FP7 Grant Agreement N° 607109

Security for smart Electricity GRIDs

D1.8 SEGRID DSO oriented Operational Security Capability Model

Due date of deliverable: 30/11/2017
Actual submission date: 01/12/2017
Work package: WP 1 Use Case & Security Goals
Number of pages: 22
Revision: Version 1
Classification level: Public

Project type:	Collaborative project – small or medium scale focused research project
Thematic Priority:	FP7-SEC-2013-1
Start date of project:	October 1st, 2014
Duration:	39 months

This project has received funding from the European Union's Seventh Framework Programme for research, technological development and demonstration under grant agreement no. 607109

Author(s)	Maarten Hoeve (ENCS)
Contributor(s)	Frank Fransen (TNO)
Reviewers	Rob Epskamp (Alliander) Nuno Emanuel Pereira (EDP) Youssef Ait Beh Ouali (ENCS)

Security Assessment	**See deliverables table for security assessment requirements**
Approval Date	30/11/2017
Remarks	None

TABLE OF CONTENTS

1 INTRODUCTION 4
1.1 The SEGRID Project 4
1.2 The objective of deliverable D1.8 5
2 SECURITY OPERATIONS CAPABILITIES 7
2.1 Vulnerability Management 7
2.1.1 Identify Host Vulnerabilities 7
2.1.2 Identify Architecture Vulnerabilities 8
2.1.3 Coordinate Resolving Vulnerabilities 9
2.2 Incident Management 10
2.2.1 Triage Anomaly Detection Events 10
2.2.2 Triage Misuse Detection Events 11
2.2.3 Coordinate Incident Response 11
2.2.4 Gather Forensic Evidence 12
2.2.5 Maintain Infrastructure 12
2.2.6 Hunt Threats 13
2.2.7 Analyse Incidents 13
2.3 Intelligence Gathering 14
2.3.1 Keep Contact with Information Sharing Groups 14
2.3.2 Process Online Intelligence 15
2.3.3 Analyse Threat Actors 15
2.4 Logging & Monitoring 16
2.4.1 Maintain Anomaly Sensors 16
2.4.2 Collect Logs 17
2.4.3 Maintain Misuse Sensors 17
2.4.4 Create Correlation Rules 18
3 RESPONSIBILITIES 19
3.1 Challenges in Setting up a Security Operations Team 19
3.2 Skills Required 20
4 ROADMAP 22

1 Introduction

1.1 The SEGRID Project

SEGRID's main objective is to enhance the protection of Smart Grids against cyber-attacks. In the coming years, the level of automation in electricity distribution grids will grow substantially. Smart meters will be deployed at home premises, and remote terminal units (RTUs) will be placed in distribution substations. The increased automation should provide a better view of how electricity flows to the medium and low voltage grids, and provide grid operators increased control to influence that flow. But the increased automation also has major consequences for the cyber security posture of the electricity grid. By nature, automation adds new attack vectors through which cyber attackers can enter and attack the networks of grid operators, and offers other possibilities to cause physical damage to the electricity grid itself. Moreover, there is an increased risk of physical damage to the electricity grid due to inadvertent events in the cyber domain of Smart Grids (e.g. malware, technical dependencies, and human errors).

From a technical point of view, it is not sufficient to only consider all the different components in a Smart Grid separately; they will together form a truly integrated system-of-systems and the Smart Grid will neither be completely owned, nor completely controlled, by a single power system operator. There will be many Smart Grid services and components that are operated by other organizations, such as public telecom networks and third party-delivered (outsourced) application services. There will potentially be many new methods for connecting with various Smart Grid applications using a diverse set of communication channels, such as local connection interfaces, distributed web access, and smart apps on smart phones. A number of new cyber security issues become critical in this context.

This new utility-wide system(-of-systems) will not come into existence overnight; the Smart Grid will be composed of a mix of old, even legacy, and new components. Therefore, we look upon the Smart Grid as a gradually evolving system in which new functionalities are added to accommodate new use cases with the challenge to maintain security, privacy and dependability of the Smart Grid as a whole. The required security solutions to cope with this situation range from improvements in current security, privacy and operational solutions and procedures to the introduction of new security and privacy paradigms.

The term 'Smart Grid' is used for a wide variety of developments related to the automation of electricity grid. Focusing on the efforts of the project, SEGRID worked on five concrete use cases, shown in Figure 1. These have been further elaborated in deliverable D1.1 of the project. The five SEGRID use cases are:

1. Smart meters used for on-line reading of consumption and technical data;
2. Load balancing renewable energy centrally;
3. Dynamic power management for smart homes, smart offices, and electric vehicles;
4. Load balancing renewable energy regionally (substation automation);
5. Automatic reconfiguration of the power grid.

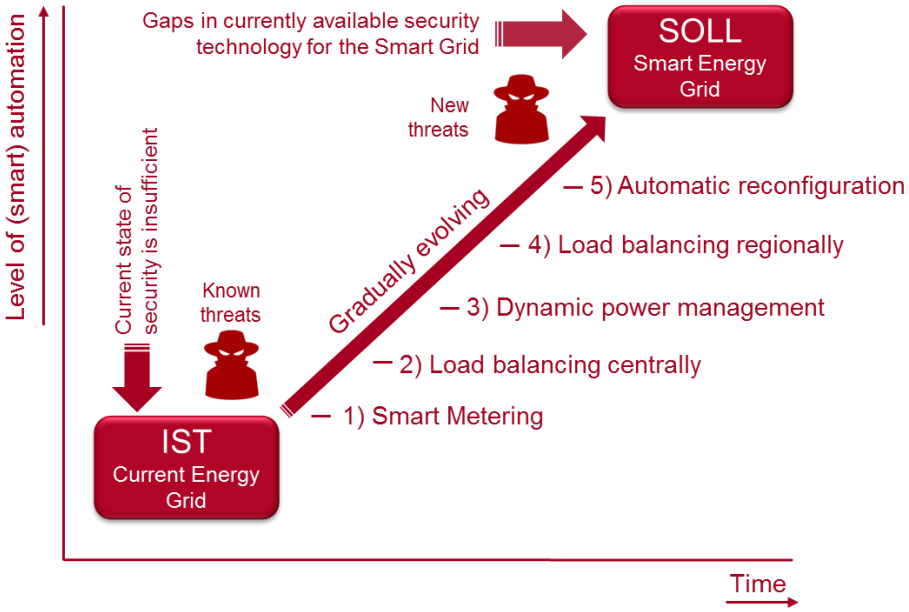

Figure 1: SEGRID storyline

The SEGRID use cases reflect important steps of Smart Grid developments until 2020 and beyond, and cover the most relevant security and privacy issues that will arise from the increasing complexity of Smart Grids.

SEGRID formulated privacy and security goals for the use cases defined (WP1) and performed a risk analysis of these use cases (WP2). Vulnerability assessment frameworks and tools to identify these risks in real systems were developed (WP3) and gaps in the currently available security solutions were identified (WP2). Solutions for some of these gaps were developed in WP4. Selected vulnerability assessment tools and novel security solutions were tested in SITE (WP5).

1.2 The objective of deliverable D1.8

Smart grid security work has initially focused on system design. The new systems for smart metering, distribution automation, electric vehicles, and demand response all require a security architecture, and the selection of security measures. At the same time, existing (legacy) infrastructures need to be redesigned to improve their security. Security officers at grid operators have therefore worked to assess the risks of smart grid systems, and select mitigating security measures.

As smart grid systems become operational, grid operators will also need to build an operational security capability. They need to be able to identify and deal with incidents and vulnerabilities on a day-to-day basis. As a first step, many grid operators are now setting up security monitoring. They are purchasing intrusion detection systems for their critical SCADA systems, and setting up a team (such as a CSIRT, CERT, or SOC) to analyze and respond to the alerts.

The responsibilities of these smart grid security operations teams will grow in the future, as it has done in other industries such as telecommunication and banking. They will take on tasks such as incident response, vulnerability scanning, forensics, managing firewalls, security assessments, and penetration testing.

This document analyzes the capabilities a smart grid security operations team will need in the future. To attract and train personnel with the right skills, grid operators require a long-term view of the capabilities the team needs. This document describes a model that defines the required capabilities, and to give a roadmap to develop these vulnerabilities, from the current pilots with monitoring to having a mature security operations team for the future smart grid.

2 Security Operations Capabilities

This section describes different security operations capabilities that a grid operator may need.

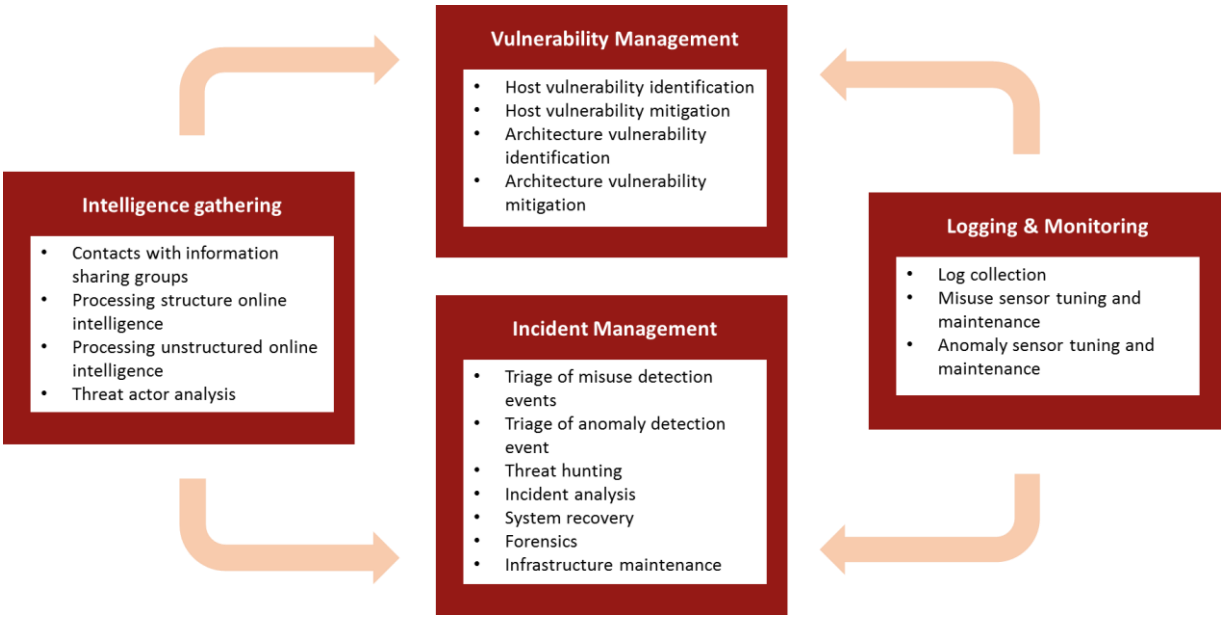

2.1 Vulnerability Management

Incident management capabilities are all capabilities that help an organization to identify and fix or mitigate vulnerabilities. The capabilities in this section are related to the Technical vulnerability management section in ISO 27002. In particular, they help to implement the following controls:

- Employees and contractors using the organization's information systems and services should be required to note and report any observed or suspected information security weaknesses in systems or services. (Section 16.1.2)
- Information about technical vulnerabilities of information systems being used should be obtained in a timely fashion, the organization's exposure to such vulnerabilities evaluated and appropriate measures taken to address the associated risk. (Section 12.6.1)

2.1.1 Identify Host Vulnerabilities

Objective	To identify vulnerabilities on individual hosts in the network.

Activities	• Perform automatic vulnerability scanning on non-legacy systems
	• Identify vulnerabilities on legacy systems through
	  ○ Reviewing installed software versions
	  ○ Active scanning on test systems
	  ○ Passive scanning
	• Coordinate regular security assessments and penetration tests

Many hosts in smart grid systems can be treated as modern IT systems. This is especially true for hosts at the boundary with enterprise IT systems, such as systems in a demilitarized zone used to exchange

data, jump servers used for vendor remote maintenance, and laptops and workstations used by engineers. Such systems should be kept fully patched. Patching can be centrally managed, and active scanning can be used to track the patch levels.

Vulnerability management gets more complex for critical operational technology hosts, such as SCADA servers, and the hosts at substations. SCADA servers and secondary substations can usually only be patched working with their vendor. Hosts at primary substations is in many cases not patched at all, because of operational constraints. Vulnerability identification should take such restrictions into account. There should be a clear, risk-based policy on when vulnerabilities require action. For instance, vulnerabilities in substations may only be patched if they can be exploited remotely. This means that in many cases host vulnerabilities must also be considered from an architectural perspective.

Vulnerability scanning, in which the vulnerability scanner probes services over the network, is considered dangerous for legacy systems. They may freeze or crash because of the unusual data sent by the vulnerability scanner. Active scanning should not be used on equipment in substations, or within the SCADA system (unless thorough testing has shown that it is possible).

There are several alternatives. First, the installed software versions can be tracked manually, and compared to information from vulnerability databases. This can be labor-intensive. Good information on installed software is often lacking, and will need to be gathered. Change management processes would have to be improved to keep the version information up to date.

Second, scans can be performed on test environments, as far as these are available. Such an approach can give a lot of information. But there is a risk that the test environment is different from the real environment.

Third, information can be gathered by passive scanning, listening to normal network traffic. This in non-intrusive, and can be used without risk on legacy systems. Information can be gathered on the use of unencrypted traffic, deprecated cryptographic algorithms, or protocols, and in some cases on software versions running through headers. Many intrusion detection systems developed for SCADA systems now support such functionality. The information gathered will however be incomplete. Hence, it is dangerous to completely rely on this approach.

The security operations team should be capable of using the alternative vulnerability identification methods. This requires not just technical skill, but also great discipline to perform the manual actions without missing vulnerabilities.

Security assessments or penetration tests by external parties should be used to complement the internal vulnerability identification activities. They can be used to validate how well the vulnerability management process works, and to provide in-depth assessments of security-critical systems or systems for which there are suspicions of vulnerabilities.

2.1.2 Identify Architecture Vulnerabilities

Objective	To identify vulnerabilities in the network and system architecture.
Activities	• Prepare network maps • Actively scan the network • Analyze network traffic • Review firewall rules

	• Review WAN security measures

The security of many operational technology (OT) systems depends on the network architecture. As discussed in Section 2.1.1, it is not always possible to keep operating systems and applications up to date. Parts of the system that cannot be updated should be shielded from the outside using firewalls, and buffers such as demilitarized zones. This capability concerns identifying possible vulnerabilities in such measures.

The activity starts with understanding the network. The security operations team needs to know which hosts are in the network, and how they communicate. It is not possible to protect hosts you do not know about. These may even become backdoors into the network.

Grid operators have a good view of the networks at the control centers. They can improve the view by actively scanning non-legacy zones, such as DMZs, and by listening to the traffic. SCADA IDS systems can automatically generate asset lists, and network maps from traffic.

The information coming from these activities, should be used to regularly review the firewall rules. The security of the whole system often depends on the firewalls being correctly configured. So, there should be checks that only allowed traffic can go through.

Grid operators often know much less about what is in the substations. Substations can stay in operation for twenty years or more. Every few years there are new designs and technologies introduced. Hence, there are often many generations of substations in the field. If no strict change management is used, the asset inventories of the substations can be out of date. The only way to understand the network is often to visit the substation, and go over all equipment and wires that are found there.

Often this task is not feasible. Instead, the architecture should be designed to not depend on the security of the inside of substations. To stop the spread of malware, the substation inside should be shielded from the central systems. Laptops that are used to maintain equipment in the substations should be kept extra secure. Measures should be taken to stop one compromised substation affecting others or the central systems. The vulnerability identification activities should then be aimed at verifying that such measures have been correctly implemented.

2.1.3 Coordinate Resolving Vulnerabilities

Objective	To coordinate the steps required to resolve identified vulnerabilities.
Activities	• Coordinate patching for off-the-shelve software • Coordinate fixing vulnerabilities for custom software (project to implement fixes by vendor, and then test the fixes) • Coordinate verification of vulnerability fixes

Getting vulnerabilities resolved will often require work from multiple parties. Often the vendor of the affected system needs to get involved. Within the grid operator, the departments maintaining the systems will need to do the updates. In some cases, procurement is needed to negotiate with the vendors. The control center needs to be informed if there are possible disruptions to operations. In some cases, service providers, such as the telecom operator or parties managing parts of the network, need to be involved.

The security operations cannot take responsibility for all this work. They do however have a responsibility in coordinating that identified vulnerabilities are fixed in a reasonable time. When a vulnerability

has been identified, they should make sure the parties responsible for fixing it are informed, and that these make a plan to fix them. They should then track if the plan is executed.

When the vulnerability is expected to be fixed, they should make sure it is verified that it indeed is. This can be done by the security operations team, using the same methods they used to initially identify the vulnerability. In some cases, it may however be required to have tests performed by the vendor of by external penetration test companies.

2.2 Incident Management

Incident management capabilities are all capabilities that help an organization to identify and respond to security incidents. The capabilities in this section are related to the Information security incident management section in ISO 27002. In particular, they help to implement the following controls:

- Information security events should be reported through appropriate management channels as quickly as possible. (Section 16.1.2)
- Information security events should be assessed and it should be decided if they are to be classified as information security incidents. (Section 16.1.4)
- Information security incidents should be responded to in accordance with the documented procedures. (Section 16.1.5)
- The organization should define and apply procedures for the identification, collection, acquisition and preservation of information, which can serve as evidence. (Section 16.1.7)

2.2.1 Triage Anomaly Detection Events

Objective	To identify security incidents based on alerts from anomaly detection sensors, and prioritize the identified incidents.
Activities	• Filter security alerts from operational alerts • Prioritize security incidents based on impact

Anomaly detection sensors are applications or hardware appliances that search for deviations from normal behaviour. The sensors keep a baseline of what is normal in a network or system. Often this baseline is initially learned updated automatically, and then tweaked manually (see Section 2.4.1). The sensor then searches for things that do not fit in the baseline.

Anomaly detection sensors are quite common in smart grid systems. The communications in operational technology (OT) networks are regular and stable, making them well-suited for anomaly-detection techniques. And because there are few known attacks, misuse detection sensors (see Section 2.2.2) are often perceived as having limited coverage. Special network-based intrusion detection sensors have been developed for OT. Many grid operators are now deploying them.

False positives have a different nature for anomaly detection sensors than for misuse detection sensors. Any alert indicates a real deviation from what is normal (leading some sensor developer to opportunistically claim their system has no false positives). The question however is if the cause of the deviation is a security incident, or something operations, such as an accidental failure, a misconfiguration, or an unusual system state or operation. Triage is in the first place concerned with answering this question.

Because of this, it is in many cases best to have the initial triage on the anomaly detection done by operational departments, such as a network operations centre, or departments responsible for maintenance on SCADA systems or substations. Experience with OT anomaly detection sensors shows that

most alerts are caused by operational problem or misconfigurations. Operational departments like having this type of information, because it helps them catch potential problems before they lead to downtime. They are also better able to identify operational causes of alerts. Only if they cannot find a good operational reason, they should escalate the event to the security operations team.

Raw alert data from the anomaly detection data can be imported into a SIEM system to use in correlation rules, as part of misuse detection activities, or to use in trend analysis in threat hunting.

2.2.2 Triage Misuse Detection Events

Objective	To identify security incidents based on alerts from misuse detection sensors, and prioritize the identified incidents.
Activities	• Filter out false positive alerts, and identify security incidents • Prioritize security incidents based on their impact

Misuse sensors are software applications or hardware appliances that search for signs of known attack. Classical examples are signature-based network intrusion detection systems, or anti-virus programs. But more advanced approaches, such as malware "detonation" solutions, also fall in this category.

Even when misuse detection sensors are properly tweaked (see Section 0), a large part of their alerts will be false positives that are not related to a security incident. Triage is meant in the first place to filter out such false positives. Once it has been established that the alert is related to an incident, the incident is usually prioritized for further analysis and response. Usually some heuristic for the impact of the incident is used. The follow-up actions then depend on the impact categories.

In larger teams, triage can be split into formally defined lines. The first line handles simple incidents according to clearly defined rule. Any alert that cannot be handled quickly, often in a fixed time limit, is escalated to a second line. The second line does more in-depth analysis of more complex incidents. They can stay on these incidents until they are closed. In some cases, specialized expertise can be called in from a third line.

Having such response lines is useful when part of the security operations is outsourced. Organizations can for instance choose to outsource the first line analysis to a managed security service provider. This only work if the responsibilities of the lines is clearly defined.

2.2.3 Coordinate Incident Response

Objective	To coordinate all activities needed to analyse, contain, and recover from an incident.
Activities	• Initiate and track incident response actions. • Communicate with affected parties • Identify containment and recovery actions from the incident analysis • Plan and track containment and recovery actions

Smaller incidents may be handled by a security operations analyst on its own, or with help of a system administrator. When an incident has a lot of impact, many parties may however get involved. Resolving the incident may involve different people in the IT and OT departments. The control centre may need to be informed if the incident affects the electricity grid. Responsible managers also need to be informed.

Regulation such as the GDPR and NIS also requires that government organizations are informed in some cases. In very large incidents, crisis management processes need to be started.

The security operations team should have the capability to coordinate all these stakeholders. Clear and careful communication about incidents is keys. Moreover, if many people are involved in responding to an incident, their actions should be planned, so that they do not interfere with each other, and tracked, so that the incident can be closed correctly.

When a security incident is suspected, steps should be taken to contain it, so that it cannot spread to other parts of the system. In some cases, hosts suspected to be compromised can be simply shut down, or removed from the network. Critical systems may however need to be kept running. In that case, the hosts may be separated from others through firewall rules and other networking measures. The actions of the host may be monitored with extra care.

After an incident, systems need to be restored to a secure, operational state. Malware or backdoors that were discovered need to be removed. In some cases, hosts need to be restored from a back-up.

The security operations team cannot do this alone. The system administrators or engineers responsible for a system should do the actual containing or restoring. The security operations team should however coordinate their actions to ensure that all problems found when analyzing an incident are resolved. They should also ensure that the containment and recovery do not interfere with the analysis and forensics on an incident.

2.2.4 Gather Forensic Evidence

Objective	To collect and preserve evidence of security incidents for legal and disciplinary action.
Activities	• Handle forensic artifacts • Prepare a forensic timeline of an event • Present forensic evidence in court

Serious incidents may lead to legal prosecution. Preparing evidence from an incident in such a way that it is legally valid in court requires special knowledge. Analysts need to be specially trained for this.

2.2.5 Maintain Infrastructure

Objective	To maintain the information processing infrastructure used for incident management.
Activities	• Maintenance of log servers or SIEM system

The data processing infrastructure used in incident response need to be maintained. Maintaining and tuning sensors is considered part of the logging and monitoring activities. This capability concerns maintaining the central systems used to collect and process security events.

In the simplest case, the infrastructure can consist of one server used to centrally collect and search the logs. Larger teams may deploy a SIEM system. Maintaining such a system takes considerable effort and expertise. Having a SIEM is therefore only recommended if the security operations team has time to maintain it.

2.2.6 Hunt Threats

Objective	To actively search for threat actor activities within the system.
Activities	• Analyze trends in statistics • Visual anomaly detection

The triage activities in Section 2.2.1 and 2.2.2 are reactive. Analysts respond to events generated by sensors. This approach works well with attack types that are already known. Sensors are however usually not capable of detecting new exploit or attack techniques used by advanced threats.

To detect these, analysts should actively hunt for threats. Rather than reacting to alerts they search for possible signs of attacks in security relevant data. Because they are searching for something new, the process is less structured than for triage. The results depend on the creativity and analytical skills of the analyst.

One approach analysts can take is to look at trends in statistics such as the number of connection or alerts, or the amount of traffic. In OT systems many of the statistics will be quite stable, showing little variation or growth. If they do show unusual trends, analysts can try to find the cause of this. In most cases, the cause will be operational. But in some cases, they may find signs of malicious attacks.

Another approach is to look for outliers, often with the help of visualization. Analysts can look at graphs of the amounts and types of traffic at different levels of aggregation. Strange phenomena often will clearly stand out in such graphs. They can also look at network graphs, or visualizations of the timing of packets. Many SCADA intrusion detection systems offer extensive capabilities to visualize network traffic. SIEM and log management systems offer tools to visualize events.

2.2.7 Analyse Incidents

Objective	To analyse identified security incidents to determine the hosts affected.
Activities	• Identify affected hosts • Identify impact of incident • Identify exploited vulnerabilities • Analyze malware used in incident

After an incident has been detected, through triage (Section 2.2.1 and 2.2.2) or threat hunting (Section 2.2.6), additional analysis is often required. It should be determined which host have been affected by the initial attack, and by steps attackers have taken before and after that. The actions taken by the attackers should be reconstructed to understand the full impact of the attack. Was confidential information extracted. Did the attackers build backdoors into the system? Were passwords or keys compromised? The vulnerabilities exploited by the attackers should be identified, so that they can be fixed.

If attackers try to hide their traces this type of analysis can become complex. Malware used in the incident may have to be investigated and reverse engineered. Attackers may also have deleted logs or other information used in the attack. So, this capability may be one of the most complex capabilities a security operations team can develop. Most teams will have only part of the capability in-house, and will need to get external support on specialist topics.

2.3 Intelligence Gathering

The capabilities in this section are used to gather intelligence about threats and vulnerabilities. The ISO 27002 standard only contains two controls related to incident gathering:

- Appropriate contacts with relevant authorities should be maintained. (Section 6.1.3)
- Appropriate contacts with special interest groups or other specialist security forums and professional associations should be maintained. (Section 6.1.4)

Gathering and using intelligence is however critical to keep up with the quickly changing security situation. Hence, we will take a wider view of this capability.

2.3.1 Keep Contact with Information Sharing Groups

Objective	To gather former and informal intelligence by participating in relevant security communities.
Activities	• Participate in vulnerability and incident information sharing groups • Establish contact with relevant authorities • Fix vulnerabilities based on information from groups • Translate information into detection signatures • Contribute vulnerability and incident analysis results in the groups

Information sharing groups are an important source of information about vulnerabilities and threats. Grid operators can participate in a variety of groups: national and European, government or industry, sector-specific or generic. The security operations teams should develop a strategy to gather all useful information from these groups, and not spend too much time in meetings.

Government groups have a special position. They have access to classified information not available in some other groups, and may have a special position in regulation. Most grid operators will want to establish contacts with these groups.

Initially the value of the information sharing groups will come from vulnerability information. Such information is easiest to act on. The software versions affected by the vulnerabilities is usually clearly identified, and mitigation options are given. To use this information, grid operators need to have an up-to-date view of the software in their systems. This is still not trivial with the many legacy systems in use.

Using information on threats is more challenging. Threat information is usually given as indicators of compromise for certain malware or attack campaigns. To use them, a grid operator must have the right sensors installed. If some networks or hosts are not monitored, the indicators cannot be used for them. To find attacks that happened in the past, historic data would have to be available. If the sensor data is there, the security operations team still needs to have the skill to put the indicators of compromise into the tools they are using (such as IDS or SIEM systems).

Finally, mature security operations teams can try to not only take information from the group, but also bring in new information. Doing so probably requires that the team is doing threat hunting (Section 2.2.6), and analyzing threats (Section 2.3.3). If new information is found in such activities, they can be presented in the group. Although this does not directly benefit the grid operator, it is vital to build a community that produces useful information.

2.3.2 Process Online Intelligence

Objective	To use online intelligence on incidents and vulnerabilities available online to improve the vulnerability and incident management processes.
Activities	• Update definitions in vulnerability scanners • Update signatures in misuse detection sensors • Process vulnerability and threat notifications • Process information from blogs and news items

Besides information from in-person groups, a lot of information on vulnerabilities and threats is available online. The information can be both in a structured form, in databases or notifications, or unstructured, in news items or blog posts. Security operations teams can benefit a lot from the capability to process this type of information.

The easiest source of online information are the vendors of the security tools, such as vulnerability scanners, intrusion detection systems, and SIEM systems. These vendors have teams that translate threat information in detection rules and signatures. Grid operators will get information on all major vulnerabilities and threats in this way.

By processing online sources themselves, grid operators can however sometimes get information on more specialized vulnerabilities and threats, and improve how quickly they can incorporate the information in their systems. The focus should be on specialized sources for smart grid systems. The security operations team should try to get notifications from the vendors of the SCADA systems and substation equipment. They may also want to follow vulnerability researchers specializing in SCADA and smart grids.

2.3.3 Analyse Threat Actors

Objective	To analyse security incidents to understand the method and motivation of the threat actors causing them.
Activities	• Determine the likely source of an attack • Analyze tools and techniques used in the attack • Determine the likely goal of an attack • Connect actions from the same threat actor to understand larger attack campaigns

Advanced security operations teams may want to attempt to find out who is behind incident they detect. By better understanding who is behind an attack, and what they are trying to achieve, they may be better able to react.

Analysts may try to find the source of incidents that have been detected in their systems. They could try to trace the IP addresses from which the attack originates. Or they may try to find clues in the tools and techniques used on who created them. Performing such an analysis requires a very high level of expertise, and knowledge about attacker techniques. Cooperation with specialists, possibly at intelligence services, may be desirable.

If an attack source can be identified with some certainty, it may be possible to guess the motivation and goals of the attacker. It may also be possible to connect seemingly unrelated incidents to see a pattern in the attack. This could help to help to more effectively counter advanced threats.

2.4 Logging & Monitoring

The capabilities in this section are used to collect the data needed to detect and analyse vulnerabilities and incidents. The capabilities help implement the controls in the "Logging and monitoring" section in the ISO 27002 standard (Section 12.4). Specifically, the capabilities can be used to implement the following controls:

- Event logs recording user activities, exceptions, faults and information security events should be produced, kept and regularly reviewed. (Section 12.4.1)
- Logging facilities and log information should be protected against tampering and unauthorized access. (Section 12.4.2)
- System administrator and system operator activities should be logged and the logs protected and regularly reviewed. (Section 12.4.3)
- The clocks of all relevant information processing systems within an organization or security domain should be synchronised to a single reference time source. (Section 12.4.4)

2.4.1 Maintain Anomaly Sensors

Objective	To configure anomaly detection sensors to reduce the rate of false positives, and make them more sensitive to signs of real incidents.
Activities	• Tweak and review baselines or models • Suppress incident types or categories

Most anomaly detection sensors are self-learning. They are first deployed in learning mode, where they made a baseline of the normal behavior of the network of system they monitor. In this way, the sensor is customized to the situation in that network or system.

After the automated learning phase, there is however usually a need to also manually tune the sensor. The main goal is to reduce false positive rates. The sensor may not be able to learn all normal behavior automatically. The sensor may for instance not see maintenance work on every substation during the learning phase. It will then generate alerts whenever maintenance is done on a substation (except if the substation happened to be maintained during learning). Human analysts can recognize the pattern of alerts cause by maintenance, and tune the sensor not to generate the anymore.

Besides tuning the sensors to reduce false positives, many sensors should also be reviewed to maintain their detection capabilities. During tuning, more and more causes of false positive are added to the baseline of normal behavior. The sensor becomes less sensitive.

Attackers may exploit this by slowly adding steps in their planned attack to the baseline, disguised as false positive. For instance, if they need a certain connection for their attack that is not in the baseline, they may first open that connection for a harmless purpose. Analysts will then likely classify it as a false positive, and add it to the baseline. When the real attack comes, the connection will no longer trigger an alert.

Analysts should therefore regularly review the baseline to check if all the behavior allowed by it, is indeed normal and needed for normal operations. This is not a simple task, as it requires a thorough

understanding of the sensor and the system being monitored. Some sensors do not even allow such reviews. But without this capability, the detection capabilities will likely degrade over time. In some cases, the same goal could be achieved by relearning the sensor after a certain period.

2.4.2 Collect Logs

Objective	To gather security relevant logs in a central location for analysis, and forensics.

Activities	• Set up a central log server
	• Set up infrastructure to collect logs from different parts of the system
	• Collect security logs of critical systems, including:
	  ○ Firewalls
	  ○ Operating systems logs of critical servers
	  ○ Application logs of critical servers

Besides alerts from sensors, many activities require logs generated by hosts in the system. Almost any host will generate log events that are useful in identifying and analyzing security incidents. The risk is rather being flooded in so much events that the systems and processes cannot handle them.

Security operations teams should have a conscious strategy and process to collect logs. Generally, it is recommended to start collecting logs based on certain use cases. If the security operations team wants to detect failed login attempts on routers, they start collecting router access logs. Logs may be collected in bulk for compliance reasons or to support big data inspired analysis tools.

In most cases, more logs will be collected over time. So, the log management process for integrating more new log sources. Often this requires collaboration with the departments maintaining the hosts. Getting data from substations or other field locations, may also require support from telecom departments. Making sure the time is synchronized on the logs sources is important to get usable data. In many cases, time synchronization is however already implemented for grid monitoring and control.

2.4.3 Maintain Misuse Sensors

Objective	To configure misuse detection sensors to reduce the rate of false positive, and make them more sensitive to signs of real incidents.
Activities	• Disable rule that generate large amounts of false positive • Connect asset information to the sensors if useful • Create custom signatures

Besides updating the misuse detection sensors with information from information groups and online sources (see Sections 2.3.1 and 2.3.2), misuse sensors often also need to be tuned to work well in a specific network or system. Signatures or rules that give many false positives may be removed. Asset information, such as the operating systems running, may be given to the sensors to get better detection. The security operations team may also create custom signatures for the sensor, for instance to monitor security policies in force in the system.

2.4.4 Create Correlation Rules

Objective	To create rules to correlate events for different sources to detect misuse.
Activities	• Tweak existing SIEM rules • Write new rules based on a risk assessment

Incidents are usually detected most effectively by combining events from different sources. Alerts from anomaly and misuse sensors may be combined with host logs to get more reliable detection of incidents. SIEM systems in particular can perform such correlations.

SIEM systems will often come with a set of correlation rules (use cases) supplied by their vendor. But to use them effectively, security operations teams need to have the capability to create their own rules. They can then either tweak the existing rules to work optimally in their environment. Or they can write new rules to detect risks that have been previously identified. Without a good understanding of the correlation engine, a lot of the potential of the SIEM system is unused.

3 Responsibilities

One of the main questions when implementing security operations is to identify who is responsible for the different activities. The capability model in Section 2 can help answer this question by providing a classification of the activities that may have to be performed, and the capabilities need to perform them.

The exact division of responsibilities will depend on the particulars of an organization. We will in this section only discuss the challenges, and the skills required for different capabilities.

3.1 Challenges in Setting up a Security Operations Team

Grid operators face many challenges when setting up security operations for the smart grid systems. The first challenge in setting up a security operations team is getting the budget. For many grid operators, the security operations team will mean a significant growth in security staff. Due to high demand in the market, the security operations staff may also require a high payment compared to other functions. Moreover, the required infrastructure and tools required also require an investment.

Grid operators are in a regulated industry. In the end, budgets are determined by the regulator, also based on the risks that grid operators face. There is an increasing awareness of cyber-security risks, both at grid operators themselves, and in government. And with the introduction of the smart grid, the medium and low voltage parts of the grid will be seen more as a critical infrastructure. Hence, it is likely that more budget will be allocated to security, so that the formation of security operations teams will become possible.

But even with budget available, it will be difficult to build a team with the right capabilities for the smart grids. Building a dedicated smart grid security operations team within the grid operator is difficult. Except in the case of largest grid operators, there will not be enough work to keep the team busy. For IT systems, most work for the security operations teams comes from dealing with the alerts from intrusion detection sensors, with enough incidents happening to keep the work somewhat interesting.

For smart grid systems, there will be both less alerts, and less incidents. There are less alerts, because most communication is machine-to-machine. The smart grid consists of a lot of computing devices with only a small number of people operating them. So, the normal behaviour in smart grids will be very regular, and not change much over time. Intrusion detection sensors can therefore be tuned to achieve very low false positive rates, for mature products often only a handful per day. There are less real incidents, because there are fewer actors interested in smart grids. Strong protection at the perimeter of the smart grid systems can keep out unskilled or untargeted attacks. The skilled and motivated attackers that can penetrate, are likely associated with nation states. They will try to stay hidden, and not cause incidents, unless there is a crisis. In these circumstances, it will be challenging to keep a dedicated smart grid security operations team motivated and focused.

Outsourcing is also difficult, because the team will require smart grid specific knowledge. It is therefore difficult to outsource the work. Managed security service providers that offer security operations are focussed on IT systems. Grid operators that have tried to outsource intrusion detection for their SCADA system have found that the service provider lacks the knowledge to analyse the alert. Hence, any alert from the sensors gets escalated to be handled by the grid operator itself.

Service providers could of course learn the general technologies used for SCADA. But because SCADA systems have a long history at each grid operator, they accumulate a lot of features specific to their user. These specifics are usually not completely documented, and are only known to the engineers of the

system. Acquiring specific knowledge for all customers would require a lot of investment in the service provider.

Over time the technologies used in the smart grid may become less specific. For newer systems, there is a trend to use standardized and open internet-of-things technology, such as Linux and web services. So as legacy systems are phased out, it may become easier for IT-focussed companies to monitor smart grid systems. On the other hand, the new smart grid systems will likely also stay in the field for a long time, becoming legacy themselves. At that point, they may again require specific knowledge to monitor.

3.2 Skills Required

To solve the challenges described above, it is important to know what type of skill is required for different capabilities. Some capabilities require mostly knowledge about the operational technology (OT) systems within the grid operator. Others require mostly security knowledge gained in IT systems. Others require both kinds of knowledge. See Table 1.

Both OT knowledge, and IT security knowledge are scarce. And people that possess both are especially hard to find. Market prices for such people are consequently high. For grid operators it is therefore necessary to not assign all responsibilities to such dual specialists. If a capability can be done by someone with OT knowledge, or only IT security knowledge, it should be assigned to that person. Table 1 can help with this.

Table 1: Type of skill required for different capabilities.

Mostly OT Knowledge	Mostly (IT) Security Knowledge	Both OT and Security Knowledge
Identify host vulnerabilities Identify architecture vulnerabilities Coordinate resolving vulnerabilities Keep contact with information sharing groups Triage anomaly detection events Maintain anomaly sensors	Process online intelligence Triage misuse detection events Coordinate incident response Gather forensic evidence Maintain infrastructure Collect logs Maintain misuse sensors Create correlation use cases	Analyse threat actors Hunt threats Analyse incidents

Vulnerability management capabilities are usually best assigned to people with OT knowledge. The maintenance of OT systems is usually not yet done through formal, well-documented processes. To effectively get the information needed to identify vulnerabilities, and fix them, requires not only knowing the OT systems well, but also knowing the people maintaining them. Often the only way to learn how a certain system is configured, is to ask the person who designed or implemented it. For an external IT analyst or employee at a managed security service provider, such information is almost impossible to get.

Experience has also shown that the anomaly sensors are best given to departments maintaining the OT systems. Most alerts they generate are operational. Hence, operational knowledge is required to analyze

them. Moreover, the alerts often help the operational departments in their normal tasks of keeping the systems running. So, often the departments are happy to have them.

Misuse detection sensors will typically only see attacks developed for IT systems. So, they may as well be operated by analysts with IT knowledge.

Capabilities that are focused on the incident response process, such as coordinating incident response actions and gathering logs, are also best given to IT security people. Such processes are commonplace in IT security operations. They can be moved to the OT domain without major changes.

The most difficult capabilities are those that require an in-depth analysis of attacks and attackers. Incidents that are currently found are usually spillovers from IT systems. So, they can be analyzed well by IT specialists. But it is likely that IT specialists will miss advanced threats, because they do not understand the OT-specific methods they use, and do not know what their targets may be. Specialists that combine security and OT knowledge will have to be hired or trained to perform such activities.

4 Roadmap

Figure 2 gives a roadmap for developing the capabilities, based on how advanced they are.

	Impact aware	Vulnerability aware	Threat agent-aware	Threat actor aware
OT engineering	System recovery Host vulnerability identification (basic) Host vulnerability mitigation (basic) Architecture vulnerability identification (basic) Architecture vulnerability mitigation (basic)	Host vulnerability identification (advanced) Host vulnerability mitigation (advanced) Architecture vulnerability identification (advanced) Architecture vulnerability mitigation (advanced) Anomaly sensor tuning and maintenance (basic) Triage of anomaly detection events (basic)	Anomaly sensor tuning and maintenance (advanced) Triage of anomaly detection events (advanced)	
Security operations team		Contacts with information sharing groups (basic) Triage of misuse detection events (basic) Infrastructure maintenance (basic) Log collection (basic)	Triage of misuse detection events (advanced) Forensics Infrastructure maintenance (advanced) Processing of Structured Online Intelligence Processing of Unstructured Online Intelligence Log collection (advanced)	Threat hunting Threat actor analysis Contacts with information sharing groups (advanced)

Figure 2: Roadmap for developing smart grid security operations capabilities.

In the roadmap, we distinguish the following development stages:

1. **Impact aware:** the organization is aware of its critical assets, the assets whose compromise will have the largest impact on the organization. Security operations are organized to protect these critical assets, and restore them if they would be compromised.
2. **Vulnerability aware:** the organization is aware of where it currently most vulnerable to attacks. This information is combined with the information on impact to identify the high-risk assets. Security operations are organized to protect these assets, and restore them if they would be compromised.
3. **Threat agent aware:** the organization is aware of the local actions of attackers, such as vulnerabilities being exploited, or payloads being used. They react to the attacker actions, giving the priority based on the risk estimates based on impact and vulnerabilities.
4. **Threat actor aware:** the organization is aware of who is behind the actions of the threat agents, and can piece together how the local actions fit together into larger campaigns. They can analyse threat actor methods and motivation, and use this to actively counter the threat actors that can do most damage to the organization.

Of course, organizations will not always fit neatly into one stage of this model. They may be between stages, or acquire capabilities from later stages earlier.