SEVENTH FRAMEWORK PROGRAMME

Theme SEC-2011.2.5-1 (Cyber attacks against critical infrastructures)

D8.6 Privacy Impact Assessment

Contract No. FP7-SEC-285477-CRISALIS

Workpackage	WP8 - Dissemination
Editor	Rens van der Heijden, Frank Kargl
Version	1.0
Date of delivery	M40
Actual Date of Delivery	M40
Dissemination level	Public
Responsible	UULM
Data included from	UT,ALL

The research leading to these results has received funding from the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement n°285477.

SEVENTH FRAMEWORK PROGRAMME
Theme SEC-2011.2.5-1 (Cyber attacks against critical infrastructures)

The CRISALIS Consortium consists of:

Siemens AG	Project coordinator	Germany
Alliander		Netherlands
Chalmers University		Sweden
ENEL Ingegneria e Innovazione		Italy
EURECOM		France
Security Matters BV		Netherlands
Universiteit Twente		Netherlands
Ulm University		Germany

Contact information:
Michael Munzert
Siemens AG
Otto-Hahn-Ring 6
81739 Munich
Germany

e-mail: michael-munzert@siemens.com

Contents

Abstract

This deliverable describes an assessment of the impact on privacy of security mechanisms developed in CRISALIS and includes recommendation how to deploy security mechanisms in a privacy protecting way. During earlier reviews of CRISALIS, we identified that some of the security mechanisms envisioned and developed within CRISALIS may actually infringe privacy of end-users, especially in the AMI scenario. This is why CRISALIS decided to conduct further analysis and contacted the PRIPARE project for assistance.

The majority of this deliverable builds on a joint workshop between CRISALIS and the PRIPARE FP7 CSA project organized at Alliander's offices in Arnhem, NL, on 2015-05-18. PRIPARE is a Coordination and Support Action within FP7 that focuses on privacy by design and privacy engineering that provided the necessary privacy expertise. The results of the joint discussions were used as the basis for the privacy recommendations in this deliverable.

In this deliverable, we first provide an overview over the AMI scenario and introduce the workshop's goals and agenda. In preparation of this workshop, we have identified two of the mechanisms developed within CRISALIS to be used for further investigation. One is the AMI Intrusion Detection System METIS developed at Chalmers, the other one the forensic toolkit FERRET developed by Siemens. Both seem promising candidates for such an analysis, as they are representative for typical security tools.

Chapter 3 provides the results of our privacy analysis followed by the derived recommendations found in Chapter 4. In summary, we have identified how to design privacy-friendly architectures for such security mechanisms that rely on general principles like separation of concerns and data minimization that – if followed – allows to operate such security mechanisms in a *need-to-know* fashion so that they can access only the absolutely necessary set of data items without granting excessive data access to the security operators.

These conclusions are summed up in a final chapter.

1 Introduction

This deliverable documents a privacy assessment of two security mechanisms designed within CRISALIS. This analysis was conducted with support of the PRIPARE EU FP7 CSA project. Based on this privacy assessment, a number of recommendations could be derived. The discussions and the derived recommendations are also described in this deliverable. PRIPARE used this joint effort to evaluate its privacy engineering methodology and has documented the gained insights and proposals for enhancements in their deliverable D3.2.

1.1 PRIPARE

The PRIPARE project is an FP7 coordination and support action, which describes its mission as follows:

> The mission of PRIPARE is twofold: facilitate the application of a privacy and security-by-design methodology that will contribute to the advent of unhindered usage of Internet against disruptions, censorship and surveillance, support its practice by the ICT research community to prepare for industry practice; foster risk management culture through educational material targeted to a diversity of stakeholders. To this end PRIPARE will:
>
> - specify a privacy and security-by-design software and systems engineering methodology, using the combined expertise of the research community and taking into account multiple viewpoints (advocacy, legal, engineering, business),
> - prepare best practices material (guidelines, instructions, patterns, success stories) for the development and implementation of products and services of ICT-based systems and use-cases in the area of cloud computing, mobile services and the management of cyber incidents,
> - support FP7 and Horizon 2020 research projects through training workshops and practical support in applying PRIPARE best practices in their environment,

- provide educational material on approaches for risk management of privacy and create awareness on the need for risk management culture among users. Material consistent with PRIPARE methodology will be structured in a modular way in order to fit to different target audiences (policy makers, users, ICT students and professional),
- identify gaps and provide recommendations on privacy and security-by-design practices, support of unhindered usage of Internet and on the creation of a risk management culture. A research agenda will be proposed.

PRIPARE consists of the following project partners:

- Trialog
- Atos
- Trilateral Research & Consulting
- INRIA
- The American University of Paris
- Gradient
- Universidad Politcnica de Madrid
- Ulm University
- Fraunhofer SIT
- Waterford Institute of Technology
- Katholieke Universiteit Leuven

1.2 Workshop

This section describes the preparation for the workshop that took place in Arnhem on 18th of May 2015 and contributed significant results to this deliverable.

1.2.1 Goals

The workshop organized by Alliander and Ulm University had several main purposes:

1. to provide an extensive overview of current and future privacy concerns in security mechanisms for advanced metering infrastructure (AMI),
2. to provide recommendations on how to approach these within the current legal framework, and
3. to demonstrate and evaluate the effectiveness of the PRIPARE methodology.

As one of the very first projects being assessed with this methodology, we received extensive assistance from the PRIPARE project partners, who guided and supervised the process. Because CRISALIS is one of the first projects to test this methodology outside the PRIPARE project, the workshop also resulted in extensive feedback on the methodology to PRIPARE; this can be found in PRIPARE Deliverable D3.2[1].

1.2.2 Workshop Agenda

Apart from the organizers (UUlm and Alliander), the workshop featured people from Chalmers, Siemens and Trialog.

Time	**Moderator**	**Topics**
8.30 - 9.00	UUlm	Welcome & introduction of participants
9.00 - 9.40	Alliander	Forensic tools and security mechanisms for AMI
9.40 - 10.00	Trialog	PRIPARE Methodology
10.00 - 10.15		break
10.15 - 12.30	Trialog	Privacy Analysis
12.30 - 13.30		Lunch
13.30 - 14.30	Trialog	Privacy Analysis
14.30 - 15.00		break & room switch
15.00 - 16.00	UUlm	Recommendations
16.00 - 16.30	UUlm	Summary & next steps

1.3 Privacy in AMI

Privacy protection is one of the most important challenges for the practical implementation and deployment of AMIs. There have been numerous discussions or privacy concerns

[1] Available under `http://pripareproject.eu/research/` as soon as the deliverable is approved and published.

over the past decade, and especially some initial proposals for AMI deployments were not implemented due to these concerns. Industry has since improved its collaboration with consumer and governmental organizations, which has led to significant improvements in this area. Notably the Article 29 Data Protection Working Party has published opinions for data protection in AMIs [3, 2, 1]. In the Netherlands, industry has additionally devised a code of conduct [10], which specifies voluntary commitments with respect to privacy protection. Within the CRISALIS project, privacy protection in AMIs has already been discussed in Deliverable D6.3.

Research has also extensively worked on privacy and security in AMI [8, 14], and several practical attacks on privacy in AMI have been demonstrated and widely published in the media [15, 16]. Due to this extensive media attention, AMI is one of relatively few areas where the public interest is very focused on privacy issues. This is part of the reason that the Article 29 Data Protection Working Party has published comparatively many articles in this area. What is missing was an investigation on the effects of deployment of security mechanisms like the ones designed in CRISALIS on privacy of AMI end-users.

In this deliverable we will therefore concentrate on the privacy issues produced by the security mechanisms designed in CRISALIS, something that has not yet received as much attention as effects of data collection for grid operation on privacy. Privacy protection is only effective when it is taken into account by all components of the system; the security infrastructure is no exception to this.

2 System Model

This chapter describes background knowledge and assumptions used for the remainder of this deliverable. In particular, we describe some details of the AMI system architecture and power distribution networks as deployed by distribution system operators (DSOs) like Alliander and the security mechanisms from CRISALIS that were analyzed. Parts of our discussion are specific to the Netherlands; however, they make a good use case, as they have one of the strictest legal frameworks regarding this topic.

2.1 AMI Overview

Before and during the workshop, Alliander provided participating PRIPARE partners with an overview of a typical AMI architecture, as well as details of data protection aspects both general and specific to the Dutch legislative situation.

2.1.1 Technical Architecture

The basic system architecture is the same as the one described in deliverable D6.3. For the purpose of this deliverable, we use a slightly simplified architecture as illustrated in Figure 2.1. This architecture consists of the following components: smart meters (SMs), the telecommunication provider that connects the SMs to the DSO, one or more head-end (HE) that can communicate with a database, and an interface to the EP. In practice the database will be included as part of a larger interface between the EP and the DSO, and this database will cache results produced by requests from the energy provider (EP) and data collected from the HEs. It is important to remark that there are two possible query architectures; pull and push. Pull refers to an architecture where the HE only asks its SMs for data like meter readings when a specific query from the EP is present, and may cache the result to allow more efficient queries. On the other hand, a push architecture is the approach where the SMs pro-actively share their meter readings and other data with the HE, such that they can be stored at the database and readily available for queries that come from the EP. Unlike many architectures discussed in literature, the system that is deployed by Alliander and others uses a pull architecture for most applications that involve access to meter readings. This means that data may

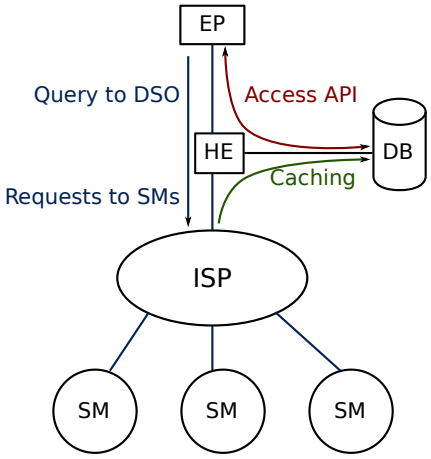

Figure 2.1: Overview of AMI architecture. This is based on deliverable D6.3 and information from Alliander. The arrows represent data flows, not direct communication: the queries are translated by the HE, not directly sent from the EP to the SMs; similarly, their responses are sent to the HE, which stores them in the DB.

not always be present and current, which could be an obstacle for security systems like intrusion detection systems, if there is no alternative way to access this data.

To limit communication costs, many DSOs design their networks in such a way that read-outs of the meters are limited to a specific time span, rather than occurring with a high frequency throughout the day. For example, Alliander limits their queries to only occur at night time: these limitations ensure that communication costs remain reasonable. This is somewhat beneficial for privacy, as the amount of data that can be transferred is also limited. However, it is theoretically possible to increase the bandwidth and time span in which data is moved from SM to HE, especially when the SM will be integrated into the consumer's home network or when its Internet connectivity is upgraded. Thus, we cannot rely on this to limit the privacy sensitive data that is made available to the DSO or EP. Indeed the literature typically considers relatively high read-out frequencies to demonstrate the privacy risk associated with an SM. However, it is important to keep in mind that the bandwidth available for data transmission is typically limited, taking privacy as another motivation to limit available bandwidth and exchanged data.

2.1.2 Legal Framework

Beyond the technical architecture, there are also several aspects in the legal framework that are relevant for the remainder of this deliverable. The most important aspect is that the Netherlands, like some other European countries, features a strict separation between the production and the delivery of energy.

In particular, the EP is not allowed to control the grid (access to homes), and similarly the DSO is not allowed to provide users with energy. This legal requirement exists as part of an effort to increase competition for energy production. As a side effect, the DSO may not request energy consumption data from the users directly but only on behalf of the EP.

For the same reason, the authority for consumers & markets (ACM) plays an important role in the legal framework where CRISALIS mechanisms may be deployed: this authority exists to ensure fair competition, and in particular also has some control over what data the DSO may collect. This is particularly important because the DSO is transparent to the user; a consumer has a contract with the EP, but does not have a lot of control to choose their DSO.

For intrusion detection systems (IDSs) deployed by the DSO this means that there is an additional obstacle to gather data and detect attacks on the network, as the DSO in particular cannot create arbitrary requests that are sent to the SM, because, legally speaking, the EP as the data controller is responsible for compliance to the privacy

legislation. The DSO is only a data processor and can thus only collect additional personal data based on very strict requirements, such as a specific suspicion of fraud.

Another important aspect is that due to European data protection regulation, any entity requires legal basis or user consent and a specific purpose before personal data may be collected. For grid management, this issue has been discussed extensively before the introduction of the large-scale roll-out of SMs in the Netherlands: the legal basis comes from a specific law that allows the DSO to collect data for this purpose. In particular, the DSO is required to take appropriate measures to maintain the grid functioning, which provides the necessary legal basis for necessary data collection, bound to that specific purpose. For our discussions we have assumed that this also includes necessary data collection to ensure security from attacks including mechanisms such as intrusion detection and forensics. However, it was noted by the data protection officer (DPO) that the legal status of intrusion detection tools in general is a gray area.

In general, we assume that collection of data by security mechanisms like the ones developed in CRISALIS is subject to data protection if personally identifying information (PII) is concerned, that is, any data that is directly linked to persons or their behavior or can be linked to persons with any realistic effort. For the AMI use case, this is especially the case when a security mechanism can access meter readings from homes.

Another important aspect is the fact that consumers can opt-out of the SM network in several ways. One way is to not install an SM in the first place, thus also avoiding the privacy risks associated with them, but also removing the advantages. In some other countries, this is not an option, because the smart meter is required. The more common way, which also occurs when a customer does not opt-in to higher frequency read-outs, is to install the SM but limit the read-out frequency to 6 read-outs per year (required for billing purposes by law). This will be a fairly common case, at least in the near future, and this also relates to the issue of bandwidth discussed in the previous subsection.

In addition to the legal framework, there is a code of conduct by Netbeheer Nederland, the organization representing DSOs including Alliander, that specifies their understanding of the current legislation. This interpretation is a relatively conservative one that is in favor of data protection. It specifies most importantly that all data coming from the meter is considered private, unless there are specific arguments that show this is not the case. Other important aspects of the code of conduct include the specification of standard privacy audits to ensure compliance of the grid maintenance process with privacy legislation in the form of a yearly audit, as well as a specification of the rights given to the DPO and safeguards that ensure his or her independence within the organization. For a practical application of mechanisms developed in CRISALIS, the implications thereof are that it is essential to include the DPO in a deployment of such systems, and that they should comply with the specified requirements.

With regards to detection of fraud (as opposed to intrusion detection in general), there are some specific mandates to allow an EP or DSO to access additional personal data when a specific suspicion exists. In the specific legal framework in the Netherlands, such detection may likely occur at the EP rather than at the DSO, because they collect the data required for billing. In any case, the DSO could be legally required to provide access to additional meter information in such exceptional cases (e.g., through a court order). Here it is important to remark that a *specific* suspicion has to exist, such that only data relating to that specific suspect consumer is analyzed. In particular, this means that a suspicion based on a set of aggregate measurements of a large group of consumers[1] is *not* sufficient reason to analyze data from each of these consumers individually: initial information on which specific consumer is suspect is required for triggering investigations.

To conclude the discussion of the legal framework, we remark that the legislation is somewhat exceptional in the Netherlands, but the general approach to privacy is similar in other European countries. In the remainder of the deliverable, we will point back to these specific exceptional cases wherever we discuss a recommendation that is specific to Dutch legislation.

2.2 Selected mechanisms

Through initial discussions we identified two specific security mechanisms developed within CRISALIS that we consider representative and that we focus the privacy analysis on. As discussed in Chapter 1, we decided to focus our discussion on the AMI use case as we foresee significantly less issues with personal data in industrial control system (ICS).

To ensure the broad applicability of our work in this deliverable, we decided to choose two widely different approaches to security mechanisms and thus ensure we can cover a good amount of terrain. This should allow us to give general recommendations despite only selecting two mechanisms; each of these is representative of a specific area. We selected METIS as an example of an intrusion detection system that can both assess meta-data and may also require to inspect packet content up to the application layer. The second mechanism is a forensic framework called FERRET as forensic investigations typically require access to a broad range of data on end-user devices.

[1] In current systems, this is not implemented, but a lot of research points to aggregation as a potential solution to privacy issues.

2.2.1 METIS

METIS is an IDS designed to detect attacks originating from a SM, for example due to compromised firmware or network attacks originating from a consumers' network. It does this by analyzing data transmitted through the network using one or more *interaction modelers* to detect anomalies in the network. These anomalies are suspicious messages, which are forwarded to the *pattern matcher* that triggers an alert if these anomalies are malicious. The interaction modeler is a Bayesian network, classifying messages based on the respective conditional probabilities; in this particular case, the probability that a specific message or message type is sent by a particular sender. The pattern matcher raises detection alerts when a particular amount of suspicious messages are received in a particular time frame. This is necessary, because the false positive rate of an individual interaction modeler is too high. The interaction modeler can be split into three individual steps: data preparation, a Bayesian Network learner and a probabilistic filter that forwards a given packet to the pattern matcher with the inverse of the probability computed for this packet by the Bayesian Network. Thus, packets that are very rare are often forwarded to the matcher, while frequent messages are very rarely forwarded. For a more detailed discussion of METIS, please refer to [6] or CRISALIS Deliverable 6.4.

For the purpose of our workshop, we considered several generalizations of METIS in our scope, in order to provide a representative discussion of intrusion detection approaches that can be used to detect a variety of attacks on AMI. One important generalization is that we also considered more generic IDSs, rather than just the stream processing approach that METIS takes. For the purpose of privacy analysis, this difference is not that significant; the important part is that PII is processed by the IDS. Therefore, we also consider systems with more complex data gathering that require full access to system data. In particular, the IDS may have access to more data from SMs, for example extracted directly by an agent running on one of the monitored hosts (e.g., one of the HEs).

This relates to the second important generalization we considered: extending METIS to include more data to be analyzed. Particularly interesting for privacy in the AMI scenario is the use of meter readings produced by the smart meter because these can potentially be used to infer information about users [9, 12]. While METIS does not do this, extending it with algorithms that do process such data is straight-forward when the traffic can be analyzed on the application layer. Some proposals for this are also discussed in CRISALIS Deliverable 6.4. One good reason to do this is that the meter readings contained within this layer can also be used to detect insider attacks, such as energy fraud.

2.2.2 FERRET

FERRET is a framework to do live and partially automated forensics, originally designed for ICSs. The main purpose of the framework is to enable an efficient analysis of compromised machines in such a network. The framework can create so-called forensic agents that perform analysis on a running compromised host, meaning that it is no longer necessary to go through complex procedures such as removing the hard drives and creating disk images. Another important task performed by these forensic agents is the aggregation and compression of data, such that bandwidth can be saved; in practical situations, only a small fraction of the disk contains relevant information, and reducing the amount of information to be transmitted saves a lot of time. In AMIs, performing forensics on a compromised SM or HE can work in a similar way; thus, we consider FERRET as an example of forensics frameworks.

The main assumption is thus that FERRET is not only applicable to ICS, but also to AMI. This is reasonable, because the need for forensic investigations also exists in this area; in addition, FERRET is reasonably generic in and of itself, simplifying an application to other areas, including AMIs.

Another important extension we discussed considers the agent, which is a software component in forensics tools that makes data available to the forensic analyst. This agent can be customized in FERRET, which inspired a discussion of policy-controlled agents that are customized for maximum privacy (and minimized data). This is already part of FERRET due to the fact that bandwidth needs to be saved; adding privacy as an additional filter would align with the design goals of FERRET. From this we also discussed the option of a type of permanent or semi-permanent agent that can perform a type of access control; the forensic analyst can send requests signed by the DPO to this agent. The agent would then automatically grant access to properly authorized requests, and discard all other requests. This removes the need to compile and deploy new code all the time.

3 Privacy Analysis

This chapter describes the privacy analysis of the two discussed use cases.

As per the PRIPARE methodology, the initial step of the privacy analysis after defining the system model was to determine the relevant feared events. As part of our discussion, we quickly determined there is exactly one relevant feared event that is specific for the CRISALIS security mechanisms:

> The leakage of different PII, in particular meter readings and behavioral information derived from these readings, through the security mechanisms or their alerts.

The remainder of this chapter describes the subsequent discussion following up on this feared event. The discussion uses the Hoepman privacy design strategies [7] as a guideline for the different measures that can be taken, and describes the background from which the recommendations were derived.

3.1 METIS and Intrusion Detection Systems

The data collection steps in the architecture that are considered for intrusion detection are the transmission of data between SMs and HEs, as well as the database contained within the DSO. Although intermediate aggregation or control stations are not currently used in the Alliander architecture, if future developments lead to such additional stations, they would also be in scope. Data is collected at these different points in the form of network traffic; it is conceivable that the IDS is allowed access to the content of such traffic. In fact, to detect attacks like energy fraud, this is essential. However, this content analysis would be critical from a privacy perspective; it could lead to the feared event with significant probability. From a privacy perspective, the fact that data isn't stored on any permanent storage is not relevant for processing; the same legal requirements apply[1]. Therefore, it is necessary to employ privacy protection measures when introducing such functionality. The communication between the SM and its corresponding HE is a securely

[1] Note that this is *only* regarding the requirements for processing: when permanent storage is used, it adds other requirements regarding how the data is stored, and how long.

encrypted connection, which protects the data from the telecommunication provider; at the HE the data is stored in a database.

Based on the assumed legal framework, even collection of encrypted network traffic legally counts as data collection, because it can potentially be deanonymized by decryption. Therefore, it is important that we provide specific arguments that this encrypted communication is irreversibly pseudonymized or anonymized. This is a difficult argument to make, especially considering the DSO possesses the key material to extract the meter readings and store them in the database. One possible solution for this is to apply the Hoepman strategy *separate* and enforce an organizational separation between a security operator and an operational department. The security operator is responsible for the IDS and has limited access to the data (namely only the encrypted traffic), while the operational department has full access to the information, bound to the purpose of grid maintenance. During our workshop discussion this appeared as one of the reasonable solutions from a practical perspective, which is why it is one of the first recommendations that was agreed upon. The security operator may even be a different legal entity like an external security company.

At some point, the aggregation of meter readings (as per the Hoepman strategy *aggregate*), was discussed as a potential solution to the issue of irreversible pseudonymization. However, it turns out that the application of this strategy here has multiple disadvantages: in many cases, the individual readings may be required in order to be able to perform detection in the first place. Apart from that, the aggregate (e.g., an average of readings) cannot necessarily be used to actually identify the malicious consumer: the legal requirements for such an invasive step are quite strong. In particular, the average of readings from several different consumers poses a big challenge; detecting that one or more of these consumers has likely cheated does not provide sufficient grounds to analyze the individual readings of all consumer readings that were aggregated here. Thus, aggregation of such values is not supporting individual fraud detection.

Another open question we discussed is whether a current IDS would have access to the meter readings produced at the SM (rather than those transmitted to the HE, which is potentially a small subset thereof). This approach would align with the *minimize* strategy, but requires that the reduced access still allows for sufficient functionality. As several experts pointed out, this access may be necessary to detect certain attacks. Therefore, the discussion focused on describing alternatives to the collection of this data in a central way, aligning more with the *hide* or *aggregate* strategies, depending on which approach is chosen.

One approach was to attempt the design of a privacy-preserving way of collecting this data, for example by employing homomorphic cryptographic primitives, or through some process of pseudonymization. The problem with this approach is that this does

not necessarily satisfy the legislative requirements as described in the legal framework. Because mechanisms like homomorphic cryptography are still mostly a research topic, the legal status is not entirely clear, but there is no specific exception from the rule that encryption is not a sufficient means for pseudonymization (because normal encryption would be reversible).

Another discussed approach was to leverage on the distributed nature of the detection system; when components of the IDS are distributed into the SM, the detection can occur in those meters without providing a direct access interface. This means that at least the process of detection will have a reduced impact on privacy, because most of the data that is analyzed never leaves the meter. However, this also has the obvious disadvantage that it is hard to guarantee correct operation of that component on the SM. In addition, when detection is done in the SM and an attack is detected, it is likely necessary to report these suspicious messages or activity to a more central location (consider the interaction modeler in METIS), which may have privacy implications. Nevertheless, this approach was the recommended way to proceed, as this also has obvious advantages with respect to performance: off-loading part of the detection into the SM means that a lot of aggregation can take place there, reducing network overhead.

At this point the discussion went towards a more detailed analysis of fraud: the question is whether this is actually in scope for the Dutch system, due to the market separation between DSOs and EPs. Legally the DSOs is not allowed to request data from the SM directly; there should always be a request from the EP that requires this information. That information can then be cached by the DSO, and possibly be used for intrusion or fraud detection. Fraud detection is especially interesting for the EP; although the DSOs have no financial interest in detecting fraud, we agreed to keep it within the scope of our discussions. It would even be conceivable that there is a security service provided to the EP by the DSO, which would also offer fraud detection services and rely on irreversibly pseudonymized data provided by the EP. However, the legality of such a system in the current legal framework is very unclear.

As a result of our discussions at the workshop, we came up with an architectural diagram that describes how to integrate intrusion detection into the architecture from Figure 2.1. An enhanced version of this diagram can be found in 3.1. The design foresees a pre-processor (PP), which filters incoming data from the SMs or the database to only forward data that is strictly necessary for specific IDS operations. It is required that the DPO approves the specific PP configuration. To emphasize the necessary separation between grid management and IDS, there is no direct connection to the IDS but all data needs to flow via the PP.

This also applies if the architecture adds interaction modelers (IMs) into the SMs. Based on the information received, the decision module (DM) decides when an alert

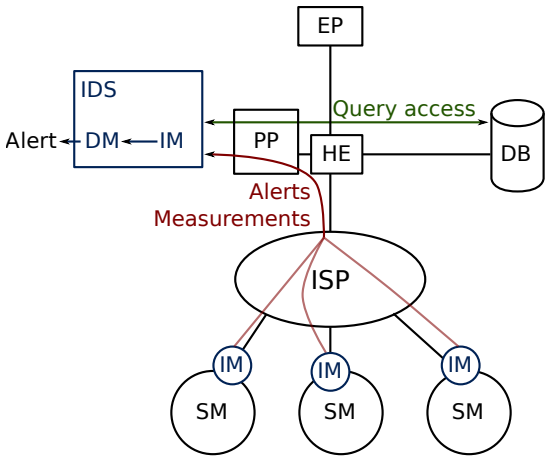

Figure 3.1: Recommended architecture for intrusion detection, based on Figure 2.1.

should be triggered.

The important observation is that we have an organizational separation between operators of the IDS and of the AMI and that the DPO of the AMI operator is in charge of approving all data flowing through the PP by ensuring data protection requirements being maintained.

3.2 FERRET and forensics

The discussion of FERRET during our workshop was somewhat more complex, because unlike METIS it was not designed for an area of use where privacy is a major concern. Rather, FERRET was designed for an area where typically having more data improves system performance, which is somewhat at odds with privacy at first glance. Nevertheless, one of the design goals of FERRET was to minimize the data collected by the forensic agent, because all this data needs to be extracted and transmitted over a bandwidth-limited network. This minimization aligns with the Hoepman *minimize* strategy and thus was one of the major points of discussion, despite being somewhat accidental. This also led to the first recommendation for FERRET: extend this data minimization module with a privacy-protecting component. Not only does this align with one of the FERRET design goals, it is also a relatively cheap measure that is potentially relatively simple to implement. A potential disadvantage is that more complex

attacks, such as data exfiltration by malicious system users through, e.g., e-mail cannot be detected this way. However, this is an issue that is mostly relevant when the monitored systems are actively used by users; this is not the case with current AMI components like SMs. Thus, at least for the short term, this is not a significant issue; when future work takes the SM in the direction of a full-blown home automation system that can also control devices in the house, this may become an issue, however.

In practice, forensic investigations often involve a lot of manual work, typically covering the analysis of log files and binaries that were detected on the compromised system. For privacy protection, it is important that the system log files on SMs do not, for example, contain detailed meter readings or other PII, such as a device history. Similarly, the analysis of binary files may involve the analysis of privacy-sensitive files; however, on a SM, privacy-sensitive files are typically not stored on the device, meaning that this is much less of a risk than for typical forensic analyses.

One advantage of forensics is that it is fairly costly to perform, meaning it will only happen for those devices that are very likely to be compromised. Although this is normally a disadvantage, it means that the necessary legal basis could already be given; a substantial suspicion against a specific customer should exist whenever sensitive data is processed. Of course, this only applies to a compromise that indicates the consumer was involved; for other cases, the consumer could be asked for permission, as her user experience is probably affected by the compromise. This would be different if a permanent forensic agent is deployed, but the current FERRET architecture requires the active deployment of an agent in order to connect the device to the framework.

Following up on the discussion of compression and aggregation in the forensic agent, the idea of using custom-built agents to improve privacy came up. Customized forensic agents normally allow the forensic analyst to collect data for a specific compromise, but they could also be designed to specifically retrieve data relevant only to this compromise, taking the minimization discussed above to the next level. However, building such customized agents may be challenging. Nevertheless, this led the discussion to continue with this idea and consider the inclusion of the DPO into the forensic process. The DPO could approve customized forensic agents for deployment thus ensuring that only necessary and – from a data protection point of view – acceptable data is gathered.

To take this idea to the next level, we discussed the possibility of permanent forensic agents installed onto the SM, which are only triggered when forensics should be performed. This approach led us to discuss the advantages of policy-based or obligation-based access control mechanism for the forensic agent. Policy-based access control is a type of access control that allows or denies access based on a policy provided by the user requesting access to specific information; similarly, obligation-based access control integrates the concept of obligations from privacy legislation into an access control mech-

anism that can enforce them. Both policy-based and obligation-based access control are mechanisms that have had a significant history within the research community [5, 11], as well as in industry [13, 4]. The additional advantage of this approach is that the DPO could approve specific policies, or specify the necessary obligations, which would significantly simplify management of the consumer's privacy protection. In addition, it would also formalize this construct more strongly, which is helpful when performing audits.

It would even be possible to integrate the consumer into the authentication process, at least when detecting compromise of the SM from other, external sources. For example, current practice when debugging specific issues with the meter, such as physical damage to the meter or software issues, are in part resolved by contacting the consumer and asking them to perform specific tasks or grant permission. This works particularly well when the issues are also interfering with the user interface, meaning that the consumer will notice the problems. As the average, honest consumer has an interest in a working meter, the DSO can to some extent rely on this as a way to get specific approval for the forensic analysis.

Another important measure that was pointed out is that both the forensic station (where the actual analysis is performed), the FERRET server and the forensic agent should require strong access control to prevent unauthorized users from accessing potentially privacy-sensitive data. This is particularly important because the forensic process always carries a risk of unintentionally compromising parts of the forensic station. Strict security policies and system isolation are tools that can be used to effectively recover from such compromises; strong access control can restrict the impact of such an attack by preventing access to most PII. The current FERRET design already foresees such access control measures.

To complete our discussion of forensics, we have also created a recommended architecture extension for a forensic framework, which is shown in Figure 3.2. Notice the similarities with 3.1: we designed the enhancements to be as similar as possible, because in practice both forensics and intrusion detection will be tasks for teams embedded within a security department or external security provider. In this graphic, there is no pre-processor, and the forensic agents (FAs) are included inside the SMs, rather than between the SM and the network. Also note that an FA can be deployed at the HE, as illustrated in the figure. The forensic framework contains a forensic station (FS), which in practice could be multiple stations if the compromise is large enough to require parallel investigation. After some time, the team will produce a forensic report – this is included to illustrate the contrast with an alert, which often takes a much shorter time to generate. Note again the organizational separation and system isolation, illustrated by the lack of connectivity between the forensic framework and the HE.

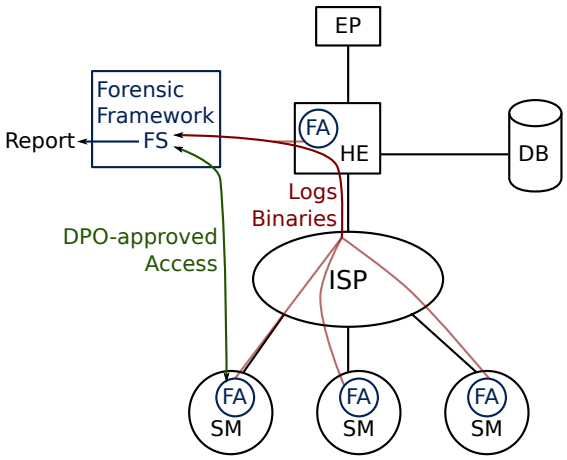

Figure 3.2: Recommended architecture for forensics, based on Figure 2.1.

3.3 Privacy for AMI

In addition to the concerns discussed above, we also identified some parts of AMI privacy that are also significant for the security mechanisms, but that are not specific to those mechanisms. We discuss these points separately here.

One of the most important concepts in privacy is consent: if the system operator can convince customers to voluntarily allow the processing of their data for security purposes, the DSO can use the data for this purpose in arbitrary detail. However, it is important to remark that this consent must be given voluntarily and explicitly: it cannot be the default. Asking for consent is especially helpful when protecting the grid against attacks from the outside, because in those cases the user may be willing to temporarily give up privacy in order to assist in the protection of the grid. Similarly, when attacks are occurring that are disruptive or damaging to the users' network infrastructure or SM, the user is often willing to assist in effective response against such incidents.

The second important aspect is compliance management, which is closely associated with privacy audits. The purpose of compliance management is to ensure that privacy controls are implemented correctly and that policies are actually executed in practice. The privacy audit is one of the most important tools to realize and formalize this process; similar to a security audit, a privacy audit is typically performed by an external organization and verifies that (1) policies are executed correctly and (2) the policies provide the desired level of privacy. In addition, it can be verified that this corresponds to the

law (or in the case of Dutch DSOs, compliance to the code of conduct (CoC)). Although privacy audits are part of compliance management, they are not sufficient to ensure compliance, as they typically take place on a yearly basis. Response to immediate security threats requires more flexibility, which is why there is an independent DPO that takes care of the day-to-day management of compliance. Organizations should consider that as the complexity of the security infrastructure increases, additional effort is required from the DPO, because the DPO has to treat grid management and security separately, because they are different purposes.

4 Recommendations

This chapter briefly reviews the recommendations that were identified by the involved partners.

4.1 General Recommendations

Although the discussions at the workshop focused on treating each mechanism individually, much of the discussion was generalizable in one direction or the other. The following lists the most important recommendations:

1. Separation of concern.
2. Exploit performance requirements to enhance privacy.
3. Decentralize where possible.
4. Ensure strong, flexible access control.
5. Use strong transport security protocols.

4.1.1 Separation of Concern

By separating between a system operator and a security provider, one can have clear interfaces between system and security mechanisms. These interfaces can be monitored and approved by DPOs to ensure that only the minimal and required set of data is available to security mechanisms.

4.1.2 Performance and Privacy

Performance is of primary concern for the practical implementation of these systems, especially because there is limited bandwidth available to transport information to other places. Wherever possible, minimization and aggregation can be employed, we can use privacy as an additional parameter to filter data transmitted to the security mechanism.

4.1.3 Decentralize

To further minimize and aggregate data, we recommend the decentralization of the security mechanisms. This avoids the very risky procedure of having to collect massive amounts of data from a variety of different sources in one place for analysis. Collecting this data is risky because it increases the attack surface, with the added disadvantage that much of the data will be filtered anyway, either for performance reasons or because it simply does not turn out to be useful.

4.1.4 Access Control

Both existing components and possible new components (such as the decentralized IDS components or forensic agents) should be subject to strong access control measures. This is true for humans accessing forensic servers but also to system components accessing other components for data access (like the forensic server accessing a forensic agent or a forensic agent accessing data on a smart meter). Although this should be standard security practice, there are also additional access control mechanisms available in the literature and in other systems that are designed with privacy in mind. In particular, they support the technical enforcement of concepts that are often used for privacy, such as obligations and policies[1]. Wherever PII is accessible, especially across organizational or geographic bounds, such access control mechanism should be evaluated and implemented as appropriate.

4.1.5 Transport Security

It is necessary to ensure that any communicating components have strong transport security mechanisms in place to prevent eavesdropping. This should be standard security practice, but as decentralization is one of the recommendations, and the data is typically transmitted over a third-party network that may or may not be connected to the Internet, it is important that transport security receives additional attention. Because of the long lifetime of an average SM, there should be predefined procedures and technical measures to ensure this can be maintained. Learning from current issues in ICSs, we also recommend these processes and measures be certified and executed regularly; this could be done as part of a privacy or security audit.

[1] More information about these can be found in Chapter 3.

4.2 Recommendations for Intrusion Detection

Specifically for intrusion detection, we consider the following recommendations to provide the most significant benefits:

1. Organizational separation of security and operations.
2. Preprocessing before data enters the IDS.

4.2.1 Organizational Separation

Companies should separate the security and the operational team organizationally. Here the security team is responsible for the secure operation of the system, for maintaining the intrusion detection system, and for intrusion response in case an attack is detected. The operational team, on the other hand, is responsible for maintaining the system itself: ensure the correct operation of the infrastructure, the SMs, and the access provided to the EP. Separating these tasks is an important step towards avoiding unnecessary access to data by members of the security team, which is especially important when the team only does monitoring. As soon as intrusions are detected based on anonymized data, intrusion response can be performed in cooperation with the DPO. The DPO can provide approval for the additional access that is often required for intrusion response. Because intrusion response can be a time-critical process, which requires some additional steps. One possible approach is to educate the DPO in typical intrusion response procedures; another could be to have standard procedures that are vetted by the DPO, or at least standardized processes to obtain permission from the DPO. That way the DPO will be able to respond more quickly and accurately, because much of the privacy implications are known in advance.

4.2.2 Always Pre-process Data

Most IDSs already contain a step that pre-processes data that enters the IDS: it is important to ensure this occurs at every place, with a module that removes privacy-sensitive data to the largest extent possible. The precise implementation of such a module may be different for each detection approach and organization; in some cases, it may make more sense to design a filter that removes irrelevant data, while in other cases a strict access control policy designed for privacy may make more sense. Nevertheless, there should always be a way to filter information, and this should be managed by the DPO, such that sufficient control can be exercised over this data. This recommendation is especially useful when implementing the decentralization recommendation from above;

it specifies how to deal with the individual components, as those detection components may still have privacy implications.

4.3 Recommendations for Forensics

As forensics is a relatively specific use case, which usually comes with a significant suspicion even in advance, we have some different recommendations in this regard. One potential challenge for any forensic system is to ensure sufficient data is available; therefore we also have recommendations with respect to (semi)permanent agents, even though this was not strictly in the scope of our initial discussions.

4.3.1 Protect the forensic station

During the forensic process, the information is typically collected in a central server or directly at a forensic station, where the analyst performs his analysis. As the analyst requires access to a lot of relevant files and logs, which could potentially contain PII, the station should be well-protected from attacks and require strong authentication of its users. In particular, advanced logging could be applied to ensure that each analyst only accesses resources related to his or her assigned projects, to avoid leakage of information to other analysts. It is also highly important to protect the infrastructure from exfiltration attacks, which may be caused by malware that is retrieved for analysis by the forensic framework. This is one of the reasons that a forensic framework and server is useful; it becomes a lot easier to isolate the forensic process from other security-critical infrastructure components.

4.3.2 Approved data requests

When deploying any kind of forensic agent, it should be ensured that this agent only accesses the data it really needs to, and transmits this information to a forensic station or server. In order to enforce this in a flexible way, there were various proposals to apply some method of data collection that relies on policies. Our recommendation is to work with such a system, creating a work flow that allows the DPO to approve forensic access to specific data. This could be realized either by specifying policies in some kind of software component within the SM by the DPO, or it could be implemented using requests that are (cryptographically) signed and thus approved by the DPO, which can be verified by the component under investigation. The biggest disadvantage of both of these approaches is that they could require extensive computational effort on the part of the SM. However, it would be an application of existing techniques, and it could be

implemented together with the policy-based access control discussed above to provide a very flexible forensics framework that can access a lot of data, while still preserving privacy as much as possible. In particular, such procedures can also be used to obtain sensitive data, if the DPO considers it necessary, or when another legal obligation exists, such as a court order.

5 Conclusion

Our effort is among the first to investigate privacy consequences of deployment of security mechanisms in AMI. Based on our discussions and the workshop results, we were able to provide a number of specific recommendations for strengthening privacy protection. The most important recommendations are (1) organizationally separating the security and operation teams, (2) decentralize the security infrastructure, similar to the grid management infrastructure, and (3) deploy strong, privacy-oriented access control mechanisms. In addition, the discussions provided directions for future research to focus on: we will briefly address these directions at the end of this chapter.

We acknowledge the support and help of the PRIPARE project in achieving these results by providing a methodological framework and privacy expertise during our effort.

5.1 Open issues

The recommendations posed by this deliverable support privacy protection, but some of them have open issues that need to be addressed before deployment. We briefly discuss some of these issues here.

First of all, an important risk that should be addressed is the potential for attackers to abuse privacy protection to hide their attacks. Our discussions highlighted that this is mitigated somewhat by the fact that a reasonable suspicion can be used as grounds for further (and possibly more invasive) investigation. However, it is likely that this in itself does not provide sufficient mitigation: the bar for reasonable suspicion is quite high. Therefore, research should focus on the design of privacy-preserving IDSs that only reveal the identity of an attacker when an attack is detected, and which has sufficiently low false positive rates to give a meaningful amount of privacy.

Another open issue is the increased risk that distribution of the security infrastructure may bring, if that infrastructure itself is not carefully protected. Although distributed detection can improve detection rates, because more security-relevant information is available, it also means that there are more components in the security infrastructure. More components typically means that the risk of compromise of individual components should not be neglected. Even when the attacker cannot use individual compromised components to attack the security infrastructure itself, it is important to consider that

some of the reports from these components may not be trustworthy. A practical example of this problem is a meter whose hardware is tampered with. Although there are solutions to this issue, tamper-resistant or tamper-evident hardware is not cheap, and if the attacker can disrupt the network connectivity of the meter, it cannot report the compromise back to the security infrastructure.

The final open issue we discuss here is how to perform threat intelligence. This is the process of gathering and sharing information about exploits and attacks that are discovered in the wild. In the current, state of the art security practices, threat intelligence is one of the most important tools that can be used to ensure an enterprise network can detect and recover from attacks quickly. By sharing signatures, process and network traces, and indicators of compromise, threats can be dealt with more effectively. However, sharing such information for AMI often means that PII, such as meter identifiers or undocumented meta-data in network packets, may be contained within this information. This makes threat intelligence difficult, and therefore an important tool for security management is lost. It may be possible to remove PII from this information, but we are not aware of current approaches that offer this functionality; it would need to be done manually, and even when the analyst has full knowledge of the semantics of the information, it may be difficult to avoid side-channels that compromise privacy. An effective means of performing threat intelligence for SM would be a good step towards improving their security.

5.2 Future work

Although most of the discussion of the AMI focuses on meter readings exclusively, the wider vision of smart grids and home automation should also be considered in future work. The design of a smart grid incorporates several new applications that may also create privacy concerns: controlling devices such as washing machines to run when power is cheapest, and providing power, for example from a consumers' solar installation, back into the infrastructure in a more efficient way. In both of these examples, the SM controls additional devices within the home, and thus will possess much more information about the consumer in order to meet his requirements. The process of returning power into the grid requires a higher degree of interactivity between the SM and the grid, which could go up to a dynamic marketplace for supply and demand, whose main purpose is to maintain grid stability when power is produced by consumers on a large scale.

Some envision that when the SM becomes sufficiently interactive, it will become a gateway not just to manage power, but also to manage various other components that are typically associated with home automation. In particular, these visions include

the regulation of sunlight, air-conditioning and heating in the house, as well as other devices that perform specific functions like vacuum cleaning robots, dishwashers and washing machines. All of these devices either influence energy usage or require a lot of energy, which makes the SM a gateway for virtually all devices to the grid. The security implications of these developments can be very interesting, and they will in large part depend on who maintains the SM and its network. If the consumer is responsible for the SM, developments may have parallels to the introduction of the home router, and the SM cannot be trusted by the grid; on the other hand, the privacy implications of an externally-controlled SM in this type of setting will be enormous, because of the tight control the SM requires over the consumers' devices.

Bibliography

[1] Article 29 Data Protection Working Party. Opinion 12/2011 on smart metering. Technical Report WP 183, Article 29 WP, April 2011.

[2] Article 29 Data Protection Working Party. Opinion 04/2013 on the data protection impact assessment template for smart grid and smart metering systems (dpia template) prepared by expert group 2 of the commission's smart grid task force. Technical Report WP 202, Article 29 WP, April 2013.

[3] Article 29 Data Protection Working Party. Opinion 07/2013 on the data protection impact assessment template for smart grid and smart metering systems (dpia template) prepared by expert group 2 of the commission's smart grid task force. Technical Report WP 209, Article 29 WP, December 2013.

[4] P. Ashley, S. Hada, G. Karjoth, C. Powers, and M. Schunter. Enterprise privacy authorization language (epal 1.2). Online, last accessed 29.05.2015, `http://www.w3.org/Submission/2003/SUBM-EPAL-20031110/`, 2003.

[5] J.-W. Byun and N. Li. Purpose based access control for privacy protection in relational database systems. *The VLDB Journal*, 17(4):603–619, 2008.

[6] V. Gulisano, M. Almgren, and M. Papatriantafilou. Metis: A two-tier intrusion detection system for advanced metering infrastructures. In *Proceedings of the 5th International Conference on Future Energy Systems*, e-Energy '14, pages 211–212, New York, NY, USA, 2014. ACM.

[7] J.-H. Hoepman. Privacy design strategies. In N. Cuppens-Boulahia, F. Cuppens, S. Jajodia, A. Abou El Kalam, and T. Sans, editors, *ICT Systems Security and Privacy Protection*, volume 428 of *IFIP Advances in Information and Communication Technology*, pages 446–459. Springer Berlin Heidelberg, 2014.

[8] J. Liu, Y. Xiao, S. Li, W. Liang, and C. L. P. Chen. Cyber security and privacy issues in smart grids. *Communications Surveys Tutorials, IEEE*, 14(4):981–997, Fourth 2012.

[9] P. McDaniel and S. McLaughlin. Security and privacy challenges in the smart grid. *Security Privacy, IEEE*, 7(3):75–77, May 2009.

[10] Netbeheer Nederland. Code of conduct for the processing of personal data by grid operators in the context of installation and management of smart meters with private consumers. Online, last accessed 29.05.2015, `http://www.netbeheernederland.nl/themas/hotspot/hotspot-documenten/?dossierid=11010056`, May 2012. English translation published 25-02-2013.

[11] Q. Ni, E. Bertino, J. Lobo, C. Brodie, C.-M. Karat, J. Karat, and A. Trombeta. Privacy-aware role-based access control. *ACM Trans. Inf. Syst. Secur.*, 13(3):24:1–24:31, July 2010.

[12] U. NIST. Guidelines for smart grid cybersecurity: Vol. 1 - smart grid cybersecurity strategy, architecture, and high-level requirements; vol. 2 - privacy and the smart grid; vol. 3 - supportive analyses and references. Technical Report 7628 Revision 1, NIST, September 2014.

[13] OASIS. extensible access control markup language (xacml) version 3.0. Online, last accessed 29.05.2015, `http://docs.oasis-open.org/xacml/3.0/xacml-3.0-core-spec-os-en.html`, January 2013.

[14] V. Tudor, M. Almgren, and M. Papatriantafilou. Analysis of the impact of data granularity on privacy for the smart grid. In *WPES '13 Proceedings of the 12th ACM workshop on Workshop on privacy in the electronic society*, pages 61–70. ACM, 2013.

[15] D. L. Ulrich Greveler, Benjamin Justus. Multimedia content identification through smart meter power usage profiles. In *Computers, Privacy & Data Protection (CPDP)*, 2012. Talk only, paper available here: `http://1lab.de/pub/ike2012.pdf`.

[16] M. Weiss, A. Helfenstein, F. Mattern, and T. Staake. Leveraging smart meter data to recognize home appliances. In *International Conference on Pervasive Computing and Communications*, pages 190–197. IEEE, March 2012.