SECURITY MATTERS

Gartner 2014 CoolVendor

Monitoring Industrial Control Systems to improve operations and security

An overview of the threats to Industrial Control Systems and the technologies to protect them

Authors

Dr. Sandro Etalle, sandro.etalle@secmatters.com
Dr. Clifford Gregory, cliff.gregory@secmatters.com
Dr. Damiano Bolzoni, damiano.bolzoni@secmatters.com
Dr. Emmanuele Zambon, emmanuele.zambon@secmatters.com
Dr. Daniel Trivellato, daniel.trivellato@secmatters.com

01 December 2013

Version 1.0

Contents

1 The need for monitoring

Industrial Control Systems (ICS) sit at the core of every industrial process - from power generation to water treatment and manufacturing. The term ICS refers to the set of devices that govern the process to guarantee safe and successful execution, and include Supervisory Control and Data Acquisition (SCADA) systems, Distributed Control Systems (DCS), and other control systems such as Remote Terminal Units (RTU) and Programmable Logic Controllers (PLC). A malfunction in any of these systems might cause the entire industrial process to fail, with serious consequences in terms of economic loss and in compromised public safety. For instance, consider disruption of an electricity transmission network; an incorrect distribution of power might affect availability to households, offices, hospitals, etc. Similarly, a faulty component that regulates the amount of chemical substances in a pharmaceutical production process might lead to entire batches of harmful compounds.

The need to monitor ICS and more generally industrial networks is advocated in many venues and has been included in several recommendations, guidelines, and standards, such as the Electricity Subsector Cybersecurity Capability Maturity Model (ES-C2M2); the North American Electric Reliability Corporation Critical Infrastructure Protection (NERC CIP); and the still in progress M/490 Directive from the European Union. There are several reasons why monitoring should be an integral part of operations, and may even be considered a competitive advantage. The marked increase of cyberattacks directed against ICS, as frequently reported by the ICS Cyber Emergency Response Team (ICS-CERT) of the U.S. Department of Homeland Security, is only one of them. Stuxnet, Duqu, Flame, and Gauss have seized all the attention of the media, but represent only a small portion of the cyberthreats targeting ICS.

Today, organizations rely heavily on third parties (i.e. contractors, vendors, etc.) to streamline their operations. Smaller organizations, such as local utilities, have consultants programming and maintaining their PLCs, while larger organizations have direct lines connecting their process control devices to vendors for 24/7 support and maintenance. In all cases connectivity comes with marvelous gains in productivity, but adds a totally new dimension of risks which must be mitigated to secure ICS networks. The shift in how industrial networks are operated and maintained makes it crucial for organizations to monitor their infrastructure not only against cyberattacks, but also against system misuse by third parties. Next to this, there is an ever-present threat coming from insiders, including disgruntled employees, human error, and network and system misconfiguration, which could affect or lead to disruption of the production process. All these threats strongly impact the bottom line revenue of an organization, and as such should receive attention of the boardroom daily.

Organizations can mitigate these threats by implementing a solid monitoring infrastructure for ICS and back-office networks. With an effective monitoring infrastructure in place, organizations are not only able detect problems at an earlier stage, but can also mitigate consequences before any real damage is done and recover more quickly from an incident, no matter the nature. Continuous monitoring and early analysis of identified issues help organizations pinpoint root causes of a problem and enforce effective countermeasures and remedial actions.

Before implementing any monitoring infrastructure or even considering a specific solution, however, organizations should assess their exposure to the risks and threats to their processes and devices. Quite often organizations overlook this phase of the security process, and rush into picking a solution that does not match their needs or expectations. This document is not intended to provide details about risk assessment procedures. Nevertheless, it is important to understand its key concepts and identify the steps that lead to the selection of appropriate monitoring solutions.

1.1 Assessing the Risk

The security process within an organization is an unending cycle that aims at continuously improving the security posture of the organization. In order to determine the next steps, estimate effort and measure improvement (or regression) over time, security personnel need to assess the organization's current security status in a number of areas.

A typical risk management process consists of three phases: risk analysis, risk evaluation, and risk mitigation. During the risk analysis phase, an organization must identify its key assets and their vulnerabilities, as well as the related threats and threat sources. Threat sources can be either external to an organization, such as foreign intelligence agencies, cyberterrorists, and hacktivists, or internal, such as poorly trained or disgruntled employees, contractors, and vendors. It is unlikely that a foreign intelligence agency would specifically target a local utility organization, whereas a multinational company might very well be targeted by cybercrime. Each adversary has different characteristics in terms of funding and capability to exploit an organization's critical assets.

To effectively analyze risk, it is important to carefully assess impact and likelihood of each threat being exercised against an asset's specific vulnerability. While assessing the likelihood of a threat is generally easier in most IT environments, it is nearly impossible in ICS networks, due to the scarce historical data available regarding past incidents. The best way available to proceed is thus to take into consideration high-impact threats and determine the cost and effort needed to recover productivity. Threats are then ranked based on their overall cost and alleged likelihood during the risk evaluation phase. The result of this phase is a prioritized list of threats to be mitigated.

Finally, an organization must decide "what to do" with the identified threats, i.e. whether they can be avoided, should be accepted, or might be reduced by means of adequate controls. In traditional risk management, there are several types of control that can be applied to reduce a risk, including: legal, deterrent, preventive, monitoring, detective, corrective, etc. The implementation of each of these controls might involve different elements of an organization, such as people (e.g. good practices), the physical world (e.g. access control), procedures (e.g. change management, disaster recovery), and technology (e.g. antivirus, intrusion detection/prevention systems). The selection of the right (combination of) controls should aim at minimizing the likelihood and impact of the identified threats, thereby maximizing the benefit for the organization.

In the rest of the whitepaper we discuss the most prominent cyberthreats to ICS and their possible impact on an organization; we also give an overview of the existing technologies that can be employed to mitigate them.

2 Threat and vulnerability analysis

A threat and vulnerability analysis is a mandatory step that should precede *all* security-related decisions, including decisions regarding which procedures to adopt and which prevention and monitoring solutions to employ. Unfortunately, more often than not, this step is skipped altogether.

In this section we analyze the possible threats to ICS. The threat analysis is performed in three steps: first, we discuss the major sources of threat to ICS (adversaries) and their capabilities; then, we characterize ICS networks, pinpointing the key components and their vulnerabilities; lastly, we present a number of example methods that adversaries can use to disrupt or sabotage the victim's industrial process. The identification and selection of the most prominent threats to an organization is instrumental in determining the most appropriate countermeasure to be employed for mitigation.

2.1 Threat sources

Threats to ICS can come from several sources, including hostile governments, terrorist groups, competitors, malicious intruders, and disgruntled or careless employees [4, 7]. These adversaries can be classified in many ways. For simplicity, here we consider the following two categories:

- Outsiders: adversaries who are not associated with the organization and acting from the outmost perimeter.
- Insiders: adversaries who are already within the target organization's perimeter.

Outsiders include governments and foreign intelligence services, hackers and hacktivists, industrial spies, (cyber-) terrorists, and organized crime. Insiders include disgruntled employees, careless or poorly trained employees, contractors and vendors.

Outsiders have varying degree of knowledge and motives. They may perform attacks for information gathering, espionage activities, or to disrupt or weaken the target's processes. They are usually motivated by political, monetary, or reputation reasons.

While outsiders typically have higher funds and resources, the insider threat should not be underestimated by organizations. Joe Weiss writes in his book "Protecting ICS from electronic threat" [8] that insiders might be the biggest concern nowadays, due to the specific knowledge of and access to the organization's infrastructure. The NIST Special Publication 800-82 [7] provides an accurate description of the insider threat for ICS: "*The disgruntled insider is a principal source of computer crime. Insiders may not need a great deal of knowledge about computer intrusions because their knowledge of a target system often allows them to gain unrestricted access to cause damage to the system or to steal system data. The insider threat also includes outsourcing vendors as well as employees who accidentally introduce malware into systems. [. . .]Impacts have ranged from trivial to significant damage to the ICS and field devices. Unintentional impacts from insiders are some of the highest probability occurrences.*" In the rest of the section we describe the different types of insiders, their capabilities and intent.

Disgruntled employees This adversary includes both current and former employees that are upset or dissatisfied with the organization. Their motivation can also be fame, greed, divided loyalty, delusion or monetary gain [5]. Depending on their level of access to the infrastructure, disgruntles employees can pose a serious threat to the employer. Disgruntled current employees will have legitimate credentials and possibly unrestricted access to the organization's infrastructure. Former employees may still have valid credentials, either because the old access rights have not been revoked, or because the employee secretly created new access credentials before the termination of employment. All current and former employees have process knowledge and may be aware of system vulnerabilities.

Careless or poorly trained employees and unintentional mistakes Employees with a lack of training, concern or attentiveness can pose a serious threat to their organization. This adversary does not have the intention of harming the organization or to cause incidents. In most cases the misuse stems from unintentional mistakes, whose effects on normal operations could remain unnoticed for a long period of time. For example, employees may unknowingly introduce malware to the internal networks by using an infected private storage media or laptop.

Contractors and vendors Contractors and vendors are third parties used by organizations for carrying out specific tasks and providing specialized services. Organizations often depend on the expertise and services from such third parties. Vendors often require remote access to offer additional (maintenance) services, such as updates and diagnostics. While organizations usually limit and control the access rights to the necessary minimum, once a remote connection to field equipment is established any misuse could take place.

2.2 The weak spots

Adversaries aiming at disrupting an organization's production processes will attempt to access, damage, or improperly use its critical assets. An asset is broadly defined as "*anything, which has value to an organization, its business operations and its continuity*" [1]. The ISA 95 standard[1] defines a multi-level model for industrial activities. Each level includes specific assets, provides specialized functions and has characteristic response times.

- Level 0: refers to the actual physical production process. It involves the machinery used for the process.
- Level 1: defines the activities involved in sensing and manipulating the production process. Includes RTUs, PLCs, and field devices such as sensors, actuators, and other I/O devices.
- Level 2: involves supervising, monitoring and governing the physical processes. The systems acting on this level are real-time controls and software, SCADA/DCS servers, and operator workstations (HMIs).

[1] www.isa.org/ISA95

- Level 3: includes activities for managing production workflow. The collection of systems acting on level 3 can be referred to as Manufacturing Operations Management Systems (MOMS).
- Level 4: focuses on managing the business-related activities of the manufacturing operations. Enterprise Resource Planning (ERP) is the primary system. Other systems acting on this level are Product Lifecycle Management (PLM), Customer Relationship Management (CRM), and Human Resource Management (HRM).

Industrial networks are typically segmented to reflect this classification. In particular, activities on level 4 and 3 are carried out in corporate (or back-office) networks, while activities on level 2 and 1 are carried out in Supervisory and Process Control Networks respectively. A simplified example of industrial network and its components is shown in Figure 2.1.

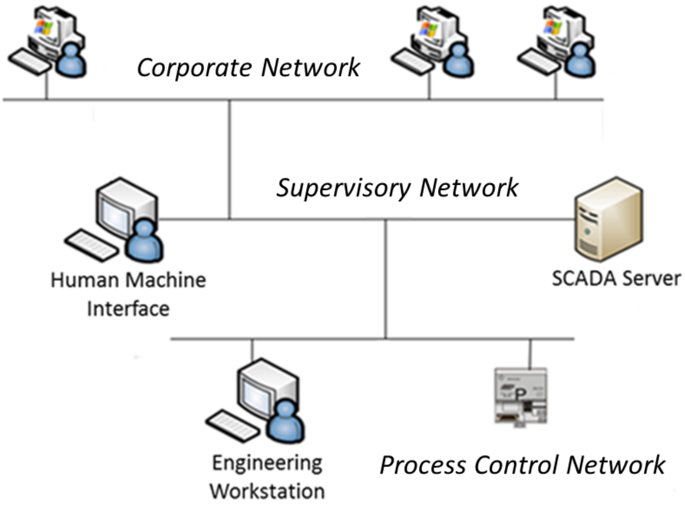

Figure 2.1: Example of industrial network

The most critical assets for an organization operating industrial processes are those located in Supervisory and Process Control Networks, since they provide full control of the victim's production process. Adversaries will attempt to take control of these assets by exploiting one of their vulnerabilities. A vulnerability is a weakness in a component's design, implementation, or operation and management that if exploited can lead to a security breach or an incorrect behavior of the component. Vulnerabilities often result from programming errors such as lack of input validation and memory safety violations. Whereas until a few years ago vulnerabilities and exploits were available almost exclusively for Information Technology (IT) components, i.e. Windows machines and Windows applications, in the last years this trend has inverted its course. Today, vulnerabilities and exploits for Operational Technology (OT) components can be purchased on an open and thriving Internet black market that grows daily.

Figure 2.2 shows the increase in the number of known vulnerabilities for ICS over the last years [2]. These vulnerabilities affect several OT components, with SCADA server and HMI being the most targeted ones. Although these components are essential for a company's production process, they are becoming easier to approach and compromise from both external and internal adversaries.

It is important to notice that the numbers in the graph report only the vulnerabilities that have been discovered and published during the indicated year, and as such they are just the tip of the iceberg. Often, vendors prefer not to disclose vulnerabilities of their components for both safety and commercial reasons.

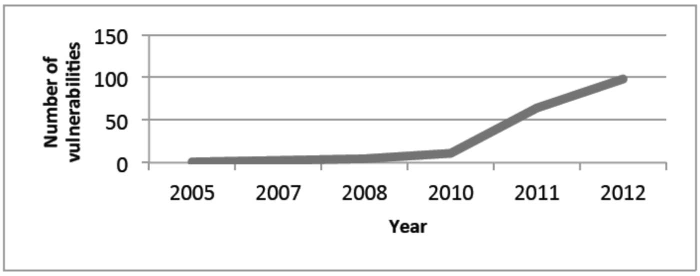

Figure 2.2: Number of known vulnerabilities for ICS

2.3 Example threats

In this section we present a number of threat scenarios involving different types of adversaries introduced in Section 2.1. In particular, we present the possible effects on industrial networks of:

1. Malicious/careless operators
2. An unexpected configuration shift
3. A targeted attack exploiting an unknown vulnerability
4. Spread of malware

Malicious/careless operators Operators and engineers must by the nature of their jobs have a thorough knowledge of all aspects of an industrial process, and are generally holders of "privileged access" to the network resources. In the normal process of their jobs they may modify a system's regular workflow, and they have access to sensitive information that could be passed on to third parties for any number of reasons. Most ICS have the potential of implementing access control mechanisms, but organizations seldom use these controls in practice, both as a measure to save on costs of their re-configuration and to achieve higher productivity and lower downtime when upgrading or updating their network. An employee may disrupt the production process by issuing unusual or abnormal commands, using parameter values that cause PLCs and field devices to crash. A trusted employee may become a threat due to financial gain or dissatisfaction, or simply due to a mistake. Whatever the motivation, the insider can cause catastrophic damage to critical infrastructure, and some have in the past.

An unexpected configuration shift Threats to ICS do not necessarily originate from adversaries, but might arise during the normal course of operations. In fact, network and system configurations tend to change over time, for instance due to on-site maintenance and hardware/software updates. Any organization dealing with critical processes tests network and systems' configuration prior to deployment, and stores a copy (a blueprint) for backup in case of disruption, so that restoring can be as quick as possible. Blueprints should be regularly updated to reflect the changes in running configurations. If for any reason blueprints are not kept up-to-date, a restore procedure could not be effective and cause additional delays and prolong downtime caused by an incident.

A targeted attack exploiting an unknown vulnerability ICS are not generally targeted by what has come to be considered mainstream hackers. The ICS attacker is often motivated, well-funded and may even be a state-sponsored organization, as research has shown for the latest successful ICS intrusions. Attacks to ICS are carefully planned and executed employing specifically developed malware, which circumvents standard security solutions by exploiting previously unknown vulnerabilities in target systems. The aim might be to steal financial or exploration data (as in the case of the Flame malware) or to disrupt operation of the provided service (e.g. Stuxnet). These attacks may be part of a larger cyber-warfare strategy.

Spread of malware One of the main threats to ICS is malware that spreads from one computer to the other, multiplying the damage. The harm resulting from malware spreading around in a corporate network can be huge, and it could become unbearable if the infection manages to reach the Supervisory or Process Control Network. Despite strict network segmentation and corporate policies, Supervisory and Process Control Networks are never truly isolated from the external world. Security experts have pointed out that the use of air gaps (or system isolation) is a "myth". External consultants and support and maintenance personnel are required from time to time to transfer data to and from those networks, for instance with USB sticks, and to carry out system software updates, thwarting network segmentation and isolation policies. Once one of the systems is infected, malware can then spread easily and often unnoticed into sensitive, high-security environments and force long periods of system downtime. A real-world example of this scenario is represented by the Shamoon[2] malware, which hit 30.000 workstations at Saudi Aramco requiring a lengthy restore process.

[2] http://www.symantec.com/connect/blogs/shamoon-attacks/

3 Mitigation solutions

After identifying the major sources of threat to ICS and illustrating some possible threat scenarios, we discuss the controls that can be used to mitigate them. In particular, here we focus on the various existing technologies that can be used for continuous monitoring of ICS networks and the detection of threats.

There exist heterogeneous and complementary monitoring solutions to mitigate threats to ICS. Each of them has pros and cons and there is no silver bullet. Depending on the output of the risk assessment, an organization should select the set of solutions that best suits its needs. Traditionally, monitoring solutions are divided into host- (or endpoint) and network-based. Most organizations employ "regular" (or slightly modified) endpoint security solutions developed for IT environments, such as antivirus and application whitelisting, to protect their SCADA/DCS servers, HMIs and engineering workstations. There is a general consensus that IT endpoint security finds straightforward application in the context of OT. These solutions, however, will not suffice when it comes to monitoring PLCs and RTUs, detecting insiders' misbehavior, and identifying misconfigurations. Network-based solutions become then the only available choice for protecting some of the most critical assets for an industrial process. In the next paragraphs we describe the existing network monitoring technologies, their pros, and their cons.

3.1 Signature-based

Signature-based detection, or blacklisting, is the most common monitoring and detection method. The idea behind this approach is simple and straightforward. Knowledgeable people, like security analysts, make a list of known malicious and suspicious network messages and devise signatures (i.e., byte patterns) to recognize them. Network traffic is then compared in real-time against these signatures; in case of a successful match, the event is reported to the security operator.

Pros: The approach is simple, widely known among the IT security community, and very effective against mainstream threats, such as well-known malware and software exploits. Usually, a signature-based solution can be deployed in a matter of hours. Some attention must be paid to the tuning phase, which is required in order to suppress false warnings that might be triggered by too strict signatures or even by the type of network traffic in a certain OT network. An additional benefit is the option of turning prevention mode on, effectively blocking (known) misbehavior.

Cons: This approach can only detect previously known misbehavior. Skilled attackers can easily circumvent detection, even by just morphing known attack payloads. Every time a new misbehavior or attack is discovered, a corresponding signature must be developed and distributed; it can take weeks before all monitoring systems are updated. In the meantime, attack payloads that have not yet been isolated, analyzed and mapped to a signature can hit without notice. For this reason, malware such as Stuxnet, Duqu, Flame and Shamoon spread across systems silently for lengthy periods of time. Stuxnet leveraged four unknown software vulnerabilities and went undetected for at least a year. Another limitation of signature-based detection is that "benign yet incorrect" network operations will also go unnoticed, unless a specific signature (or check) for it has been defined. At the present state of affairs, it is impossible to have a set of signatures that provides a reasonable degree of protection for ICS networks. More generally, blacklisting is not suitable for the ICS world, where attackers are more motivated to remain silent, and where attacks are targeted to specific systems.

3.2 Rule-based

Rule-based monitoring, also called network whitelisting, consists of identifying and detailing legitimate and acceptable network activities using rules, and blocking or alerting when non-matching behavior is observed. The effectiveness and usability of rule-based solutions largely depend on how *accurate* the underlying analysis of network traffic is. At one extreme, there are *low accuracy* solutions that work like firewalls, controlling whether network traffic matches, for example, specific combinations of IP addresses and TCP ports. At the other extreme it is possible to devise fine-grained rules, based on deep packet inspection. Solutions featuring deep packet inspection are capable of understanding (part of) the underlying protocol used by the application being monitored and of blocking/detecting the invocation of certain undesired message types (*medium accuracy* solutions), or even the use of field values which do not respect pre-determined constraints (*high accuracy* solutions).

The configuration of a rule-based solution is usually performed manually and involves expert personnel that are knowledgeable of the underlying processes. This is because rules are tailored to protect a specific environment, rather than a "general" infrastructure like signature-based solutions do. In other words, while the same signature-based system can be deployed on different networks with minor adjustments, a rule-based system requires ad-hoc configuration for each network it is deployed on.

Pros: Rule-based solutions allow only network communications that have been explicitly whitelisted. As a result, they offer protection from a wide range of malicious events and behavior, including both known and new and unknown threats.

Cons: The higher the accuracy of the solution, the higher the chance that it will block an attack, regardless of whether it is known or not. Accuracy (and thus effectiveness), however, comes at a cost. An important drawback of rule-based solutions is constituted by their high configuration cost, i.e., the cost of detailing all the whitelisting rules. The more accurate the desired level of analysis, the higher the setup cost. Whitelisting all message types and message field values used in a production site is a daunting task that does not fit most budgets or time schedules. This can be only partly mitigated by employing a coarse-grained configuration that fails to take advantage of the full power of deep protocol inspection, resulting in a marked reduction in threat detection. Additional concerns and costs come from the fact that, each time a system (e.g. a PLC) is reconfigured, updated or re-programmed, the rule-based solution that protects it must be likewise (manually) updated.

3.3 Non-signature based

In order to enhance signature-based and ruled-based approaches, computer scientists have been researching for more than a decade on automated techniques and algorithms to detect misbehavior without requiring previous knowledge of it. These approaches are usually called *non-signature based* (or anomaly-based). Today, most non-signature based solutions employ artificial intelligence algorithms, like neural networks, to first learn what are the normal, legitimate "states" of an environment and afterwards raise an alert when they observe anomalous network traffic. These algorithms have been successfully applied to other fields in the past. For instance, closed circuit television (CCTV) systems use artificial intelligence algorithms for recognizing a suspect among groups of people passing through border security.

When applied to network monitoring, non-signature based solutions come in two inherent flavors, depending on whether they carry out qualitative or quantitative analysis. There is also a third, orthogonal category, called "sandbox" analysis. Finally, an innovative technology has recently been proposed, called Deep Protocol Behavior Inspection (DPBI).

Qualitative analysis Qualitative analysis of network traffic is based on the assumption that malicious traffic will look (significantly) different from regular and legitimate traffic. Solutions based on this approach observe network traffic for a given amount of time (the learning phase), build one or several models to describe what is expected to be regular traffic, and alert when the subsequently observed traffic is "too different" from these models.

Pros: In theory, this approach could detect the most harmful and sophisticated attacks, as for instance those aimed at executing arbitrary binary code on targeted assets by exploiting unknown software vulnerabilities, deviating the underlying process by sending offending commands (or parameters).

Cons: Despite some promising results seen in initial studies, serious shortcomings became obvious when this kind of solutions were deployed in real environments. When processing real network traffic, this approach falls short of identifying malicious behavior with sufficient accuracy [3]. The few qualitative solutions available are based on the assumption that it is possible to seize network traffic's intrinsic characteristics by using a generic model. Recent studies [3, 6] prove not only that this assumption is often wrong, but also that these solutions generate an extreme percentage of false alerts even when processing regular attack-free traffic, making them too inaccurate and far too expensive to manage. An additional issue is related to tuning and customization. Tuning is essential in a dynamic environment such as a computer network. Because qualitative solutions mostly work like "black boxes", leveraging complex mathematical and statistical models, it is difficult, if not impossible, for users to tweak working parameters to improve effectiveness.

Quantitative analysis (flow-based) Quantitative solutions build a model of network flows by taking into consideration aggregate characteristics such as number of bytes exchanged between network devices, number of new connections, etc., over a time window. When subsequently observed flows differ from the model, the system raises an alert. The built model is typically automatically adjusted over time to incorporate changes in the network configuration.

Pros: The technology behind quantitative analysis is well established and has been revisited over time to deal with new attacks. Quantitative solutions are usually easy to tune and tweak for enhancing detection accuracy.

Cons: In practice, quantitative solutions can detect only misbehavior that generates spikes in data volumes and network communications, such as a denial of service, horizontal and vertical (port) scans, and brute-force attacks. This represents a severe limitation, because most sophisticated attacks, or even erroneous commands issued by an operator, would remain undetected, as they do not generate a spike in network flows.

Sandbox-based analysis This approach combines host- and network-based analysis techniques. The idea behind it is to capture "interesting" data from network traffic, such as email attachments, binary files, PDFs and Office documents, and open or execute those files inside a controlled environment (the sandbox). By enforcing certain controls over sensitive process, data and configuration files (e.g. the Windows Registry), it is possible to detect when a captured file could harm the targeted assets, for instance by dropping a botnet client.

Pros: This approach is very effective when it comes to detect advanced threats spreading through files, even when exploiting new and unknown software vulnerabilities. It does not require signatures to work, although vendors typically include those to shorten detection time.

Cons: Sandbox analysis is only available for protecting the Windows OS family, and is strongly oriented to malware detection. Hence, an attack carried out by using, for instance, an infected laptop and aiming at exploiting software vulnerability in a PLC would go unnoticed, and so would an erroneous command issued by an operator, because no malicious file is exchanged in the process. Another issue is that attackers are aware of this monitoring technique and started including obfuscation mechanisms to prevent the sandbox from detecting the threat. For instance, a certain malware embedded in a pdf document could trigger only if the system user was named "joe", or the organization's name was "ACME". This is just a simple example, more complex and sophisticated obfuscation techniques have been seen in the wild.

Deep Protocol Behavior Inspection (DPBI) DPBI is an innovative non-signature based technology developed by SecurityMatters, which combines the effectiveness of high accuracy network whitelisting with the ease of use of a self-configuring system. DPBI performs the most in-depth possible inspection of network communications, understanding both syntax and semantics of ICS protocols. Network communications are analyzed down to the values exchanged by network devices, enabling the detection of any data-based threat (e.g. unknown malware spread, illegitimate or unusual commands or values) on the most commonly used ICS protocols, thereby providing protection for all network segments. In addition, the self-configuring technology enables to automatically define the whitelisting rules by observing and learning the normal traffic within a network (communication patterns, protocols, message types, and values exchanged). The result of this learning process is an accurate view of all network activities, which allows to immediately detect any threat to the industrial process.

Pros: The combination of high accuracy network whitelisting and self-configuring technology enables not only catch the most advanced attacks to ICS, but also to provide full situational awareness and continuous monitoring of network and systems' behavior, allowing to detect, e.g., system misconfiguration and configuration changes. The self-configuring technology also allows to minimize re-configuration costs of the monitoring solution following changes to the network and systems' configuration. Finally, DPBI is devised for "passive" network monitoring, i.e. to observe network traffic without blocking communications. This allows for a higher level of detection without putting system performance and availability at risk by erroneously blocking legitimate traffic.

Cons: Despite providing protection from a wider range of threats and pinpointing their root cause with high level of precision, DPBI does not directly "classify" the threats it detects, as a signature-based solution would do. For instance, a signature-based solution could immediately report an "SQL injection" if a byte pattern resembling such attack is observed, whereas DPBI would simply report the suspicious sequence of bytes. This is because detection by DPBI is not obtained by matching network traffic with a database of known malicious payloads (i.e. attack signatures), but rather by looking for abnormal network activities, and reporting what parts of the communications were abnormal and how. Notice, however, that this detection and classification capability of signature-based solutions is limited to attacks for which a signature is devised, whereas DPBI enables detection of any kind of illegitimate behavior.

3.4 Technology comparison

Table 3.1 summarizes the detection capabilities of the different technologies available by showing which of the threats presented in section 2.3 they would be able to detect.

	Malicious/ careless operators	Unexpected configuration shift	Targeted attack	Spread of malware
Signature-based	✗	✗	✗	!
Sandbox-based	✗	✗	!	✓
Low accuracy whitelisting	✗	✗	!	!
Medium accuracy whitelisting	!	!	✓	✓
DPBI/High accuracy whitelisting	✓	✓	✓	✓

Table 3.1: Detection capabilities of the existing network monitoring technologies
Key: ✗= No detection; != Partial detection; ✓= Full detection

Within the table, "Partial detection" means that the solution might detect the threat depending on the way it is implemented. For example, signature-based solutions would detect malware spread if the malware was exploiting known vulnerabilities or a signature for the malware had already been devised.

4 Conclusions

ICS networks sit at the core of every industrial process. Recommendations, standards and guidelines increasingly advocate the importance of implementing an adequate infrastructure for continuously and accurately monitoring the activities in these networks, to protect them from the growing number and range of cyberthreats they face.

Threats to a company's industrial processes can originate both in the corporate network and in the Supervisory and Process Control Networks. Threats originating in the corporate network are mainly represented by malware that attempts to spread to "lower" network segments to compromise their components (e.g. SCADA servers and PLCs) and disrupt the production process. The most notable and imminent threats, however, are those originating in systems located within the Supervisory and Process Control Network, as these systems have direct control of field devices, sensors, and actuators. These threats include advanced malware (e.g. downloaded from USB sticks or through malicious software/firmware patches and updates[1]), direct data injection attacks and system misuse carried out by disgruntled employees, contractors and vendors, but also unintentional mistakes by operators.

There are a few existing technologies that can be employed to monitor and protect ICS networks. These technologies are complementary, as they are devised to detect different types of threat. Signature-based solutions and sandbox analysis systems are designed to protect corporate networks from known threats and unknown malware respectively, and to prevent the spreading of those threats to other network segments, but leave Supervisory and Process Control Networks mostly unprotected and susceptible to all the threats originating inside those networks. Whitelisting solutions on the other hand are the ideal choice to protect Supervisory and Process Control Networks, as they provide protection from any activity that is not explicitly denoted as legitimate and authorized.

Figure 4.1 shows an example ICS network implementing a comprehensive monitoring infrastructure to achieve the best defense-in-depth. The infrastructure consists of:

- A signature-based solution to stop known threats originating in the corporate network
- A sandbox analysis system to detect unknown malware and prevent it from spreading within the corporate network and to the Supervisory and Process Control Network

[1] http://www.informationweek.com/security/cybercrime/flame-hits-windows-update-7-key-facts/240001490

- A whitelisting solution to protect OT components from advanced threats and detect any undesired network operation
- Flow-based analysis to protect field devices from denial of service and brute-force attacks
- A SIM/SIEM to integrate and correlate the input of all these solution and provide monitoring personnel with a unified view of the current security status of an organization

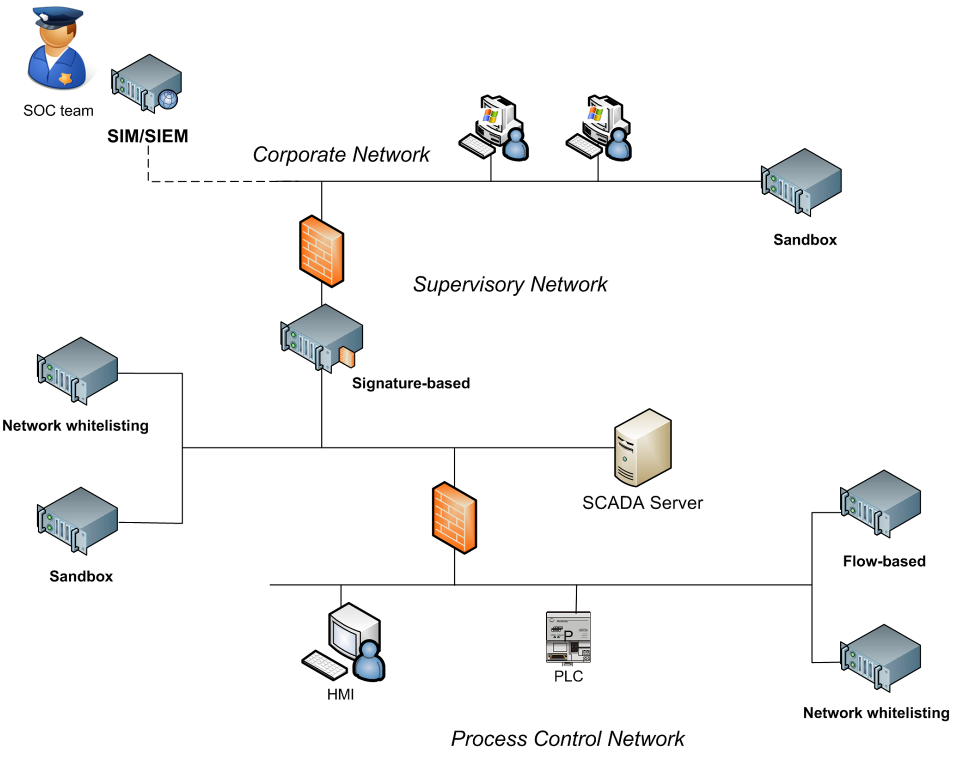

Figure 4.1: A comprehensive monitoring infrastructure to protect all segments of an industrial network

To conclude, we point out that, in order to be effective, any security solution needs to be coupled with appropriate security procedures within the organization.

4.1 Acknowledgments

This work has been partially supported by the project "Cyber Security: a Fundamental Basis for Smart Grids" (reference no. TKISG02020). We thank Christina Höfer and Fred Streefland for their useful insights.

4.2 About SecurityMatters

SecurityMatters is a spin-off from the Dutch Technical Universities. Its research team works with many different projects throughout the EU and USA and delivers game-changing network monitoring and intrusion detection technology to make their customers more secure and in control.

Some History When Damiano Bolzoni and Emmanuele Zambon left Italy and their jobs to join the research group of Sandro Etalle in the Netherlands, they had the declared intention to "create the most innovative and effective intrusion detection system of the world and build an enterprise to bring it to the market". After years of research in which several different technologies have been devised and tested, DPBI saw the light. It was a truly new, truly effective way of realizing network monitoring and intrusion detection. SecurityMatters was founded in 2008 and incorporated in December 2009. Under Emmanuele's lead, SecurityMatters development team built SilentDefense around DPBI. Version 1.0 was ready and first placed on the market in September 2011. In 2012, SecurityMatters won the SecurityMatters wins the COMMIT Wetenschapsvalorisatieprijs.

4.3 About the Authors

Sandro Etalle (1965) Sandro Etalle was 19 when he founded his first high-tech company. Meanwhile, he finished his studies at the conservatory (flute) and in mathematics (cum laude). In 1992, he left the business world to pursue an academic career. In 1995, he received a PhD from the University of Amsterdam, and is now also full professor and head of the computer security group at the Technical University of Eindhoven. He is one of the co-founders of SecurityMatters. As scientist, Etalle wrote more than 100 articles cited more than 1500 times.

Cliff H. Gregory (1948) Prior to his work in the private sector, Cliff H. Gregory served 28 years in the U.S. Navy, including assignments developing cyber security programs. His experience in Agile software product and process development, as well as business process thought leadership give him a unique vision of how to improve the safety and security of IT organizations. A published author and accomplished public speaker and trainer, Cliff holds a PhD in Math from the Dublin Institute of Technology (DIT) and a number of IT and security certifications. During his more than 35-year career in both government and private industry, Cliff has focused on Information Assurance and Systems Integrity.

Damiano Bolzoni (1981) Damiano Bolzoni graduated in 2005 in Computer Science from the Ca' Foscari University of Venice with a thesis on anomaly-based network intrusion detection system, together with his fellow classmate Emmanuele Zambon. During his master studies he worked for the Information Risk Management division of KPMG Italy. He then moved to the Netherlands to pursue a PhD at the University of Twente. He was member of the Italian Under 20 Athletics team and took part in several international competitions.

Emmanuele Zambon (1980) Emmanuele Zambon has been involved in computer security since 2003. In 2005, he also had graduated in Computer Science from the Ca' Foscari University of Venice with a thesis on Intrusion Detection Systems. During and after his studies he has been employed in the Information Risk Management group of KPMG and in Telecom Italia as a consultant, doing ethical hacking, network vulnerability/assessment, and software development/performance tuning. In 2006, Zambon followed Bolzoni to the Netherlands to work together on the development of a new technology (the core of which now forms SilentDefense). Zambon received his PhD in IT Risk Management in 2011 from the University of Twente. He is now working part-time as a post-doc at the University of Twente, doing research on intrusion detection for industrial process automation networks, and working part-time for SecurityMatters.

Daniel Trivellato (1983) Daniel Trivellato pursued his Master's degree in Computer Science at the Free University of Bozen-Bolzano, Italy, graduating cum laude in 2007. In 2012, Daniel received his PhD in computer security from the Technische Universiteit Eindhoven, where he worked on the design and implementation of innovative access control solutions for dynamic, distributed, heterogeneous systems. His work was carried out under the supervision of prof. dr. Sandro Etalle. Since 2012, Daniel works as a project leader at SecurityMatters.

Bibliography

[1] National Technical Authority for Information Assurance. Technical risk assessment. HMG IA Standard No. 1 Issue No: 3.51, UK Cabinet Office, 2009.

[2] Gleb Gritsai, Alexander Timorin, Yury Goltsev, Roman Ilin, Sergey Gordeychik, and Anton Karpin. Scada safety in numbers. Report to Congressional Requesters v1.1, Positive Technologies, 2012. http://www.ptsecurity.com/download/SCADA_analytics_english.pdf.

[3] D. Hadziosmanović, L. Simionato, D. Bolzoni, E. Zambon, and S. Etalle. N-gram against the machine: On the feasibility of the n-gram network analysis for binary protocols. In *Proceedings of the 15th International Symposium on Research in Attacks, Intrusions and Defenses (RAID 2012), Amsterdam, The Netherlands*, volume 7462 of *Lecture Notes in Computer Science*, pages 354–373, 2012.

[4] United States Government Accountability Office. Critical infrastructure protection. Report to Congressional Requesters 05-434, GAO, 2005. http://www.gao.gov/new.items/d05434.pdf.

[5] Charles P. Pfleeger. Reflections on the insider threat. In Salvatore J. Stolfo, Steven M. Bellovin, Angelos D. Keromytis, Shlomo Hershkop, Sean W. Smith, and Sara Sinclair, editors, *Insider Attack and Cyber Security - Beyond the Hacker*, volume 39 of *Advances in Information Security*, pages 5–16. Springer, 2008.

[6] Y. Song, M. E. Locasto, A. Stavrou, A. D. Keromytis, and S. J. Stolfo. On the infeasibility of modeling polymorphic shellcode. In *Proceedings of the 14th ACM conference on Computer and communications security*, CCS '07, pages 541–551. ACM, 2007.

[7] K. Stouffer, J. Falco, and K. Scarfone. Guide to industrial control systems (ics) security. Special Publication 800-82, National Institute of Standards and Technology, 2011. http://csrc.nist.gov/publications/nistpubs/800-82/SP800-82-final.pdf.

[8] Joseph Weiss. *Protecting Industrial Control Systems from Electronic Threats*. Momentum Press, 2010.