a real threat for networked automation systems today.

This part of the tutorial explains the different types of malware and suggests how an automation system can be defended against them. A case study is presented where some of the suggested mechanisms are applied in a real-world automation system.

Protecting their networked automation systems against viruses, worms, and related threats (so-called malware) is one of the foremost security concerns many enterprises must deal with. As the incidents discussed in the first part of this tutorial [1] have shown, this is indeed a serious issue. Damage caused by malware may include:

- The obstruction of communication between hosts on the LAN because of high loads on the network.
- Host operating system crashes and forced shutdowns.
- Permanent damage to application and data files stored on the hard disks of those hosts.
- Malware can also send out information from the system to outside destinations or open a backdoor into the system, allowing remote control of this host by an attacker.

Any of these symptoms of a malware infection are unacceptable in a manufacturing and process control environment.

This article looks at different strategies that can be used on technical and procedural levels to defend against such threats.

What types of malware are there? And how do they infect a system?

There are four main types of malware, though actual implementations may have characteristics of multiple types [2].

A *virus* is a program that infects a host as part of a legitimate data transfer, and is generally attached to another message or file. In the pre-network days of computing, hosts were infected with viruses via floppy disks. In particular, the virus was often situated in the part of a disk that is read and executed whenever the computer was started (so called boot sector viruses). Nowadays, viruses spread mostly via executables or documents containing so-called active content (macros, scripts) sent over the network, for example as attachments to emails, instant messaging, or as part of web pages. Infections via portable storage media do happen, but more rarely.

The action a virus executes when it is invoked is called its payload. A virus payload may display a message on the screen or it may completely take over and manipulate the display. It may change or delete random files or even the entire file system, or it may crash the system. Its most characteristic action is that, in addition to other malicious actions, it spreads by attaching itself to other programs and documents on the host.

A virus relies on one or even multiple user actions, ie, starting a program, surfing to the web page or inserting a portable storage medium to be propagated and/or activated. Of course, in most cases the user is not aware that such actions will activate the virus.

1 The geographical spread of Slammer within 30 minutes of its release. The diameter of each circle is a function of the logarithm of the number of infected machines, so large circles visually under-represent the number of infected cases in order to minimize overlap with adjacent locations. For some machines, we can determine only the country of origin rather than a specific city. [Figure and caption taken from [3]].

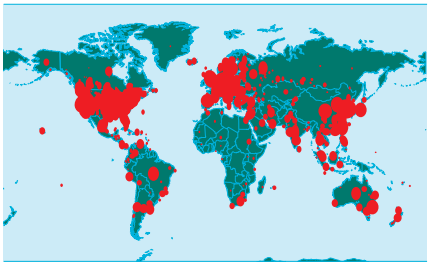

A *worm* is very similar to a virus but the major difference is that it does not require a user action for propagation. A worm can be regarded as a self-propagating virus, copying itself from host to host via the network. Typical worms use the e-mail system to propagate – they mail themselves from the infected host to recipients taken from the user's e-mail address book, or they directly scan other hosts on the LAN or network addresses on the Internet for potential entry points and then open a connection to complete the transfer.

As worms are not activated by user actions they are often located in the host's memory and can be removed by rebooting. To date, few worms have had a payload that has caused permanent damage to the infected hosts, though, of course, this is not a guarantee that future worms will not be destructive. Worms often overload the networks around the infected host as they spread to other destinations 1 and Table 1.

A *Trojan* (like the proverbial horse) is an application that camouflages itself as a harmless piece of software but in fact has an additional malicious functionality. A Trojan can be regarded as a virus without a propagation component. A typical function of a Trojan is to provide a backdoor into a system for remote control by an external attacker. For this purpose, the Trojan wants to remain undetected for long periods and therefore has no interest

Table 1 Major worm/virus epidemics in recent years [4]

Name	First seen	Means of propagation	Size of epidemic and damage type
Sasser	30 April 2004	Vulnerability in Windows LSA service	>= 1,000,000 hosts infected (repeated shutdown/reboot)
Blaster	11 August 2003	Vulnerability in Windows DCOM	>= 500,000 hosts infected
Slammer	25 January 2003	Buffer overflow in SQL Server	>= 75,000 hosts infected (network saturation)
CodeRed	19 July 2001	Buffer overflow in IIS	>= 350,000 hosts infected (network congestion)
ILOVEYOU	4 May 2000	VB script attachment processed by Outlook	>> 10,000,000 hosts infected
Melissa	26 March 1999	Mail attachment; on opening sending itself to Outlook addressbook entries	>= 100,000 hosts infected

in causing obvious malicious effects noticeable by a user.

In fact certain Trojans, so called root kits, modify significant parts of the operating system to prevent detection. They may replace, for example, operating system utilities that list files on the hard disk or show open network connections by versions that hide the Trojan activity.

A special kind of Trojan is spyware. Spyware applications are active in the background and transmit certain information from the infected host, such as software applications installed on the host, or web pages visited, to outside recipients. Like Trojans, *spyware* tries to stay undetected over extended periods of time.

How to prevent an infection?
Generically speaking, malware infections can be prevented by limiting their means for propagation and communication.

To prevent virus infections, the different ways a virus infected file enters the system must be controlled. Fortunately for automation systems, it is possible to eliminate most of these attack paths [5]:

E-mail or instant messaging (IM/IRC) attachments:
there should be no incoming e-mail or IM/IRC traffic to hosts on the automation system network. This can be enforced by not installing any e-mail or IM client applications on the computers, or disabling those that are part of the operating system. Additionally the corresponding protocols should be blocked at the firewall(s) between the automation network and outside networks should be blocked. If there are operational reasons for receiving e-mail and using IM facilities in the control room, a separate computer can be used which is not connected to the automation network but directly to the enterprise intranet. This computer may also have access to outside communication networks, for instance via dial-up modem or ADSL. Also on this computer the e-mail or IM client applications should be configured so as not to execute any active content. In addition filtering applications should be used to remove attachments, scripts and other dangerous ingredients from any incoming message.

Web browsing:
like incoming e-mail, web browsing to arbitrary sites on the Internet from any host in the automation network should be prohibited. As de-installing a web browser is often not seen as a realistic option, filtering at the firewall and VPN setup can be used to ensure only a small number of pre-approved websites – such as reference documentation on the enterprise intranet or a vendor support site – can be reached. This also mitigates the threat that users will work around the email client restrictions mentioned in the previous section by using web based email clients. Again, the use of separate hosts not connected to the automation network is suggested for web access. All downloaded content should be filtered for active and potentially malicious elements. Ideally, only pure HTML should reach the client application.

Malware infections can be prevented by limiting their means for propagation and communication.

Drive sharing:
drive sharing between hosts inside and outside of the automation networks should generally be prohibited. This can be enforced by firewalls. More secure protocols exist for transporting data in and out of the automation network 2.

Portable media (CD, DVD, disk, memory stick, etc.):
processes, procedures, and supporting physical means (eg, drive locks) in the plant should ensure that all portable media are scanned for malware with up-to-date scanning engines and signatures before they are brought into contact with a host on the control network. Extra protection might be provided by multiple scans using different anti-virus products. Of course the host used for malware scanning should not be connected to the control network.

Data transfer applications (FTP, PC Anywhere, etc):
direct transfer of files into the control system by any of those means should be restricted to a minimum. An intermediate staging server should be used to scan all incoming data for malware. Digital signatures can help verify that the file really originates from the assumed, and trusted sender, and that it wasn't modified between the virus scan and import into the control system.

Portable computers:
every portable computer should have the most up-to-date operating system patches installed as well as a current malware scanner running. Recently, several security vendors have started to offer solutions to verify a computer fulfills certain requirements before it can connect to the network.

Externally accessible service:
every host should be "hardened". This means that unrequired services and user accounts should be removed/disabled, and user access rights and policies are set such that every user (including system accounts) only has the minimum level of privileges necessary for operations.

The dataflow architecture of the system should ideally be designed in such a way that no requests are needed from external PC clients to the control network. In this case the firewall can be configured to block all connections from the outside, including worm scans. If ports must be opened, there are several options to protect them against malware:

- If the enterprise can enforce application settings on both the client and the server hosts, ports can be configured differently than the well-known numbers targeted by worms.
- VPNs can restrict the set of permitted client hosts.
- Within the control network further firewalls and personal firewalls can be used to create containment zones.
- Non-essential services with known weaknesses should only be offered on one of several redundant servers to reduce the effect of an infection.

An additional measure to avoid attacks through an externally accessible

service is to install all updates that remove known security vulnerabilities in that service. This is not necessarily the best method because it requires a considerable effort to install each patch. Also, it is advisable to defer installation until the automation system supplier has tested and qualified the patches. Despite vendor testing at a reference installation, it can not be guaranteed that the patch will not interfere with a specific system. Most importantly, patching can not prevent "zero-day attacks", ie, attacks via vulnerabilities which are unknown until the attack occurs. Because of this, any anti-malware policy should not rely on patching alone.

How to detect an infection?

Even before it affects the proper functioning of the automation system, it may be possible to discover the infection and thus initiate countermeasures before any damage is done. Available means of doing this include host-based anti-virus tools, network and host intrusion detection systems, general network health monitoring tools that are perhaps even integrated into the control system HMI, 3 [6], as well as relatively recent technologies to detect specific worm activity in a network [7]. However, certain hosts, directories, or applications may have to be excluded from virus scanning and intrusion monitoring for performance reasons.

In any case, such mechanisms will only be useful if there is somebody who continuously monitors these systems and is trained and equipped with tools to respond immediately to contain the infection and recover from it.

How to recover from an infection?

To plan and execute an appropriate response to an infection, certain information must be collected:

- Which functions/hosts are needed most urgently and have to be recovered with highest priority?
- What down time is acceptable for the different system functions?
- Which persons have to be alerted? How quickly are they available? Who may replace whom?

Based on information about functions, interdependencies, and data flows, it may be possible to design the system architecture such that it has predefined isolation points. If the network is cut at those points, the remaining islands are at least partially functional.

Redundant functions can be distributed over several of these containment zones.

The following technical means support the rapid recovery of plant control:

- A warm standby backup host isolated from the network for functions that allow basically no outage.
- Recent and complete system backups on swappable and bootable hard disks for the immediate restart of infected systems after the infected hard disk has been removed.
- Recent and complete system backups on other media ("ghost images") for the immediate reinstallation of affected hosts.
- Local, non networked backup control mechanisms, such as manual valves as well as procedures on how to man and operate these.
- Communication means for emergency staff that are independent of the IT network, e.g. two-way radio.
- Internet access from a separate host independent of the enterprise network, eg, via ADSL or dial-up to a different provider, so that information and updates from the Internet can be obtained even while the main network is down.

Once these technical measures are in place, the following sequence of steps should be adhered to when an infection occurs:

- Isolate hosts that are known to be infected.
- Separate the different containment zones at the prepared isolation points.
- Identify and remove the vector of the infection (eg, connection to outside network).
- Once the spread of malware is stopped, connect standby backup hosts that were kept offline until that point.
- Identify the malware.
- Restore the infected hosts according to the prepared strategy (hard disk, ghost image, rebuild from the original media, disinfection). Follow expert guidance on how to remove the

2 Prototype of a tool developed in ABB Corporate Research for securely importing files into a control network.

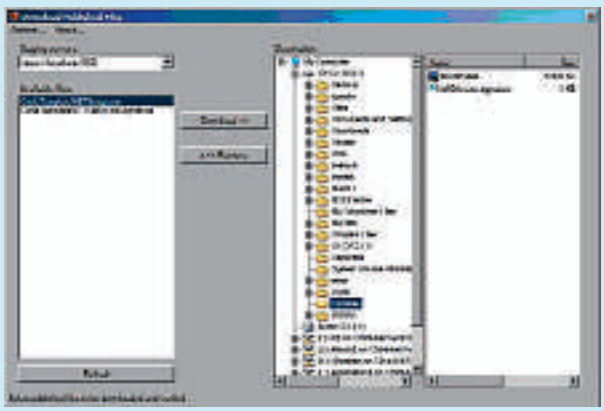

3 Prototype of a tool developed in ABB Corporate Research to alert the process control system operator to anomalies in network behavior, e.g. caused by a worm.

identified malware. If in doubt, it is better to rebuild the system – to ensure that all pieces of the infection are really removed – rather than to rely on malware removal tools.
- Take backup hosts offline again.
- Reconnect restored hosts.
- Take measures, like changing certain procedures, to prevent the reoccurrence of the same infection.
- Reconnect to the outside network.

Case study

This case study describes the malware protection strategy for a slightly simplified version of a real ABB customer automation system. The plant control system consists of an HMI and a performance data collection and reporting subsystem.

The customer wanted access to the reporting server from PC clients in his intranet, which is also connected to the Internet. The main security concern for the customer is the threat of worm infections, especially via the reporting service which must be exposed through the firewall.

The following description highlights some of the more interesting technical malware protection mechanisms suggested for this situation 4. For brevity, procedural and standard technical measures are not described.

The operations part of the control system is isolated from the rest of the network. Any files imported into this part via portable media must first be scanned at the anti-virus station. The anti-virus station is located in the control room and is only connected to the intranet. It has a personal firewall and can only "pull" updates from the net into the anti-virus station. The reporting system is connected to the controller network via its own server. The controller network uses a non-TCP/IP/Ethernet fieldbus and is therefore unlikely to be a worm vector.

A firewall is used to block requests from the intranet to all ports on the reporting server except those needed to obtain performance data. To further reduce the attack surface, only a subset of the PCs on the intranet – those that are supposed to run the performance reporting client application – are allowed to send data to the open ports in the firewall. This is achieved using an IPSec based VPN between these hosts and the firewall. In this case, the VPN is mainly used to ensure host authentication, not data confidentiality.

Summary

Protection against malware is mostly a "people issue". Clear policies explaining what is permitted, procedures to ensure that permitted actions are executed correctly, and training to make users aware of policies, procedures, and the reasons behind them are important factors that help protect a system against malware infections. A policy with respect to data flows between the automation system and the outside, for example, should be defined such that many of the infection vectors previously mentioned are automatically disabled, even if this causes some inconvenience for users. Technical means such as firewalls, virus scanners, content filters, and intrusion detection systems are available to support and enforce procedures.

Martin Naedele
ABB Switzerland, Corporate Research
martin.naedele@ch.abb.com

Rolf Vahldieck
ABB Automation Products, Germany
rolf.vahldieck@de.abb.com

4 Network diagram of a malware-protected automation system.

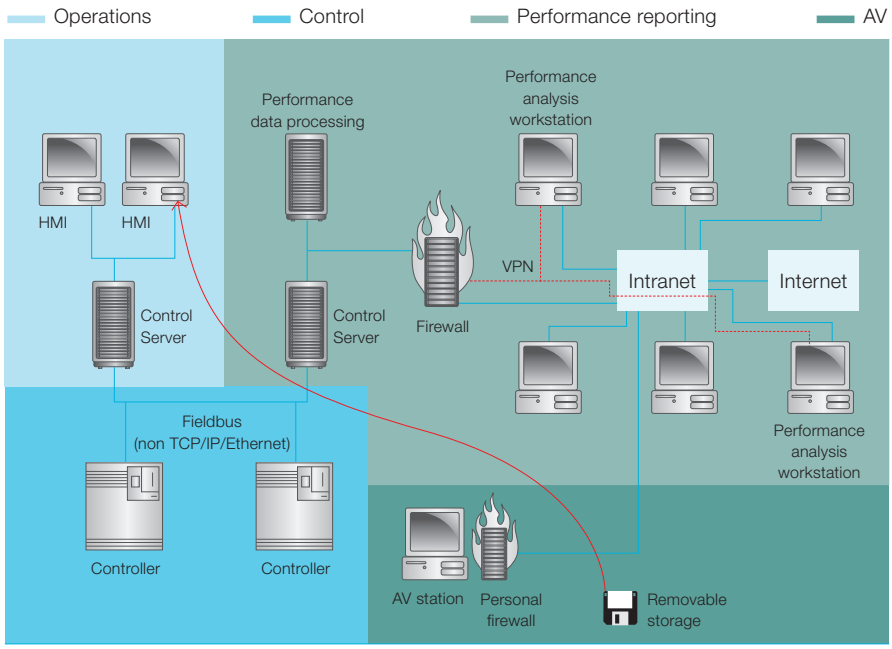

References

[1] Naedele, M. , Dzung, D.: IT security in industrial plants – an introduction, ABB Review 2/2005.
[2] T. Chen, J-M. Robert, "Worm epidemics in high-speed networks," IEEE Computer, June 2004.
[3] M. Naedele: Sicherheitsstrategien für automatisierte Produktionssysteme, in: D. Burgartz, R. Röhrig [Eds.] Information Security Management, TÜV Verlag, to be published in 2005.
[4] Naedele, M., Biderbost, O.: Human-Assisted Intrusion Detection for Process Control Systems, 2nd Int. Conf. on Applied Cryptography and Network Security (ACNS), Tunxi/Huangshan, China, June 2004.
[5] Riordan, J., Zamboni, D.: Billy Goat Detects Worms and Viruses, ERCIM News No. 56 January 2004, http://www.ercim.org/publication/Ercim_News/enw56/riordan.html.
[6] N. Weaver, V. Paxson, S. Staniford and R. Cunningham, A Taxonomy of Computer Worms, Proc. ACM CCS Workshop on Rapid Malcode, October 2003.
[7] D. Moore, V. Paxson, S. Savage, C. Shannon, S. Staniford and N. Weaver, Inside the Slammer Worm, Security and Privacy, July/August 2003.

Further reading

Richard Harrison: The Antivirus Defense-in-Depth Guide, Microsoft, 2004

Industrial information system security Part 3

Standards for securing industrial automation systems

Martin Naedele, Dick Oyen

Part 2 of this three-part tutorial on information system security in industrial networks explained the different types of malware and suggested how an automation system could be defended against them.

This final installment looks at various initiatives that have been started over the last couple of years by different groups to create standards and other forms of guidance to secure industrial automation systems. An overview of a number of those initiatives and their work products is presented, and the approach taken by IEC TC65 WG10 to produce technical blueprints for securing certain control system scenarios is explained.

Currently there is a flood of information available on information system security in general. On top of this, there are some automation vendor white papers that explain certain aspects of locking down their systems. However, there is a serious lack of impartial and easily accessible guidance on how to systematically secure automation and control systems against electronic attacks.

There is no doubt that standards in this area would be very beneficial for both automation users and automation vendors, thus enabling them to:

- Estimate the effort required to implement, maintain, and operate security mechanisms and processes.
- Specify the security objectives for their plant and the security functionalities and measures that have to be provided by vendors and system integrators.
- Compare the threat coverage and cost of offered security solutions.
- Implement and operate cost efficient security mechanisms across multiple plants and locations in the enterprise.
- Be sure that their defense against IT-based threats corresponds to state-of-the-art solutions.

Vendors and system integrators will be able to:

- Anticipate security requirements and develop corresponding functionality.
- Create security architectures that may be reused across multiple projects and customers. This reduces costs in proposal writing, engineering, and the purchasing of third party security devices and applications.

Major standardization initiatives

The following survey of industrial security standardization initiatives is adapted from [1] and [2].

ISA S99

The intention of the ISA (Instrumentation, Systems, and Automation Society) Committee SP99, "Manufacturing and Control Systems Security"[1] is to create guidance documents and a standard (S99) on introducing IT security to existing industrial control and automation systems.

ISA is entitled to produce standards for the process industry with national validity in the US. Many ISA standards are used internationally as best practices or, such as S88 and S95, adopted as international standards.

Standards outlining how to systematically secure automation and control systems against electronic attacks would be very beneficial for both automation users and automation vendors.

SP99 started its work in 2002. As a first step, it produced two technical reports that were published in spring 2004. The first report "Security Technologies for Manufacturing and Control Systems" [3] is a comprehensive survey of what is state-of-the-art in security technologies and mechanisms, with comments on their applicability for the plant floor. It covers: authentication and authorization; filtering/blocking/access control; encryption and data validation; audit, measurement, monitoring, and detection; operating systems and software; and physical security. Each technology is evaluated with regard to the following questions: Addressed security vulnerabilities; typical deployment; known weaknesses; use in an automation environment; future directions; recommendations; and references.

The second report "Integrating Electronic Security into the Manufacturing and Control Systems Environment" [4] presents recommendations for a security architecture and describes the administrative issues and processes for introducing a security management system in industrial plants. The approach in this report is inspired by ISO/IEC 17799 [5]. It contains sections on developing a security program, policies, risk assessment, audits and testing, developing, selecting, and procuring, countermeasures, as well as examples for policies and forms.

Since the summer of 2004 SP99 has been working on the S99 standard. S99 focuses on:

- Retrofitting security mechanisms in existing plants with commercially available components without actually prescribing a specific architecture.
- The processes to operate the underlying management system and administrative processes.

The actual security architecture and processes will likely be customized for specific plants.

IAONA

The Industrial Automation Open Networking Alliance (IAONA) is an interest group of industrial communication system users and manufacturers. Its Joint Technical Working Group Security[2] has developed a Security Data Sheet which is intended to serve as a template for automation system and

Table 1 Secutity initiatives

CIDX (http://www.cidx.org/CyberSecurity/) creates procedural security guidance for the chemical industry. Its work is aligned with ISA SP99. CIDX is mostly active in North America.
NAMUR (http://www.namur.de/en/694.php) provides guidance on secure usage of networking technology for the process industry. NAMUR is mostly active in Germany/Europe.
NERC (http://www.nerc.com/) is the North American self-regulation authority for power utilities. Compliance with NERC 1200 and successor CIP 002…009 standards on security management with their strong focus on processes and documentation is compulsory for North American power utilities.
CIGRE (http://www.cigre.org), the International Council on Large Electric Systems addresses IT security considerations in a number of its working groups.
PCSRF (http://www.isd.mel.nist.gov/projects/processcontrol/), the Process Control Security Requirements Forum, promotes security certification of future control system components according to ISO/IEC 15408 ("Common Criteria"). It is driven by the US National Institute of Standards and Technology, the US national ISO/IEC 15408 certification authority.
PCSF (http://www.pcsforum.org/), the Process Control System Forum, was established 2004 as a meta-initiative to promote information sharing between all the other initiatives on the topic.

device vendors to document the security and communication related features and requirements of their individual products. This information can serve as valuable input for the automation security architect as he designs and configures the necessary security mechanisms for the plant. The benefit of such a Security Data Sheet is that it collects, at a single location, concise security relevant information that is otherwise often hard to obtain from vendor literature.

IEC

In early 2004 the IEC Technical SubCommittee 65C (Digital Communications), through its working group WG13 (Cyber Security), started to address security issues - within the IEC 61784 standard – for field buses and other industrial communication networks. These issues are outlined in a new part 4 entitled "Digital data communications for measurement and control – Profiles for secure communications in industrial networks".

What became evident during this work was that security issues in the automation system cannot be solved by protecting communication alone and by looking only at the field level. Instead, the working group started to specify state-of-the-art secure realizations of certain common automation networking scenarios, such as dial-up remote access. These descriptions, called requirement sets, contain a product independent specification of technical mechanisms in the context of a best-practice security architecture, as well as guidance on the configuration and operation of these mechanisms. The approach is described in greater detail below.

Consequently, the work of the group was moved to TC65 WG10 to align the actual and necessary work with the IEC committee mandate. The completed standard IEC 62443, entitled "Security for industrial process measurement and control – Network and system security" is expected in 2006. The final voting for international validity will take place during the first half of 2007.

Some other security initiatives are briefly described in Table 1.

Security management on the plant floor according to ISA S99

The ISA SP99 technical report TR99.00.01 [3], "Security Technologies for Manufacturing and Control Systems" provides guidance on the applicability of a broad and inclusive range of security technologies. Its advice comes from the combined experience of security experts from automation system vendors and users. As the information presented is analytical in nature, it is not a normative standard against which compliance can be measured. The reader determines the applicability of the information to the specific case. It is an excellent document for those starting to determine security measures and those with experience. TR99.00.01 continues to be updated but its content will not be covered by the S99 standards.

1 Cyber Security Management System.

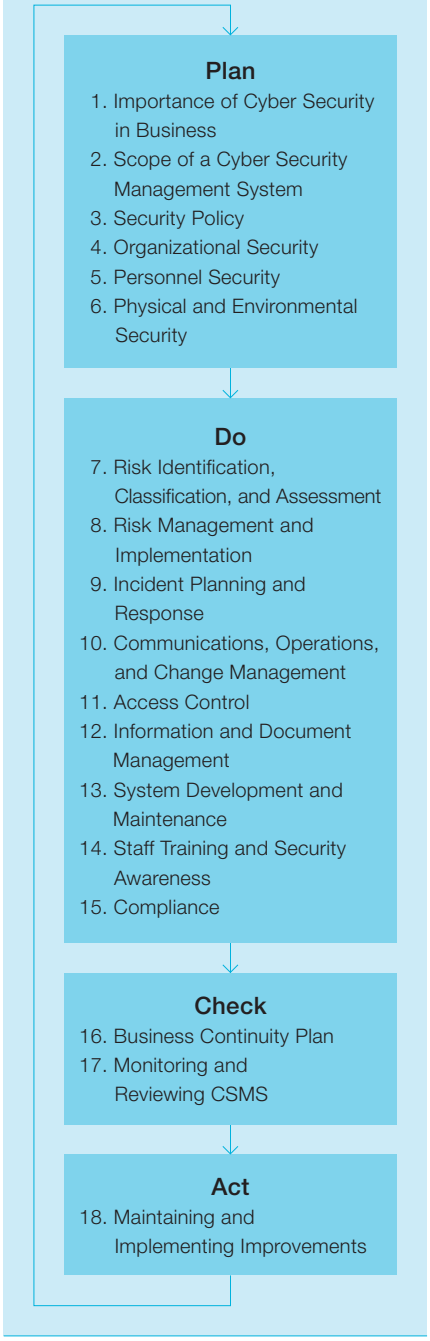

As of October of this year, drafts of two of the four parts of the S99 standard are almost ready for public review.

Part 1 defines terms and describes the models used in discussing security in automation systems. Part 2 advises how a cyber-security management system (CSMS) can be established. There are 18 key elements in a CSMS which are structured in a life cycle that is constantly repeated through four phases: Plan, Do, Check, and Act. The CSMS is provided by the Chemical Industry Data Exchange (CiDX) [7], which adapts the four phases of the British Standard BS 7799-2:2002 [6] to automation systems and defines the 18 key elements. The CSMS and its elements are shown in 1. Its cyclic nature is implicit in step 18 in which the security program itself is modified according to lessons learned in the course of the preceding elements.

Plan phase:

Security planning begins with making a business case so that top management can set a clear top-level policy that mandates the security program.

Organizational Security is planned and this takes into account all of the departments and people that are involved with the control system. It identifies roles and establishes responsibilities relative to security.

Security relates to people: those who have assets to protect, those who are expected to protect them, and those who might compromise those assets. Personnel Security defines personnel policies to estimate and maintain the trustworthiness of those who are given greater access to the assets. Physical and Environmental Security must also be planned. Cyber security is

based on an assumption that there are substantial (not absolute) barriers against physical attack.

Security risks are identified, classified, and assessed in the planning phase. Detailed instructions about how to do this is provided in S99 and the material that it references.

Do phase

Risk assessment leads directly into the Do Phase. Using the risk assessment, security resources can be efficiently applied to real vulnerabilities.

Procedures are established that plan the response to potential incidents. Response planning must include when it becomes necessary to notify government officials of a significant threat to the community.

Overall management policies and procedures are established to cover communications, system operations, and change management.

Access control defines the privileges that accompany specific roles. It also defines the procedures that limit people's access to activities and information to which they are privileged. Authentication means are determined which will ensure that a particular user (person or software) has the necessary access authorization.

Information and Document Management identifies the security classification of data and specifies safeguards. Security issues of developing and maintaining the system are also handled by policies and procedures.

Staff must be trained in the relevant security procedures and all personnel should undertake regular refresher courses on general security precautions. Compliance of departments and personnel to the security policies and procedures must be measured through continuous monitoring, and periodically, through audits. Compliance must also take into account external requirements such as those of customers, contractual partners, and regulatory agencies.

Check phase

The Check Phase includes developing a Business Continuity Plan and following it. This plan establishes how the company will operate through incidents that result in serious damage, plant outage, and possibly community catastrophe. The lessons learned from the CSMS activities are reviewed in the Check Phase.

Act phase

As the CSMS is a cyclic process, the security program is revised according to the review in the Check Phase. S99 Part 2 provides more details as well as 19 steps for establishing a

2 Modular security architecture. Each can be mapped to certain components of the automation system and its network.

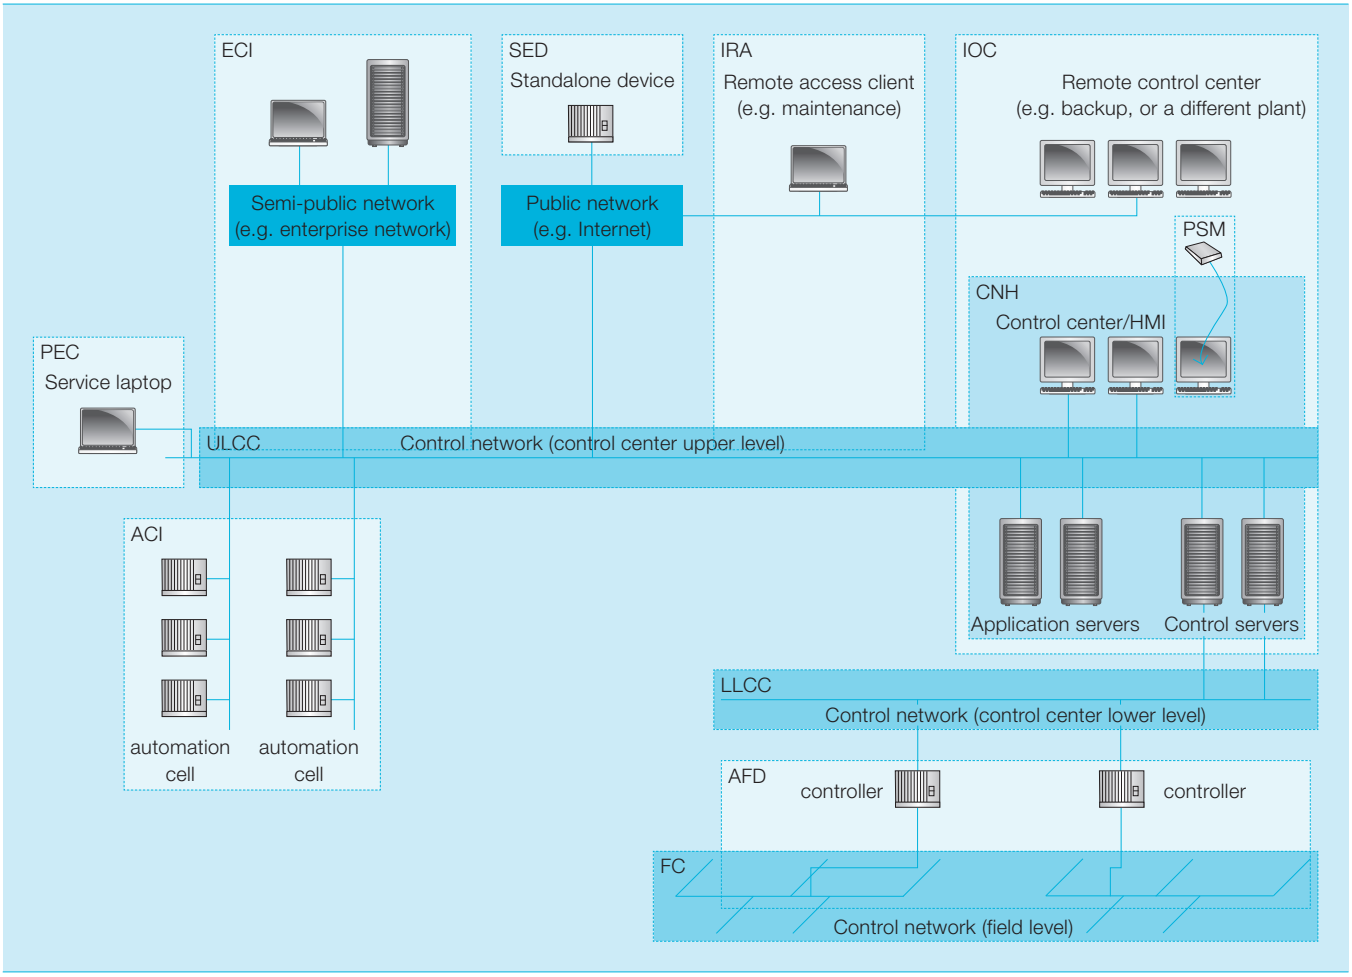

CSMS. Part 3 aims at providing guidance on how to operate the CSMS.

Technical security architecture based on IEC 62443

IEC 62443 mainly addresses – on a system level – technical aspects of the security architecture, and thus compliments product oriented initiatives like IAONA, and process guidance provided by SP99 and NERC.

With the ongoing standardization efforts for Industrial IT security processes and architectures, plant managers have a real chance to implement state-of-the-art and cost-efficient information system security.

The basic idea of the IEC approach is that of a modular security architecture. Each module corresponds to a certain usage or communication scenario and can be mapped to certain components of the automation system and its network 2. Each module is represented by a requirement set specified in the standard. Some of the requirements, as well as the physical or logical components they refer to, are common to multiple modules. Security architecture modules can and should be combined to suit the specific usage and threat situation of an automation system. The standard will provide guidance on the priority of modules for situations where a complete implementation of the standard is not possible due to budget limitations for initial implementation and ongoing maintenance.

The requirements will be formulated in a way that can be used as the basis for Requests for Proposals (RFPs) for data communication standards, and offers, as well as security audits. They should, at the same time, allow for different technical solutions. One goal is that it will be possible to meet the requirements of the standard using products and technologies that are commercially available today. The requirements can also be applied to current and legacy systems and they can be scaled down for systems where an analysis has indicated they represent a low risk for both the enterprise and society.

The working group foresees the following modules:

Enterprise – control net interconnect (ECI): ECI defines the security architecture for non-real-time dataflow between a control network and an enterprise network, preferably unidirectional out of the control network.

Interactive remote access (IRA): IRA details the security architecture needed so that parts of the control system can be accessed remotely (ie, via telephone dialup or Internet) for perhaps engineering or expert diagnosis.

Inter control center connect (ICC): ICC describes how communications between fixed control centers over public networks can be secured.

Stand-alone embedded device (SED): SED outlines the security requirements for an automation device that is not contained in a security zone and for which a full-blown security perimeter would not be cost efficient, eg, a pole-top Intelligent Electronic Device (IED).

Portable engineering computer (PEC): PEC details how a control system can be protected against threats originating from portable computers that may be moved back and forth between public networks and the control system

Portable storage media (PSM): An automation system may be exposed to malware infections through storage media like memory sticks or CDs. PSM explains how this can be prevented.

Automation cell interconnect (ACI): ACI outlines the security architecture required for protected communication between automation cells within a control network.

Upper Level Control center (ULCC): Part of a control network is connected to operator workstations, “historians ”, application servers and connectivity servers. ULCC datails network oriented security mechanisms specific to this part.

Lower Level Control center (LLCC): LLCC outlines network oriented security mechanisms in the part of the control network connected to controllers and PLCs.

Field Control (FC): FC outlines network oriented security mechanisms in the part of the control network connected to field devices.

Control network host (CNH): CNH explains how automation workstations and servers for operations and engineering can be secured against attacks from insiders and malware, for example.

Automation field device (AFD): AFD explains how field devices and embedded controllers can be secured.

Each module describes: a use case to which it applies; threats that are addressed or not addressed; the underlying assumptions; the requirements; and the party (automation vendor, system integrator, or plant owner) responsible for meeting each of the requirements. The core part of each module is the requirement set and it contains between 20 and 50 requirements, depending on the module.

Each requirement consists of a normative statement, optionally including scale-down alternatives, a rationale, and in many cases one or more application notes. The rationale is an essential element, as it enables the reader to make an informed decision about the importance and applicability of the requirement. The application notes provide technical guidance on how the requirement could be realized.

The IEC 62443 standard describes the “what” and “why” of the security architecture, but the “how” is specific to an individual site and system and is therefore left to the engineering judgment of the plant experts and the automation/IT integrator.

Summary

With the ongoing standardization efforts for Industrial IT security processes and architectures, specific to control and automation systems, plants managers now have a real chance to

Footnotes

1) http://www.isa.org/MSTemplate.cfm?MicrositeID=988&CommitteeID=6821

2) http://www.iaona.org/home/jtwg-se.php

Tutorial

implement state-of-the-art and cost-efficient information system security.

The standardization initiatives described above have so far been characterized by a general recognition that pragmatic solutions are needed to serve the industry, as well as very constructive collaboration among automation vendors and end users so that this objective is achieved.

ABB is a major contributor to various security standardization initiatives. The company offers products and solutions that are compliant to evolving standards, and provides assistance to its customers in applying these standards to specific plants and sites.

Martin Naedele
ABB Switzerland, Corporate Research
martin.naedele@ch.abb.com

Dick Oyen
ABB US, Corporate Research
dick.oyen@us.abb.com

References

[1] Naedele, M.: Standardizing Industrial IT Security – A First Look at the IEC approach, 10th IEEE International Conference on Emerging Technologies and Factory Automation (ETFA 05), Catania, September 2005

[2] Dzung, D., Naedele, M., von Hoff, T., Crevatin, M.: Security for industrial communication systems, Proceedings of the IEEE, Vol. 93 (6), June 2005, pp 1152–1177

[3] ISA SP99: Security Technologies for Manufacturing and Control Systems, Instrumentation, Systems, and Automation Society, ISA-TR99.00.01-2004, March 2004,

[4] ISA SP99: Integrating Electronic Security into the Manufacturing and Control Systems Environment, Instrumentation, Systems, and Automation Society, ISA-TR99.00.02-2004, April 2004,

[5] ISO: Information technology – Code of practice for information security management, ISO/IEC 17799:2000, December 2000,

[6] British Standards Organization: Information security management systems – Specification with guidance for use, BS 7799-2:2002, September 2002

[7] Chemical Industry Data Exchange (CiDX): Guidance for Addressing Cybersecurity in the Chemical Sector, Version 2.0, December 2004.

INDEX 2005

1/2005: Pioneering spirits

3/2005: Sustainability

2/2005: University and industry cooperation

4/2005: Innovation – The DNA of business