Targeted Attacks against Industrial Control Systems: Is the Power Industry Prepared?

Maria B. Line[†][*], Ali Zand[‡], Gianluca Stringhini[§], and Richard Kemmerer[‡]

[†]Norwegian University of Science and Technology
[*]SINTEF [‡]UC Santa Barbara [§]University College London

maria.b.line@item.ntnu.no, *{zand,kemm}@cs.ucsb.edu*, *g.stringhini@ucl.ac.uk*

ABSTRACT

Targeted cyber attacks are on the rise, and the power industry is an attractive target. Espionage and causing physical damage are likely goals of these targeted attacks. In the case of the power industry, the worst possible consequences are severe: large areas, including critical societal infrastructures, can suffer from power outages. In this paper, we try to measure the preparedness of the power industry against targeted attacks. To this end, we have studied well-known targeted attacks and created a taxonomy for them. Furthermore, we conduct a study, in which we interview six power distribution system operators (DSOs), to assess the level of cyber situation awareness among DSOs and to evaluate the efficiency and effectiveness of their currently deployed systems and practices for detecting and responding to targeted attacks. Our findings indicate that the power industry is very well prepared for traditional threats, such as physical attacks. However, cyber attacks, and especially sophisticated targeted attacks, where social engineering is one of the strategies used, have not been addressed appropriately so far. Finally, by understanding previous attacks and learning from them, we try to provide the industry with guidelines for improving their situation awareness and defense (both detection and response) capabilities.

Keywords

Cyber situation awareness; Incident management; Industrial control systems; Information security; Interview study; Power industry; Preparedness; Targeted attacks

1. INTRODUCTION

Information security incidents happen at an alarming rate and the cost of these incidents is considerable for industries and governments [7]. Incidents can be the result of human error, technology glitches, malicious actors, or any combinations of these. Malicious actors have various skills and motivations. On one side of the spectrum, script-kiddies try out their public exploitation tools on a large set of targets in hopes of finding unpatched systems that are vulnerable. On the other side of the spectrum, highly skilled attackers with predetermined targets and a considerable amount of budget and time, try their custom private zero-day tools that were developed for a specific target – also known as targeted attacks.

Targeted attacks are on the rise [9], but compared to their less-sophisticated counterpart, they are much less frequent. Unfortunately, when they occur they can cause catastrophic damage to the affected systems and/or organizations. Because low budget script-kiddie attacks use and reuse a limited number of well-known techniques, they can be prevented, stopped, and/or mitigated by off-the-shelf security products. Given this situation, critical industries are relatively well protected against these types of attacks. In contrast, targeted attacks are more challenging to deal with because of their uniqueness in each case. They may use an array of advanced attack techniques, ranging from social engineering to finding and exploiting existing vulnerabilities in the implementation and configuration of a particular organization's systems. When organizations perform risk assessments, targeted attacks will typically, and in most cases rightfully, end up in the outer corner of the risk matrix; categorized as *lowest probability, highest damage potential*. This leaves the organization with a challenge: the organizations realize that targeted attacks actually might happen, but, since the probability is still quite low, they prioritize preventing the higher-probability incidents. Also, if an attacker, or a well organized group of attackers, with the right amount of resources really wants to strike a specific organization, there are limited, if any, mechanisms that could prevent it.

The challenge with defending against targeted attacks is twofold:

- A targeted attack is typically quite difficult to detect. Intrusion detection systems (IDS) have difficulties identifying attacks with a wide time frame. Also, social engineering is a very powerful tool for mapping an organization and getting access to information. These attacks are completely out of scope for IDSes, which rely on well-known patterns.

- Even when the current state-of-the-art is able to deal with some aspects of the attack, organizations have not yet caught up with the latest techniques.

Targeted attacks against the energy industry have already happened [3]. This demonstrates that there are well organized groups of attackers out there, with the motivation, resources, and competence to cause serious damage. Even though such attacks are not yet common, the industry needs to be well prepared to meet this threat. Learning from past incidents is one efficient way of preparing, but the next attack is likely to consist of new and unexpected elements as well.

The energy industry is currently facing major changes to its industrial control systems. The implementation of new technologies and the use of commercial off-the-shelf hardware and software increases the chances of introducing new vulnerabilities and makes these systems attractive targets for attackers [18]. As many as 59% of all incidents reported to the Department of Homeland Security in 2013 occurred in the energy sector [3], which demonstrates that attacks against this industry is not just a futuristic scenario.

The human factor is of utmost importance where the technology fails. However, in the world of ever-changing computer systems, new applications, and new security threats, educating all users to be well functioning perimeter controls for an organization is, at the very least, challenging. Still, the human system operators need to be able to interpret alerts, put pieces of information together, and know about possible attacks and their consequences. This ability is referred to as Cyber Situation Awareness (CSA) and can, to some degree, be supported by automatic tools.

The following are our contributions in this paper:

- We conduct, to the best of our knowledge, the first systematic assessment of cyber situation awareness in power distribution system operators.
- We provide a new taxonomy of targeted attacks and use it to provide insight into the importance of different aspects of cyber situation awareness for defending against these types of attacks.
- We provided a list of prioritized suggestions for the power distribution system operators in order to increase their cyber situation awareness.

For practitioners, the taxonomy for targeted attacks should be valuable for educational purposes and for prioritizing preventive measures. For the research community, our documentation of gaps between the practices used in industry and current research efforts outline future research needs from an industrial point of view.

The paper is organized as follows. An overview of related work is provided in Section 2. Section 3 introduces a taxonomy for targeted attacks and classifies publicly well-known attacks based on this taxonomy. Cyber situation awareness is described in Section 4, where a categorization is given and a set of questions for assessing CSA is presented. The research method used in our assessment is described in Section 5 and findings from the interviews are summarized in Section 6. Section 7 discusses the organizations studied with respect to level of preparedness for targeted attacks, while Section 8 provides concluding remarks and suggests further work.

2. RELATED WORK

With the increasing trend of state-sponsored cyber espionage and possible cyber terrorism, a need for addressing more advanced targeted attacks and advanced persistent threats (APT) has emerged. One key component of the ability of an organization to deal with targeted attacks is the organization's knowledge about the environment, the threats, and the events. This knowledge is called cyber situation awareness. Cyber situation awareness helps an organization predict and anticipate future attacks and detect, prevent, mitigate, or respond to them in an efficient and timely manner.

To measure the preparedness of an organization for dealing with targeted attacks one needs to know the organization's level of cyber situation awareness, defenses, and incident response capabilities. Previous work has addressed the defense against targeted attacks in different ways. Some researchers have analyzed famous targeted attacks, while others have focused on measuring and improving cyber situation awareness or incident response.

Targeted attacks. Thonnard et al. [25] presented the first large-scale study of targeted attacks that affect corporate networks. In particular, in their work they analyzed a number of targeted attacks where e-mail was used as the means of getting a foot into the victim's network. Multiple researchers studied in detail single targeted attacks, such as Stuxnet [14, 17]. Tankard [24] presented a taxonomy of the different steps of targeted attacks, identifying common traits such as spearphishing and lateral movement. In this paper, we provide a more comprehensive taxonomy of targeted attacks, identifying common elements of notable attacks that have happened in the recent years.

Incident management. Current practices in incident management in industrial control organizations were investigated by Line et al. [19]. They identified the lack of a common perception among control room operators of what an information security incident is. They pointed out that a big challenge in the control room is recognizing that an incident is actually occurring and determining the most appropriate first responses. Also, they found that general emergency preparedness exercises are regularly performed, but IT-based scenarios are rarely used as a basis for such exercises. Furthermore, exercises that include both IT and control room personnel are not common. Control room operators know their systems in extensive technical detail, but they are not well trained in information security issues; understanding threats and possible symptoms of an attack is therefore a competence that needs to be strengthened. However, this study did not investigate what technical security mechanisms were used to monitor and detect attacks in the control systems, and the researchers did not relate their findings to the concept of cyber situation awareness nor address the specific challenge of targeted attacks. We are not aware of other studies discussing information security practices related to industrial control systems.

Cyber situation awareness. According to Barford et al. [8] situation awareness can in general be described as a three-phase process: situation recognition, situation comprehension, and situation projection. CSA systems are intended to support a human operator in understanding ongoing incidents. Tadda [23] provides an overview of metrics developed for measuring the performance of CSA systems. He specifically points out the need for research in measuring the level of situation awareness achieved by human operators, and he indicated that it would require quite different means than measuring the performance of a computer system.

Goodall *et al.* [15] use data mining techniques to extract relations between cyber mission, cyber assets, and users. Grimaila *et al.* [16] suggest a cyber-damage assessment framework. Doupé *et al.* [13] conducted a live hacking competition that introduced the concept of cyber missions and cyber assets into the exercise. Teams acting based on better cyber situational awareness were rewarded. Dai et al. propose a reference model to address the unique cyber situation awareness needs of real-world missions [11]. Paul et al. provide a taxonomy of questions necessary for a cyber security analyst to acquire cyber situation awareness [20].

In this paper, we attempt to assess how prepared the power industry is for targeted attacks by studying their level of cyber defense, cyber situation awareness, and incident response.

3. TARGETED ATTACKS

In recent years we have witnessed a growth in the number of attacks that target specific entities, such as government agencies and large organizations. Examples of these attacks include the Night-Dragon attack against energy companies [2] and Stuxnet, which targeted the Iranian nuclear program [17]. The way in which targeted attacks are carried out is very different from large-scale threats that affect a large number of victims. First of all, since these attacks target a specific organization, miscreants can tailor their attacks to that organization's infrastructure and software configuration. As a result of this, targeted attacks often use vulnerabilities and exploits that have never been observed before (i.e., *zero-day* attacks) [25]. Second, targeted attacks can be performed over a long period of time, and usually involve multiple steps.

Although single instances of targeted attacks have been studied in detail by researchers, we are not aware of any categorization of the general characteristics that are common to multiple instances of targeted attacks. In the following, we provide a first taxonomy of the steps that compose a targeted attack. We hope that, although this is a first step, this taxonomy will constitute a foundation for future work in the area and will help researchers in defining the common traits of different targeted attacks as they happen in the wild.

3.1 A Taxonomy for Targeted Attacks

By reviewing the description of notable targeted attacks that have been analyzed by researchers and by the press, we identified common characteristics that describe such attacks. In particular, we identified four typical characteristics of a targeted attack. The first characteristic is the purpose of the attack. Then, we identify two characteristics that describe how the attack is conducted: the initial attack vector and the lateral movement strategy. Finally, we characterize targeted attacks based on the location of the command and control server used to coordinate the infected machines that are leveraged in the attack. In the following, we describe these characteristics in detail.

Purpose of the attack. Traditional large-scale attacks have a variety of purposes, which span from leveraging infected machines for an economic gain (e.g., sending spam) to stealing personal information from the victims. In the case of targeted attacks, we identify two main purposes that drive attackers. The first purpose that motivates attackers is the *exfiltration* of sensitive information from the target organization. Notable targeted attacks focused on either stealing corporate secrets from the victim organization [2] or exfiltrating source code for the organization's products [1]. The second purpose is *sabotage*; attackers might try to sabotage the operations of a competitor or of an enemy government [17]. A third possible purpose for a targeted attack could be *extortion*: attackers could make their victim network unavailable and ask for a ransom. However, since this type of attack would require a sabotage action first, we consider it as a subset of the sabotage category.

As we will explain later, the purpose of a targeted attack can affect the way in which the attack can be detected: exfiltration requires information to be sent outside the organization, and it can potentially be detected by defenses on the edge of the network. On the other hand, sabotage actions could be detected by anomaly detection systems deployed inside the network.

Initial attack vector. As we mentioned, targeted attacks are typically composed of several steps. In the first step, attackers need to get a foot into the victim's network. This can happen with or without requiring an action by the victim organization's employees. We define an initial vector to be *automatic* if the attacker can compromise a machine (e.g., a desktop or a server) in the victim network without any action needed from employees in the organization. As an example, we consider a *drive-by download* attack [21] that exploits a vulnerability in one of the employee's browsers and automatically downloads malware as an automatic attack. Similarly, an attack that exploits a vulnerability in one of the organization's webservers is considered an automatic vulnerability. A *manual* initial attack vector, on the other hand, requires an interaction from somebody inside the company to succeed. Examples include *spearphishing* emails that ask an employee to download and install an executable or malware that is disseminated through infected USB drives that need to be plugged into an internal computer to get installed. If an attack requires both automatic and manual interaction, for example it requires the victim to fall for a spearphishing scam and click on a link to enable the target web page to exploit the user's computer, we still consider this attack as "manual."

Similarly to what happens for the purpose of attacks, the different *modus operandi* used by the attackers influence the way in which an organization can detect the attack. Automatic attack vectors can be detected by intrusion detection systems that monitor the network behavior within the organization. Conversely, attacks that are triggered by a user action, such as manually installing a malicious program, have to rely on host-based measures, such as signature-based antivirus systems.

Lateral movement. After establishing a presence in the victim's network, attackers usually aim at compromising more computers. The reason for doing this is to gain access to computers that have higher privileges than the one that the attackers were able to compromise during the initial attack [2]. These operations are known as *lateral movement* [1]. Think, for example, of an attack that first compromises the office manager's computer and then tries to obtain access to the CEO's computer, with the goal of stealing corporate secrets.

Similarly to the initial attack, lateral movement can happen in an *automatic* or *manual* fashion. However, because these attacks are not coming from the outside, companies need to deploy defenses that monitor connections within the organization to detect them, rather than only relying on systems that monitor connections to and from the outside.

Location of the command and control server. Attackers need a command and control server (C&C) to give orders to their infected machines. We define two possible locations for the C&C used in a targeted attack: *inside* the victim's network or *outside* the network. In the first case, attackers use one of the computers that they compromised as the command and control server, while in the second case they instruct their infected computers to connect to one or more remote servers to retrieve their orders.

3.2 Real-life Targeted Attacks

A number of attacks targeting critical infrastructure industries are well-known through the media and information security analysts. This section summarizes such attacks with respect to their assumed purpose, strategies used, and consequences, for the cases where they are publically known. Table 1 reports a brief description of notable targeted attacks given the taxonomy we provided in Section 3.1.

NightDragon. This attack was geared towards the harvesting of sensitive information related to competitive proprietary operations and financial details regarding field bids and operations, and it targeted mostly energy companies [2]. The attack would initially exploit one of the company's servers by compromising one of the company's web servers, or by infecting one of the company's employees' computers through a spearphishing email. The exploited server would then be used as a C&C server for the operation, and additional machines within the company would be compromised when visiting it. Finally, the infected machines would be instructed

[1] Despite what the name might suggest, lateral movement is not limited to the attacker gaining access to hosts similar to the ones originally compromised (or hosts in the same subnetwork), but it also refers to hosts that have higher privileges.

Attack Name	**C&C Location**		**Initial Attack**		**Lateral Movement**		**Final Step**
	Internal	External	Automatic	Manual	Automatic	Manual	
NightDragon	✓		✓	✓	✓		Exfiltration
Operation Aurora		✓	✓		✓		Exfiltration
Careto		✓		✓			Exfiltration
Stuxnet		✓		✓	✓	✓	Sabotage

Table 1: Taxonomy of targeted attacks. We characterize the targeted attacks based on four elements: (i) the location of the command and control server (internal or external to the organization), (ii) the way attackers first set foot into the network (automatic — by exploiting a vulnerability, or manual — by leveraging social engineering techniques), (iii) the way in which attackers perform lateral movement within the company's network (manual or automatic), and (iv) the purpose of the attack (data exfiltration or sabotage).

to disable the company proxy and send sensitive data to a foreign server controlled by the cybercriminal.

Operation Aurora. This attack targeted Google, with the goal of stealing the company's corporate secrets [1]. The attack started by infecting computers by using a browser vulnerability, and it instructed the infected computers to connect to an external C&C server. The infected machines would then find more vulnerable computers and automatically exploit them. Finally, sensitive information from the company would be sent to a foreign server.

Careto. This attack was created with cyber espionage in mind, and its goal was to exfiltrate sensitive information from the target company [4]. First, attackers would send spearphishing emails to the company's employees, pointing them to malicioud web pages. Infected machines would then be instructed to connect to an external C&C server and send sensitive documents to the cybercriminals.

Stuxnet. The purpose of Stuxnet was to reprogram industrial control systems of a specific type and hide any changes [5, 6, 14]. It targeted centrifuges at the Fuel Enrichment Plant in Natanz, Iran. The authors of Stuxnet are unknown. The attack started by having a consultant working for the target company insert an infected removable device into one of the company's computers. The malware would then stealthily propagate to other computers on the network by either infecting other removable drives, by automatically exploiting vulnerabilities, or by copying itself in the company's network shares. Infected computers were then instructed to connect to an external C&C. Once the infection propagated to a computer that had access to the control systems, it altered the operation of these systems in a sabotage attack.

4. CYBER SITUATION AWARENESS

Cyber situation awareness [8] has attracted considerable interest from the security research community in the recent years as a way to help proactive cyber defense, especially in the face of targeted attacks. In short, CSA involves understanding the network environment, the missions, the resources, and the threats. In other words, one should know 1) her own missions and resources and how they depend on each other, 2) about the network events, and 3) the threats, how they threaten the missions, and how they can be mitigated or prevented.

Cyber situation awareness is an umbrella term for a set of capabilities and techniques that is difficult to measure. This section presents the theoretical background we used for designing our interview guide.

CSA consists of the following capabilities [8]:

- Comprehension of the current situation
- Understanding the impact of attacks
- Understanding how situations evolve
- Understanding attacker's behavior
- Understanding the causes of the current situation
- Being aware of the quality of the collected information (confidence or trustworthiness)
- Predicting plausible futures

We used this list of capabilities as a guide to design our interview guide. We specifically tried to determine the existence or lack of each capability.

Several traditional security tools and methods have been used to improve an organization‘s cyber situation awareness. The following are the commonly-used and well established traditional cybersecurity methods [8]:

1. Intrusion detection and alert correlation
2. Damage assessment (e.g., using dependency graphs)
3. Intrusion response
4. Taint and information flow analysis
5. Attack trend analysis (e.g., finding the source of intrusions)
6. Vulnerability analysis (e.g., using attack graphs)
7. Causality analysis and forensics

In addition, we wanted to know the relationship between the organization's CSA capabilities and its ability to counter targeted attacks. CSA techniques may have the following properties in regard to dealing with targeted attacks:

- How does the technique help? Detect, prevent, react (limiting the damage, fixing it, or preventing future attacks), predict, forensics
- Where is the technique used? Before the initial infection, after the initial infection and before propagation, after propagation, after the final step
- What is a deviation? An attack behavior or out-of-ordinary behavior? (misuse/anomaly detection)
- How is the attacker modelled? Based on the attacker behavior or based on the system?
- What kind of input is required? Network traces, OS/service logs, high-level information about the organization. Is the technique white box or black box?
- How high level is the output? Is it actionable? Is it automatable? (Automatic/semi-automatic/manual)
- How integrated is the technique with the work flow of the organization (monitoring systems, training, roles)?
- How connected is the technique with the organization missions? Is it assigned to different resources proportional to the criticality of that resource?
- How does the technique improve its performance and precision? Learning (supervised or unsupervised)? Manual update (whitebox/blackbox)?

We aimed at answering the above questions with our interview guide. A mapping between the CSA capabilities and our questions is provided in Appendix B.

5. RESEARCH METHOD

We conducted a survey using qualitative interviews [22]. The objective was to assess the cyber situation awareness for targeted attacks against control systems. We wanted to investigate to what degree distribution system operators (DSOs) in the electric power industry are prepared for such attacks and to measure their abilities to detect these attacks and appropriately respond. Furthermore, we wanted to identify knowledge gaps that should be given attention in future preparedness exercises. Six large Norwegian DSOs participated in the study[2].

5.1 Data Collection and Analysis

We used semi-structured interviews, which were guided by a predefined set of questions, but they allowed for additional, unplanned questions or a change in the order of questions [22]. The interview guide was developed based on a categorization of elements comprising cyber situation awareness, as presented in Section 4. One fellow researcher and one expert from a supplier of control systems assisted in evaluating the questions. The interview guide is presented in Appendix A.

The interviews were carried out during two weeks in April 2014. Due to large geographical distances between the interviewer and the respondents, the interviews were performed as online meetings using Microsoft Lync[3]. All interviews were voice recorded and transcribed, the study was registered at the Data Protection Official for Research[4], and all respondents signed a consent agreement. All DSOs but one (F), requested that the researcher performing the interviews sign a non-disclosure agreement.

Confidentiality issues prevented the interviewing researcher from sharing detailed transcriptions with fellow researchers. Therefore, a short summary of each interview was written, in which the participating organizations were anonymized. The other researchers were then able to participate in discussing the main findings and indirectly contribute to the data analysis. These short summaries excluded the need for detailed coding and analysis of the data material, as they provided sufficient overview and the insight needed for writing up the results. The detailed transcriptions were still used by the main researcher.

5.2 Industrial Case Context

The six DSOs were selected because they are among the largest DSOs in the country. An overview of the participating DSOs, hereby denoted A-F for anonymization purposes, and respondents is shown in Table 2.

We originally asked for the control room managers, but as the interview guide was distributed along with the request for participation, some of the DSOs identified other persons who were better able to answer our questions. Giving the interviewees the possibility of doing some preparation was a means to improve the quality of the responses. We did not expect all respondents to be fully oriented on all technical security mechanisms, as their roles and responsibilities may vary depending on the organization, and the knowledge we asked for may be distributed across a number of personnel in each organization.

DSOs A, C, F outsourced their IT operations, including the network infrastructure for the control systems, to an external IT supplier that is 100% owned by the corporation. The DSOs have dedicated personnel in the control room responsible for maintaining the control systems. DSO A also has two persons responsible for information security in the control systems. The respondent from A was part system manager and part information security manager for the control systems. The respondent from C was an information security advisor for the corporation and employed by the IT supplier. The respondents from F were the control room manager, the control system owner (from management), and one IT-technical manager from the IT supplier.

DSOs B, D, E operate their own IT services and network infrastructure. DSO B has an internal department for IT operations, including the network infrastructure for the control room. The control systems are maintained by dedicated personnel in the control room. The respondent was the information security manager of the corporation, which includes the distribution grid and power production. DSO D has a dedicated group that maintains the control systems and the network infrastructure for both the administrative and the control systems. Administrative software and services are outsourced to an external supplier. The respondents were three persons from this group: the manager and two engineers. In DSO E two dedicated groups are responsible for maintaining the control systems and the network infrastructure respectively, and the respondent was the leader of a coordinating committee for these two groups where IT and IT security matters were discussed. Hence, only DSO D has the same personnel maintaining both the control systems and the network infrastructure. In all other DSOs these two responsibilities are shared between two different groups.

All DSOs have dedicated personnel for maintaining their control systems. In addition, they all have service agreements with their supplier for the control systems, which includes assistance in case of failures, annual reviews of the systems, and critical patches whenever necessary.

6. FINDINGS

In this section we present findings from the interviews with DSOs on cyber situation awareness and the use of supporting technical tools. We would like to remind the reader that the scope of this study was limited to the industrial control systems and did not include the whole organization. All findings are therefore related to the procedures and practices for the control systems only.

6.1 Perception of Threats

Observed attacks. None of the respondents have ever detected any attacks in their control systems, neither targeted attacks nor general malware infections. DSO F has, however, detected malware on the inside of the network perimeters that protect both the administrative systems and the control systems, but only administrative computers were infected. The respondent from A said that they consider targeted attacks as having a low probability of occurring, as there has been no such attacks against the power industry in the country so far. As soon as the first attack of this type is observed, they would consider the probability of additional attacks to be much higher than today. Still, he stated that they need to prepare in advance, as they might not have time to implement new measures when the threat escalates.

Worst case. The worst case scenario envisioned by all respondents is hackers gaining access to the control systems and controlling the power switches. This would allow the attackers to cause power outages in large areas within minutes. There are several layers of security mechanisms that must be bypassed in order to gain such access, and it requires extensive technical knowledge of protocols and the software used in the control systems as well. Hence, the respondents considered the probability of such an attack to be quite low. Still, the potential damages are catastrophic. The respondent

[2] They all serve more than 80,000 power consumers and are considered large in Norway.

[3] http://products.office.com/en-us/lync/

[4] Equivalent to the Institutional Review Board (IRB). URL: http://www.nsd.uib.no/personvern/en/index.html

DSO	Role of interviewee	Outsourced IT services	Infosec coordinator
A	Responsible for infosec and daily operations of control systems	Yes; network infrastructure to external supplier	no
B	Infosec manager of corporation	No; but network infrastructure maintained by separate dept.	yes
C	Infosec advisor in IT supplier (100% owned by corporation)	Yes; network infrastructure to external supplier	yes
D	Three interviewees: Responsible for daily operations of control systems (manager and two engineers)	No; and network infrastructure and control systems maintained by same group	no
E	Responsible for infosec and daily operations of control systems. Leads a coordinating committee for IT matters for network infrastructure and control systems.	No; and network infrastructure and control systems maintained by different groups	yes
F	Three interviewees: Control room manager, Owner of control system (Corp.mgmt.staff), and Responsible for network infrastructure for all systems	Yes; network infrastructure to external supplier	no

Table 2: The distribution system operators (DSOs) and respondents participating in the interview study

from DSO A specifically mentioned advanced persistent threats; someone spending time to get to know their systems and understand patterns, in order to disable physical defences in the grid and precisely control the switches.

Potential attackers. The respondents acknowledged the increased risk of being hit by general attacks due to increased connectivity and an increased use of commercial off-the-shelf products. Furthermore, terrorists and foreign nations and intelligence were considered potential attackers by DSOs A, C, D, E. However, they considered themselves as not being large enough to be attractive targets for attackers. At the same time, the three largest DSOs did admit that they have customers that might be attractive targets, and, therefore, they themselves might be exploited in an attempt to strike these customers. Also, DSO C was worried about their suppliers being attacked, as that could put the DSO itself in danger too. The supplier holds extensive knowledge about their systems and also has the possibility to connect remotely.

6.2 Preparatory Activities

Procedures. None of the DSOs have any written procedures for responding to incidents in the control systems, except for DSO B, who has procedures on how to respond to malware infections. Both DSO B and F mentioned that physically disconnecting the control room from the network would be a possible and efficient reaction. The respondents from D admitted to the lack of procedures and stated their need to work on this. At the same time, they rely on their employees to have the experience and knowledge to enable them to respond properly if needed. Also DSO A has identified their own lack of procedures in a comprehensive risk assessment recently, and they are currently working on fixing this and other identified shortcomings as well. General incident response procedures exist for the whole organization of DSOs E and F, but not specifically for the control systems[5].

Risk, criticality, and dependency assessments. All the DSOs have performed dependency and criticality assessments to determine the appropriate levels of redundancy and other security mechanisms. Risk assessments are performed at least annually, as is required by national regulations. However, all DSOs reported that any system changes trigger additional risk assessments of some kind. Risk assessments can be performed informally in a small group or more extensively by external consultants, depending on the scope and need. None of the DSOs mentioned that they have had external technical security assessments, such as penetration testing, of their control systems, but DSO D intends to do this in the near future.

Awareness. Specific security awareness training for personnel operating the control systems has not been performed by any of the DSOs. This lack of training was identified by DSO A as one of their major missing pieces, and they intend to put more effort into this now. General awareness raising activities were however reported by DSOs D, E, F. The respondent from E stated that this is quite a challenging area, but they try to do their best to build a security culture, as they acknowledge that technology can never compensate for lack of culture. DSOs B, C, and D believed that the level of awareness has increased recently, as the field of information security has matured and the mindset among management has changed. There is in general a good understanding of the need to protect the control systems from malware infections and other threats. The need to have different computers for different purposes (Internet access vs. control system activities) is also evident.

Training. Two of the DSOs (B, D) reported that they have performed one table-top exercise on responding to an IT security incident. DSO A has performed one exercise on unavailable control systems due to a fire, where the standby control room had to be put into operation. The other three DSOs (C, E, F) have never performed exercises based on an IT security incident. Furthermore, none of the DSOs have based any exercises on what they state as the worst case scenarios. The respondent from A stated that as long as they lack a decent continuity plan, they do not have the necessary basis for performing exercises. They are currently working on setting up a complete test system, which they could use for such exercises, as it is impractical to perform exercises on critial systems where operational disturbances are not acceptable. This was also given by DSOs B and D as one of the reasons for not performing realistic preparedness exercises on IT security incidents. DSO F intends to include IT-based scenarios in their new five-year plan for preparedness exercises, and the respondent from DSO C hoped for more focus on IT-based exercises as their newest hiring is someone to be responsible for overall security, including preparedness planning and exercises. All respondents were aware that the au-

[5] This might be the case for the other DSOs as well, but was not investigated in this study due to the limited scope of control systems.

thorities require the DSOs to include IT security incidents in their preparedness exercises since summer 2013, when the new version of the regulations came into force.

6.3 Technical Security Mechanisms

All respondents described similar infrastructures, policies, and security mechanisms: complete documentation is in place, several layers of firewalls, and a DMZ zone between the administrative systems and the control systems. Furthermore, connecting external PCs and other devices to the control systems is prohibited, and this rule is claimed to be enforced at all times. All DSOs have a service agreement with their supplier of control systems, which includes an annual review of the systems with necessary upgrades, patching of critical vulnerabilities whenever needed, and assistance 24/7 in case of failures.

Patching regime. All DSOs usually install patches themselves. On some occasions there is a need for the control system supplier to install critical patches. In these cases, the supplier uses a dedicated computer in a physically controlled room on their own premises and connects remotely via a dedicated VPN channel. This VPN channel is opened on request only, and it is closed at all other times. The supplier recommends a certain patch to be installed, this action must be approved by authorized personnel at the DSO, and then the VPN channel is opened for a specific time period. However, the dedicated computer is used to serve more than just one organization, although most likely it is not used for other purposes[6]. One of the DSOs (E) expressed a slight concern regarding this practice and stated that whether they should pose a requirement on their supplier to have a dedicated computer for this DSO only is under continuous consideration.

Encryption. The use of encrypted communication from control room to components in the field is rather limited. In the cases where the DSO owns the infrastructure, they do not see the need for encrypting the communication. Also, the majority of the components in the field were not designed to support encryption. The respondent from DSO E anticipated this will change in the future, when the use of, and need for, encryption will increase and that new equipment will support encryption. The respondent from DSO A stated that communication with the most critical components in the field, on lines owned by a third party, is encrypted. He added that the security of the communication should be in accordance with the level of physical protection in the field.

Monitoring and detection. Firewalls are deployed by all DSOs separating them from external networks; in some cases also on a selection of hosts inside the control systems and between layers in the control systems. DSO B and D reported that their firewalls also include IPS functionality, while DSO C has plans for implementing IPS in the near future. DSO A plans to increase the number of host-based firewalls and implement monitoring systems as well to strengthen their ability to monitor and detect attacks. DSOs B, C, and D implement standard logging functionalities on servers and routers in the control systems.

The firewalls and other logging systems generate alerts in case of suspicious events. However, the practice of how such alerts are followed up varies greatly between the DSOs. None of the respondents reported having systems that correlate logs from different sources. DSO D explicitly expressed that this is something they would like to have. One of the respondents from DSO F stated that they systematically investigate all alerts, but he was from the IT supplier and refers to detections in the network infrastructure, not the control systems. However, they operated the network perimeter that surrounds the control systems. DSO B and E explicitly stated that something must trigger their attention for the logs to be investigated. Otherwise, the alerts may just as well be ignored. The respondent from DSO C expressed his concern that alerts are not being correctly interpreted when they occur.

Antiviruses were deployed on all servers in the control system at DSO C and D, while DSO A, B, and F do not run any antivirus software in their systems. In DSO E they run all patches and other data to be inserted into the control systems through an antivirus control. None of the respondents mentioned having systems that make it possible to isolate a compromised server and/or computer.

A national CERT for the power industry is about to be established. It is a collaborative effort between the largest DSO, the operator of the regional grid, and the largest power producer in the country. The intention of this CERT is to monitor network traffic in order to detect coordinated attacks against this industry and to perform advisory services to its members.

Assessment of sufficiency The respondents were asked whether they felt that the most critical parts of the control system were sufficiently protected. All but the one from DSO A expressed that they are satisfied with the current security level. In DSO A they have recently performed a thorough risk assessment and are currently working on implementing the identified measures, which is why he is not satisfied at this point. His main concern is that they have no methods or tools for monitoring and detecting attacks except from limited use of logging in the firewalls and servers. The respondent from DSO E stressed the need for a certain balance between security and functionality. He acknowledged that there is always room for discussing possible improvements of the security mechanisms, but he is satisfied with their current implementations based on the need for this balance.

7. DISCUSSION

In this section we discuss the level of preparedness for targeted attacks of the organizations studied. Furthermore, we provide a comparison between known targeted attacks and the defenses deployed by the DSOs. Finally, a list of prioritized recommendations to DSOs is provided.

7.1 Level of Preparedness

Our findings show that the DSOs have similar cyber situation awareness capabilities. They significantly depend on their perimeter security systems.

The organization structure, whether parts of the IT/network/control systems are outsourced or not, does not seem to make a difference in how they perceive threats or select security mechanisms.

Missions, resources, and dependencies. All DSOs were well aware of their missions and resources and the interdependencies between the organization resources and missions. One reason can be the relatively small size of the IT systems. Also, regulations force the power companies to perform risk assessments and dependency analyses.

Threats. Respondents were able to speculate the impact of a possible targeted attack. One of the respondents from DSO D stated that it used to be a problem that control system operators stored movies, music, and other personal data on the control computers, but this is not an issue any more. The level of security awareness has been raised dramatically in recent years, among both operators and management. There is a comprehensive understanding of the need for securing and protecting the control systems against general information security threats.

[6] This is an assumption made by one of the respondents. We have not talked with representatives from any of these control system suppliers.

Network events. The DSOs do not have a comprehensive and timely view of the important events happening in their network. None have considerable security monitoring inside their network. They employ minimal intrusion detection, and none have an integrated alert correlation system. Although the static nature of the tasks in a DSO makes it the perfect candidate for an anomaly detection system, none of the DSOs have implemented or acquired such a system. Lacking the capability of tracking the network situation, in the case of an attack, these companies are also unable to track the attacker's behavior. This significantly limits their capabilities for identifying the attacker and performing forensics.

Policies. All six DSOs suffer from a lack of intrusion response policies and practices. We specifically find it surprising that the DSOs under study have a lower level of intrusion response preparedness than an average company [26]. We speculate that one reason for this is the fact that the cost of incident response practices for a power company are far more than that of an average company. Another reason might be that they consider the probability of attacks to their industrial control systems as rather unlikely. The respondent from DSO E pointed out that physical attacks would be a lot easier to perform than a sophisticated targeted attacks, which would require extensive technical expertise. This is also supported by other respondents as well; as long as the industry has not experienced any cyber attacks, they consider them to be very unlikely. They are more familiar with the risk of physical attacks, or at least physical failures, and this is what they use as scenarios in their preparedness exercises and what they base their ultimate countermeasures on; shutting down the control room. However, a physical attack requires physical presence. A cyber attack could be carried out from all over the world. Also, there is the issue of accountability. Performing a physical attack requires more courage and spirit than a cyber attack, where one can safely sit behind a screen and perform damage to objects that are far away. Hence, comparing the probabilities for these two very different types of attacks is indeed impossible.

All six companies perform vulnerability analyses regularly. A factor contributing to this is that regulations require them to do this at least annually.

Intrusion response. If an attacker is detected inside the control systems, the universal response stated by all respondents is that they can just shutdown the systems but still maintain the power supply, at least for a while. However, this calls for the need to detect and understand that malicious activities are occurring. Targeted attacks are typically not possible to detect just by having humans looking at the systems.

Most DSOs never ran an exercise involving a computer incident as a threat: only DSO B reports that they ran it, but it was an isolated event that does not occur regularly. We believe that simulating such an attack is critical to improve the awareness of computer threats by the DSOs, and we advocate for them to implement regular exercises, similar to what they do with physical outages and incidents.

7.2 Detecting Known Targeted Attacks

We compared the technical and educational security countermeasures described by the various respondents during the interviews, to assess the level of preparedness of the various DSOs to potential targeted attacks. In particular, we compared the targeted attacks described in Section 3 with the defenses employed by the DSOs, to see how and at which stage a targeted attack against them could potentially be detected and blocked. The goal was to identify possible elements of a targeted attack that could be used to improve the defenses of the DSOs.

We summarize the defenses employed by the different DSOs in Table 3. For each step of a targeted attack, and each mode of propagation (manual or automated), we look at the technical or policy countermeasures put in place by the DSO that could have potentially detected and blocked the targeted attack at that stage. The possible defenses include using a Firewall, Intrusion Prevention System (IPS), a host-based anti-virus product, a company policy, physical separation, or awareness generated by a security exercise.

While all DSOs use different DMZs for the administrative and the control network, only DSO D has the control room completely separated and disconnected from the rest of the network. We believe that this *modus operandi* is the best defense for targeted attacks against the control system, and we encourage other DSOs to use physical separation. We understand, however, that having complete separation is challenging, because, for example, it prevents the supplier from connecting to the control machines to install updates.

Because of the logical separation between the different parts of the network, most DSOs rely on detecting attacks by using a firewall. Three of them use an off-the-shelf IPS to detect potential threats, but IPSs are known for generating many alerts, and DSO B admits that often they do not have the man power to go through all the generated alerts. In general, a firewall and an IPS work in detecting attacks that have been observed before, but struggle in fighting attacks that have been engineered specifically to target a company, such as the targeted attacks that we described in Section 3.

Regarding manual attacks, such as inserting a removable storage device in one of the control room computers, most DSOs have strict policies about this. We believe that developing and enforcing strong security policies is a good countermeasure against the spread of targeted attacks. Conversely, many DSOs rely on anti-virus programs to check whether the files downloaded on their computers are malicious. Similar to what we said with IPSs, these off-the-shelf countermeasures work in protecting the company against traditional threats, but they fall short in blocking targeted attacks.

7.3 Recommendations

With the emerging threats of sophisticated targeted attacks, our study shows that the power distribution system operators need to put much more effort into detection and response. Based on the current practices among the DSOs and the nature of targeted attacks, we advocate that strengthening the capabilities of the human operators is more important and will have greater impact than investing in more advanced technical tools at this point. The following is a list of prioritized recommendations with the aim of increasing their cyber situation awareness:

1. Perform regular emergency preparedness exercises involving IT attacks. It is a paradox that all respondents state the worst case scenario as one that would have severe consequences - attackers gaining control of power switches and being able to cause power outages in large areas - but still, none of the DSOs have performed any preparedness exercises based on this scenario. New regulations since July 2013 require the DSOs in this country to include IT attacks in their exercises. It remains to be seen how long it will take the DSOs to comply with these requirements. We provide this as our primary recommendation and appreciate that the authorities share the view of the importance of this matter.

2. Prepare for withstanding social engineering attacks. Targeted attacks tend to use social engineering as one of the strategies to collect information and build trust. This was not brought

DSO	C&C Location		Initial Attack		Lateral Movement		Final Step	
	Internal	External	Automatic	Manual	Automatic	Manual	Exfiltration	Sabotage
A	F	F	F		F		F	
B	I*	F, I*	F, I*	P	F, I*	P	F, I*	E
C	F	F	F	A	F	A,P	F	
D	I	I	I	A	I, S	A,P	S	
E	F	F	F	A	F	A	F	P
F	I	I	I	P	I	P	I	

Table 3: Summary of the potential countermeasures employed by the different DSOs to the various steps of a targeted attack. The countermeasures can be a firewall (F), an IPS (I), an anti-virus product (A), a company policy (P), awareness generated through an incident exercise (E), or physical separation (S).

up by the respondents as a concern. Although general information security awareness campaigns are run by some of the DSOs, they do not consider social engineering. With the current and emerging threats towards control systems, we recommend that the DSOs run preparedness exercises where social engineering is one of the ingredients.

3. Implement physical separation between industrial control systems and other networks. We are aware that the current development goes in the opposite direction, from having completely isolated industrial control systems to connecting them to the corporate and other networks. The reasons for this include efficiency and functionality. Information security should never stop business, hence we see this recommendation as unreasonable. Still, from a security point of view, this is a very efficient way of preventing attacks.

4. Deploy anomaly detection. With the current mechanisms, attackers could be inside the systems for a long time without anyone detecting anything. The operator's standard response of shutting down the system will not be performed if no malicious activity is detected.

5. Use regulations as a means of ensuring improvements. The DSOs strive to comply with national regulations in general. The authorities have a major responsibility in posing appropriate and sufficient requirements, which aid the DSOs in prioritizing the right tasks. Although there is a risk of ending up with too many regulations, the most important measures should be stated as requirements to enforce implementation.

A well-justified question is whether an average DSO should be able to resist all kinds of targeted attacks. As the respondent from DSO E pointed out, if a well-organized, powerful attacker group sets out to attack them, they will succeed, since the DSO does not have the resources, nor the aim, of being able to withstand such an attack. This concerns the balance of protection and accepted risk, which is the core decision in the overall information security management process.

With smart meters there will be a two-way communication with all power consumers. There are reasons to believe that this communication will include the control systems in one way or the other, i.e. by the use of a Distribution Management System (DMS). Also, the smart meters will bring along large amounts of personal data. These two factors might make the DSOs even more interesting as targets for attacks. We did however not cover the consequences of the introduction of smart meters in detail in this study.

7.4 Validity of the Study

Construct validity. We carefully designed the interview guide with the aim of covering the CSA capabilities as presented in Section 4 while at the same time being usable in the setting of an online meeting not lasting longer than approximately one hour to avoid respondent fatigue [10]. The interviewees may be biased, either consciously or unconsciously [12]. As the interview guide was distributed in advance, some preparations could be done. Furthermore, a trust relationship between the researcher and the interviewees is key to obtain honest responses. This trust was already established before this study due to previous collaborations[7]. Our impression is that the interviewees provided honest answers, revealing vulnerabilities, weaknesses, and the need for improvements.

It was rather challenging to cover all details in our questions, as responsibilities and expertise are distributed among a number of persons and roles: control room manager, control system operator, control system supplier, operator of network infrastructure, IT supplier, and information security manager. To get more complete answers, we could have interviewed a selection of personnel from each organization, including suppliers. Due to the limited time available, this was not possible for this study. It would also be difficult to get the participating organizations to commit this much time and personnel.

All interviewees were provided with a draft of this paper and given the opportunity to comment on the results. This is referred to as member checking [22], and is a strategy for reducing researcher bias. As only one researcher did most of the analysis of the interviews, this was especially important. We received one comment regarding the role and responsibilities of one of the respondents, and the description was revised accordingly.

External validity. The goal of a qualitative study is usually to investigate a specific case and provide a deep understanding of it, rather than generalization. Our detailed description of the industrial case provides the basis for considering its results' applicability to other settings. The participating organizations are among the largest from the specific industry, and should hence be expected to be in the lead. It takes a considerable amount of work to build trust with organizations in order to achieve the level of access to DSOs that is necessary to perform this type of study. This trust comes after years of interaction. Unfortunately, the downside of this is that because building trust is time consuming, one often has to be satisfied with a smaller and possibly narrower sample. Still, we consider our selection of the six DSOs to be representative for all large DSOs in Norway. When it comes to qualitative studies, generalizability is strengthened by an increase in the number of studies.

[7] Some of the interviewees were new to the researcher (the one from A and two from D), but there was already a trust relationship established between the organization and the researcher.

8. CONCLUSION AND FURTHER WORK

We assessed the preparedness of organizations for detecting and responding to targeted attacks towards industrial control systems. Distribution system operators from the electric power industry (as appealing targets for such attacks) participated in our interview study. Our findings indicate that they are not well prepared for this type of threat. They are not even ready for more traditional information security threats. They lack tools for monitoring and detecting attacks, as well as systematic approaches to follow-up on logs and alerts. Each organization is responsible for securing their control systems and ensuring continuous operation of the power supply. Hence, they need to be well-informed about current and emerging threats and to develop the necessary capabilities for addressing these.

Regulations work. Although all interviewed DSOs lacked the incident response capabilities, they all had periodic risk assessments and vulnerability and dependency analyses. We think that this strong inconsistency in the level of defense capability is due to the inconsistency in the laws regulating the cyber defense capability requirements for DSOs. While all DSOs were obligated to conduct periodic dependency, vulnerability, and risk analyses, the requirements for incident response capabilities are quite vague. This great difference in practice due to differences in regulations shows that regulations work and that the right laws can improve the state of cyber defense significantly. Based on this observation, we suggest that passing regulations obligating DSOs to have active security monitoring systems, anomaly-based intrusion detection systems, extensive system wide logs, extensive and correct use of cryptography, and regular incident response practices can significantly improve the current state of affairs.

Misconception and wrong threat perception. We noticed that the interviewed DSOs downplayed the probability of a successful cyber attack. Two DSOs specifically said that conducting a physical attack would be easier for the adversary. We find this notion contradictory with the Department of Homeland Security report claiming that the power industry reported the highest number of incidents in 2013 [3]. We argue that a possible factor contributing to this perception might be the higher than average security level in these companies. This security might act as a double-edged sword. While it makes it more difficult for cyber criminals to successfully break into the system, it also gives a false or excessive sense of security to the employees. We suggest that systematically providing security awareness training to the technical staff of DSOs, to make them aware of the actuality of the threats, can significantly improve their perception of the importance of cyber defense mechanisms.

Further work. We believe that studies like ours provide good-precision estimations of what is current practice in organizations, and we would like to encourage similar studies to be carried out in the future. Stronger evidence is needed for making conclusions about the industrial control systems in general. Furthermore, there are several parties that have responsibilities related to incident management, preparedness, and cyber situation awareness. Outsourcing is one main reason for this, and organizational structures is another. Including the suppliers of IT systems and control systems in future studies would enrich our knowledge and understanding of the challenges and needs, both for practitioners and researchers.

9. ACKNOWLEDGMENTS

This work was supported by the Office of Naval Research (ONR) under Grant N000140911042, the Army Research Office (ARO) under grant W911NF0910553, Secure Business Austria, and the Norwegian Research Council, grant numbers 234644/F11 and 217528.

The authors would like to thank the DSOs and specifically all of the respondents that contributed to this study. Further, we appreciate the feedback provided by Senior research scientist Martin G. Jaatun (SINTEF) and Cyber security manager Olav Mo (ABB) during the development of our interview guide.

10. REFERENCES

[1] Operation Aurora. http://en.wikipedia.org/wiki/Operation_Aurora, 2010.

[2] Global Energy Cyberattacks: "Night Dragon". Technical report, McAfee, 2011.

[3] ICS-CERT Monitor, Oct/Nov/Dec 2013. `https://ics-cert.us-cert.gov/sites/default/files/Monitors/ICS-CERT_Monitor_Oct-Dec2013.pdf`.

[4] Unveiling "Careto" - The Masked APT. Kaspersky Lab, February 2014.

[5] D. Albright, P. Brannan, and C. Walrond. Did Stuxnet take out 1000 centrifuges at the Natanz enrichment plant? Technical report, Institute for Science and International Security (ISIS), 2010.

[6] D. Albright, P. Brannan, and C. Walrond. Stuxnet Malware and Natanz: Update of ISIS December 22, 2010 Report. Technical report, Institute for Science and International Security (ISIS), 2011.

[7] R. Anderson, C. Barton, R. Böhme, R. Clayton, M. Eeten, M. Levi, T. Moore, and S. Savage. Measuring the Cost of Cybercrime. In *11th Workshop on the Economics of Information Security (WEIS'12)*, 2012.

[8] P. Barford, M. Dacier, T. Dietterich, M. Fredrikson, J. Giffin, S. Jajodia, S. Jha, J. Li, P. Liu, P. Ning, X. Ou, D. Song, L. Strater, V. Swarup, G. Tadda, C. Wang, and J. Yen. Cyber SA: Situational Awareness for Cyber Defense. In S. Jajodia, P. Liu, V. Swarup, and C. Wang, editors, *Cyber Situational Awareness*, volume 46 of *Advances in Information Security*, pages 3–13. Springer US, 2010.

[9] D. Batchelder, J. Blackbird, D. Felstead, P. Henry, J. Jones, and A. Kulkarni. Microsoft Security Intelligence Report. Microsoft, 2014.

[10] P. Ben-Nun. *Respondent Fatigue*, pages 743–744. Sage Publications, Inc., 1st edition, 2008.

[11] J. Dai, X. Sun, P. Liu, and N. Giacobe. Gaining Big Picture Awareness through an Interconnected Cross-Layer Situation Knowledge Reference Model. In *International Conference on Cyber Security (CyberSecurity) 2012*, pages 83–92, Dec 2012.

[12] T. Diefenbach. Are case studies more than sophisticated storytelling?: Methodological problems of qualitative empirical research mainly based on semi-structured interviews. *Quality & Quantity*, 43(6):875–894, 2009.

[13] A. Doupé, M. Egele, B. Caillat, G. Stringhini, G. Yakin, A. Zand, L. Cavedon, and G. Vigna. Hit 'em Where it Hurts: A Live Security Exercise on Cyber Situational Awareness. In *Proceedings of the Annual Computer Security Applications Conference (ACSAC)*, Orlando, FL, December 2011.

[14] N. Falliere, L. Murchu, and E. Chien. W32. stuxnet dossier. *White paper, Symantec Corp., Security Response*, 2011.

[15] J. R. Goodall, A. D'Amico, and J. K. Kopylec. Camus: Automatically mapping Cyber Assets to Missions and Users. *MILCOM 2009 - 2009 IEEE Military Communications Conference*, pages 1–7, Oct. 2009.

[16] M. Grimaila, R. Mills, and L. Fortson. Improving the Cyber Incident Mission Impact Assessment Processes. In *4th Annual Workshop on Cyber Security and Information Intelligence Research*, 2008.

[17] R. Langner. Stuxnet: Dissecting a cyberwarfare weapon. *IEEE Security & Privacy*, 2011.

[18] M. B. Line. Why securing smart grids is not just a straightforward consultancy exercise. *Security and Communication Networks*, 7(1):160–174, 2014.

[19] M. B. Line, I. A. Tøndel, and M. G. Jaatun. Information security incident management: Planning for failure. In *8th International Conference on IT Security Incident Management and IT Forensics (IMF)*, pages 47–61, May 2014.

[20] C. Paul and K. Whitley. A Taxonomy of Cyber Awareness Questions for the User-Centered Design of Cyber Situation Awareness. In L. Marinos and I. Askoxylakis, editors, *Human Aspects of Information Security, Privacy, and Trust*, volume 8030 of *Lecture Notes in Computer Science*, pages 145–154. Springer Berlin Heidelberg, 2013.

[21] N. Provos, P. Mavrommatis, M. Rajab, and F. Monrose. All Your Iframes Point to Us. In *USENIX Security Symposium*, 2008.

[22] C. Robson. *Real world research*. John Wiley & Sons Ltd., 3rd edition, 2011.

[23] G. P. Tadda. Measuring performance of Cyber situation awareness systems. In *11th International Conference on Information Fusion*, pages 1–8, June 2008.

[24] C. Tankard. Advanced persistent threats and how to monitor and deter them. *Network security*, 2011.

[25] O. Thonnard, L. Bilge, G. O'Gorman, S. Kiernan, and M. Lee. Industrial Espionage and Targeted Attacks: Understanding the Characteristics of an Escalating Threat. In D. Balzarotti, S. Stolfo, and M. Cova, editors, *Research in Attacks, Intrusions, and Defenses*, volume 7462 of *Lecture Notes in Computer Science*, pages 64–85. Springer Berlin Heidelberg, 2012.

[26] C. Witchall and J. Chambers. Cyber incident response: Are business leaders ready? The Economist Intelligence Unit (EIU), 2014.

APPENDIX

A. INTERVIEW GUIDE

Cyber situation awareness: Targeted attacks towards industrial control systems

An information security incident is commonly defined as something that compromises confidentiality, integrity, and/or availability of information. In this interview we are focusing on targeted attacks rather than technical failures. Furthermore, only the industrial control systems in the DSO are in question. The administrative IT systems are outside the scope of this study.

General

1. What is your role in the organization?
2. How many operators work in the control room?
3. How many power subscribers do you serve?
4. Can you estimate the number of computers and running applications?
5. Did you ever observe an attack to your control systems? Or malware?
   (a) Would you consider any of them as targeted?
   (b) Are you aware of any successful attack in your control systems?
   (c) How were these detected?
   (d) Were you able to identify the attacker(s)?
6. Do you think that anyone could be interested in attacking your systems? Who could this be?
7. Do you have any customers that could be potential victims for targeted attacks?
   (a) Could such an attack also hit your organization?
8. What is the worst possible consequence of a targeted attack?

Policies

9. Have you performed any criticality assessment of resources (computers, applications, information items, other) in the control systems?
   (a) Do you know about dependencies between certain resources?
   (b) Do any resources become more critical at specific points in time, or do they always keep the same level of importance?
   (c) Do you feel that the level of protection for the most critical resources is appropriate?
   (d) How is the IT defense different for the critical resources vs. non-critical ones?
10. Do you perform regular cyber security assessment?
11. How do you deal with the reported vulnerabilities; what kind of patching regime do you have?
12. Do you have any documentation of the technical security mechanisms on the control systems?
13. How are the control systems connected to the administrative network in your organization (i.e., one-way flow of data, VPN, no connection at all)?
14. How do you deal with employees and/or external consultants bringing their own computer, or other devices like USB memory sticks and similar, into the control room?

Preparedness

15. Do you have response procedures for cyber attacks?
    (a) How are they different from other failure response procedures? (Graceful degradation, restoring backups, limiting access, removing malware?)
16. Are control room operators made aware of the threats that they can encounter in their day-to-day job?
17. Have you ever performed exercises based on a scenario of targeted attacks towards the control systems?
    (a) Why/why not?
    (b) If any, were they table-top exercises or more realistic action-based exercises?
    (c) Do you have regular simulated attack practices?
    (d) Do you practice the worst-case scenarios?
18. What would be a beneficial way of training for responding to targeted attacks towards your control systems?

Technical security mechanisms

19. Do you encrypt critical data items while in transfer and stored?
20. Do you have off-site backups?
21. Do you only have network-edge defenses (e.g., IPSes), or do you also have detection mechanisms that can detect malicious activity inside the network?

    (a) Are such defenses host-based (antiviruses) or do they look at network traffic too?

22. Which specific defenses do you have? For each mechanism, use the following keywords to guide the conversation:

    (a) What is the purpose of this mechanism?

    i. Does it detect attacks?

    ii. Does it prevent attacks?

    iii. Does it react to attacks?

    iv. Does it predict attacks?

    v. Does it give more information about an attack that has already happened?

    (b) Input

    i. Which type(s) of input is needed (e.g., network, OS, service, or organization level logs)?

    ii. Where does the input come from (e.g., is it automatically deduced by the device, entered manually, or the output of another device)? Is input needed initially or continuously? How often is it entered? How many man-hours per week are required? How sophisticated is the input? Does it need to be configured, or is it a black-box system?

    iii. How high-level is the input? Is it human readable or raw?

    (c) Output

    i. Which type(s) of output is generated (i.e., attack alerts, network/OS/service/organization level logs)?

    ii. Where does the output go (i.e., to a human analyst or to another system)?

    iii. How high-level is the output? Is it human readable or raw? What is the size of the output?

    iv. Is the output actionable? Is it connected to an automated system? Does the action need human intervention?

    (d) Integration with the workflow and organization missions

    i. Is the system constantly running, do you run it when something happens, or is it run periodically?

    ii. Does the system need a human analyzer to be run, or do you run it and leave it be? If the system needs human configuration/input/intervention/analysis, how often does it happen? Is there a position/duty in the organization associated with it?

    iii. Is the tool/technique applied to the organization‘s most critical resources or to the whole organization? Are the most critical assets more protected, or monitored more often? How is the application of the tool different for a not-so-important resource and a mission-critical resource?

    (e) Internal model

    i. Is the model static or dynamic? Does it learn and change through time? Does it need initial/ongoing configuration? Does it learn (supervised/unsupervised; i.e., does it need human intervention for learning or does it learn on its own?)?

    ii. Does it learn the systems normal behavior? Does it learn the attacker's goals? Does it predict the next steps of the attacker?

    (f) Efficiency

    i. Does it work satisfactorily, or do you see any needs for improvements?

    ii. What is the amount of information that a security administrator (or all of them) should look at manually and daily? (Either in bytes, or lines, or pages)

    iii. What is the amount of information generated daily by the security logging tools?

    iv. What is the number of attacks reported daily/monthly/annually (either false positive or true positive)?

    v. How many of them, after manual inspection, turn out to be true?

B. THE CSA MAPPING

The questionnaire was purposely designed to cover in-place defenses and policies, incident response capabilities, and cyber situation awareness. Table 4 shows the relation between these areas and the questions. Each row represents one of the capabilities understudy and the numbers in each row are the number of the questions that are trying to evaluate the associated capability.

- General: general information about the organization
- CSA-Comprehension: comprehension of the current situation
- CSA-Impact: impact assessment
- CSA-Evolution: understanding how attacks evolve
- CSA-Behavior: attacker behavior analysis
- CSA-Causes: attack causal analysis
- CSA-Confidence: confidence in the acquired information
- CSA-prediction: prediction of future attacks or future steps of an attacker
- Defenses: the technology-based cyber defenses in place
- Policies: the policy-based cyber defenses in place
- Response: the incident response capability.

Category							
General	1	2	3	12			
CSA-Comprehension	4	5	21	22			
CSA-Impact	8	9	10	22			
CSA-Evolution	22						
CSA-Behavior	6	7	16	22			
CSA-Causes	22						
CSA-Confidence	22						
CSA-Prediction	22						
Defenses	5	9	13	14	19	21	
Policies	9	10	11	14	17	19	20
Response	11	15	16	17	18		

Table 4: Mapping between CSA capabilities and the questions.