

Functional Requirements Under Security PresSuRE

Stephan Faßbender, Maritta Heisel and Rene Meis

paluno - The Ruhr Institute for Software Technology – University of Duisburg-Essen, Germany
{firstname.lastname}@uni-due.de

Keywords: security analysis, problem frames, requirements elicitation

Abstract: Recently, there has been an increase of reported security incidents hitting large software systems. Such incidents can originate from different attackers exploiting vulnerabilities of different parts of a system. Hence, there is a need for enhancing security considerations in software development. It is crucial for requirements engineers to identify security threats early on, and to refine the threats into security requirements. In this paper, we introduce a methodology for Problem-based Security Requirements Elicitation (PresSuRE). PresSuRE is a method for identifying security needs during the requirements analysis of software systems using a problem frame model. Our method does not rely entirely on the requirements engineer to detect security needs, but provides a computer-aided security threat identification, and subsequently the elicitation of security requirements. The identification is based on the functional requirements for a system-to-be. We illustrate and validate our approach using a smart grid scenario provided by the industrial partners of the EU project NESSoS.

1 Introduction

Recently, there has been an increase of reported security incidents hitting large software systems. Beside, for example, reputation damage, loss on market value and share, and costs for legal infringement (Cavusoglu et al., 2004; Khansa et al., 2012), fixing the security defect causing the incident is costly. Fixing a defect when it is already fielded is reported to be up to eighty times more expensive than fixing the corresponding requirements defects early on (Willis, 1998; Boehm and Papaccio, 1988). Thus, security issues should be detected as early as possible for a system-to-be. Therefore, it is crucial for requirements engineers to identify security threats, and to refine the threats into security requirements. But eliciting good requirements is not an easy task (Fire-smith, 2003), even more with regard to security, as most requirements engineers are not security experts in the first place.

In this work, we propose a method called problem-based security requirements elicitation (PresSuRE), which guides a requirements engineer through the process of eliciting a set of security requirements in collaboration with the stakeholders of the system-to-be and security experts. PresSuRE has several benefits. It does not require the requirements engineer to have a security background. It does not require any preliminary security requirements and

security relevant information. It lowers the effort by providing tool support for semi-automated modeling and an automated security analysis. And PresSuRE is completely guided by a detailed process.

PresSuRE is based on the same idea of deriving information flows from functional requirements like the problem-based privacy analysis (ProPan) (Beckers et al., 2013a), but changes the analysis to be suitable for security. The analysis and elicitation is based on a complete set of functional requirements for a system-to-be. The method is accompanied with tool-support¹. PresSuRE is based on the problem frame notation introduced by Jackson (Jackson, 2001). We decided to use problem frames because they have a semi-formal structure, which allows a automated analysis. Furthermore, they are system-centric and embody a description of the environment of the system-to-be, which is important for a security analysis. And the overall system-to-be is decomposed using the requirements, which makes the usage of problem diagrams scalable even for large systems, because the complexity of single problem diagrams is independent of the system size and the number of problem diagrams scales linearly. Thus, they are suitable as input for a semi-automated analysis, as they have a predictable structure, underlying semantics, and support focusing on parts of the system-to-be.

¹<http://www.uml4pf.org/ext-pressure/installation.html>

We briefly describe the case study (Section 2) we use for the running example and the validation. The problem frame notation is explained in Section 3. Section 4 introduces the running example, which is used for the rest of the paper. The PresSuRE method is then explained in Section 5, and in Section 6 the method is validated. In Section 7 related work is discussed, and the final conclusion is drawn in Section 8.

2 Case Study

To illustrate the application of the PresSuRE method, we use the real-life case study of smart grids. As sources for real functional and quality requirements, we consider diverse documents such as “Application Case Study: Smart Grid” provided by the industrial partners of the EU project NESSoS², the “Protection Profile for the Gateway of a Smart Metering System” (Kreutzmann et al., 2011) provided by the German Federal Office for Information Security³, and “Requirements of AMI (Advanced Multi-metering Infrastructure)” (ope, 2009) provided by the EU project OPEN meter⁴.

To use energy in an optimal way, smart grids make it possible to couple the generation, distribution, storage, and consumption of energy. Smart grids use information and communication technology (ICT), which allows for financial, informational, and electrical transactions.

We define the terms specific to the smart grid domain and our use case in the following. The *smart meter gateway* represents the central communication unit in a *smart metering system*. It is responsible for collecting, processing, storing, and communicating *meter data*. The *meter data* refers to readings measured by smart meters regarding consumption or production of a certain commodity. A *smart meter* represents the device that measures the consumption or production of a certain commodity and sends it to the gateway. An *authorized external entity* can be a human or an IT unit that communicates with the gateway from outside the gateway boundaries through a *wide area network (WAN)*. The *WAN* provides the communication network that interconnects the gateway with the outside world. The *LMN (local metrological network)* provides the communication network between the meter and the gateway. The *HAN (home area network)* provides the communication network between the consumer and the gateway. The term *consumer* refers to end users of commodities (e.g., electricity).

²<http://www.nessos-project.eu/>

³www.bsi.bund.de

⁴<http://www.openmeter.com/>

3 Problem-Oriented Requirements Engineering

Jackson (Jackson, 2001) introduced the concept of *problem frames*, which is concerned with describing, analyzing, and structuring software development problems. A problem frame represents a class of software development problems. It is described by a *frame diagram*, which consists of domains, interfaces between them, and a requirement. Domains describe entities in the environment. Jackson distinguishes the domain types *biddable domains* that are usually people, *causal domains* that comply with some physical laws, and *lexical domains* that are data representations. Whenever we have influence on the design of a domain it is a *designed domains*. To describe the problem context, a *connection domain* between two other domains may be necessary. Connection domains establish a connection between other domains by means of technical devices. Examples are video cameras, sensors, or networks. Note, that one domain can have more than one type, for example a domain can be a connection and causal domain at the same time.

Interfaces connect domains, and they contain *shared phenomena*. Shared phenomena may be events, operation calls, messages, and the like. They are observable by at least two domains, but controlled by only one domain, as indicated by the abbreviation of that domain and “!”. For example, the shared phenomenon *MeterData* in Figure 1 is observable by the domains *SMGProvider* and *PersistentMeterData*, but controlled only by the domain *PersistentMeterData* (abbreviation *PMD*).

The objective is to construct a *machine* (i.e., software) that controls the behavior of the environment (in which it is integrated) in accordance with the requirements. Problem-oriented requirements analysis relies on a decomposition of the overall problem into sub-problems, which are represented by *problem diagrams*. Problem diagrams contain the requirements belonging to the sub-problem. When we state a requirement, we want to change something in the environment. Therefore, each requirement *constrains* at least one domain in the environment. A requirement may also *refer* to several domains in the environment of the machine.

The problem frames approach distinguishes therefore between the *requirements (R)*, the *domain knowledge (D)*, and the *specification (S)*. The requirements describe the desired system after the machine is built. The domain knowledge represents the relevant parts of the problem world. The specifications describe the behavior of the software in order to meet the require-




Figure 1: Problem Diagram RQ 5 : Provide Meter Data

ments.

We describe problem frames using UML class diagrams, extended by stereotypes, as proposed by Hatebur and Heisel (Hatebur and Heisel, 2010). Figure 1 shows a problem diagram in UML notation. The biddable domain (UML class with stereotype `«biddableDomain»`) *Consumer* controls the *request consumption data* command (Name of the UML association with the stereotype `«connection»` between the classes *Consumer* and *HAN*), which is observed by the causal connection domain *HAN* (UML class with stereotype `«causalDomain, connectionDomain»`). The *SMGProvider* controls the phenomenon *readData*, which is obtained from the lexical domain *PersistentMeterData* (UML class with stereotype `«lexicalDomain»`). Additionally, the *SMGProvider* provides the data. The *PersistentMeterData* controls the *meter data* it contains. The *HAN* forwards the data and commands it observes. The requirement RQ 5 (for a textual description see Section 4) constrains the *HAN* and refers to the *Consumer*, and the *PersistentMeterData*.

4 Running Example: Billing

We chose the use case *Meter Reading for Billing* given in the documents of NESSoS and the open meter project to exemplify our method. This use case is concerned with gathering, processing, and storing meter readings from smart meters for the billing process. Beside the billing use case there are 13 use cases described for the minimal features of a smart meter




Figure 2: Problem Diagram RQ 1: Receive Meter Data

gateway in total, which we all considered for our validation. The functional requirements for this use case are defined as follows:

(RQ 1) Receive meter data The smart meter gateway shall receive meter data from smart meters.

(RQ 2) Process meter data The smart meter gateway shall process meter data from smart meters.

(RQ 3) Store meter data The smart meter gateway shall store meter data from smart meters.

(RQ 4) Submit billing data The smart meter gateway shall submit processed meter data to authorized external entities.

(RQ 5) Provide consumption data to consumer The smart meter gateway shall provide meter data for consumers for the purpose of checking the consistency of bills.

The problem diagram for RQ 5 was already shown in Figure 1 and explained in Section 3. Figure 2 shows the problem diagram for RQ 1. The causal domain *smart meter* controls the *send data* command, which is forwarded by the *LMN* and finally observed by the machine domain *SMGReceiver*. The *SMGReceiver* controls the phenomenon *writeTemporaryData*, which stores the received information in the lexical domain *temporary meter data*. Additionally, the *SMGReceiver* can *request data* which is forwarded by the *LMN* to the *smart meter*. The causal connection domain *LMN* forwards the data and commands it observes. The requirement RQ 1 constrains the *temporary meter data* and refers to the *smart meter*. These two problem diagrams are sufficient to understand the rest of the paper, nevertheless we use all five functional requirements for our method.

Note that we will even simplify this example in the following. We will not elaborate on all security elements but restrict ourselves to one example for each element, for example one asset, to improve the comprehensibility for the reader and for reasons of space. Nevertheless, for the validation we elaborated the full case study in means of 27 requirements and all possible assets and attackers.

5 The PresSuRE Method

The PresSuRE method consists of four phases and nine steps, which we will explain in the following (Figure 3 gives an overview. Whenever we refer to the figure we use a different font).

5.1 Model functional requirements

We assume that the functionality of the system-to-be is described completely, coherently and unambigu-

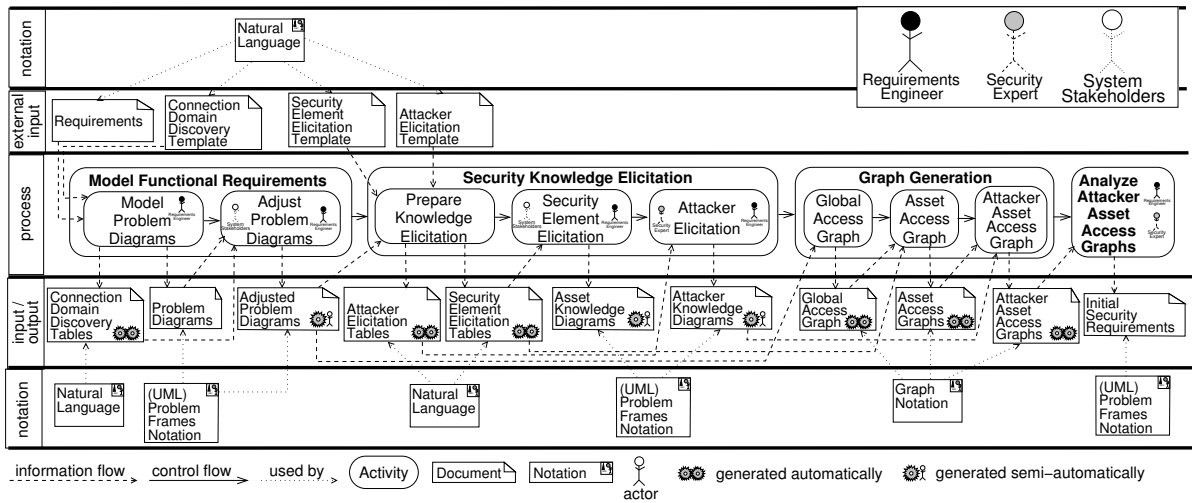


Figure 3: PresSuRE Method

ously. The functional requirements are a good starting point for a security analysis as the requirements engineer is used to deal with them, they are often already well defined, they already contain everything which has to be protected, and they also contain the entry points for possible attack vectors an adversary can use. Note that the results PresSuRE delivers are correct and sufficient with respect to functional requirements used. Hence, having incomplete or bad requirements has an impact on the results. But how to tackle this problem is out of scope of this problem. The modeling of the requirements is achieved by performing the following two steps.

Model Problem Diagrams In the first step of the PresSuRE method, the functional requirements have to be modeled using the problem frame notation. This can be done by the requirements engineer alone, based on a textual description of the functional requirements. The result is a set of problem diagrams as well as an automatically generated connection domain discovery table. *The functional requirements and corresponding problem diagrams are presented in Section 4.*

Adjust Problem Diagrams As setting up problem diagrams allows some degree of freedom, adjustments might be needed to prepare the problem diagrams. For the PresSuRE analysis, connection domains are specifically important. Many attacks use such connection domains as entry points, such as a wireless LAN for sniffing important data, a display which is visible to an attacker, or even a postman, who can be socially engineered. But as connection domains are not of central relevance for fulfilling

Table 1: Connection Domain Discovery Table

Domain 1	Domain 2	phenomena	Q1 ⁺	Name ⁺	Q2 [*]	Name ^o	Q3 ^o	Name ^o	PD [*]
Authorized-External-Entity	SMG-Submitter	SMG-SI{submit-Data} AE-EI{request-Billing-Data}	yes	WAN	no	-	no	-	RQ4-Submit-Meter-Data

⁺ Q1: Information transmitting domain involved? ^{*} Q2: Domain displaying information involved? ^o Q3: Domain storing information involved? ^{*} PD: Related problem diagrams

the functional requirements, they are often left out. Hence, one has to make sure that all connection domains are explicitly modeled.

For each connection between domains, the requirements engineer and the system stakeholders have to check if there is a connection domain in between. Typical connection domains, which should be modeled, are domains transporting information (such as network domains or a postman), domains which show information (such as (semi)- public displays), or domains which store information for exchange between domains (such as a network-attached storage).

While requirements engineers are able to analyze, understand, and explain the problem diagrams, and know which new connection domains might be introduced, they often lack the domain knowledge, which reveals already established connection domains. This knowledge is available through the stakeholders of the system-to-be. Hence, the requirements engineer guides through the discovery and models the results, while the system stakeholders provide the needed information.

The requirements engineer and the system stakeholders use a table containing the connected domains

pairwise, the phenomena in between and a standard questionnaire, which helps to elicit the missing connection domains. Table 1 shows such a table. The result of this step are adjusted problem diagrams, which are modeled by the requirements engineer using semi-automated wizards. *For our example, using the table and answering the questions, we see that our problem diagrams have to contain WAN, HAN and LMN as connection domains. This information is already reflected by the problem diagrams shown in this paper.*

5.2 Security Knowledge Elicitation

Before starting the security analysis, some security-specific knowledge has to be elicited. This information is crucial for the success of the analysis, as in most cases the functional requirements do not contain enough information for considering security thoroughly. The knowledge about assets in the system-to-be and attackers which might tamper with the system has to be made explicit. As this knowledge is not or only partially available for requirements engineers, they have to collaborate with the stakeholders of the system and security experts.

Prepare Knowledge Elicitation Even though, the functional requirements do not contain the information for security analysis, they do already contain some information, which is the starting point for eliciting the additional domain knowledge. We use security element elicitation tables, and attacker elicitation tables to elicit this information. The tables are automatically generated from the problem diagrams.

Identify assets, authorized entities and rights

The baseline questions for this step are “What has to be protected?” (*asset*), “Who is eligible to access the asset?” (*authorized entities*), and “Which actions are allowed for a stakeholder regarding an asset?” (*rights*). This is in line with common security and threat analysis methods (ISO/IEC, 2009a; Howard and Lipner, 2006; ISO/IEC, 2009b; Jürjens, 2005). We use the previously generated security element elicitation tables to elicit this information. The instance for persistent meter data is shown in Table 2. Such tables are completed by the stakeholders of the system-to-be using the following description, while the requirements engineer models the results.

Assets Identify those domains, which have to be protected. Every domain beside the machine is an asset candidate. Most likely one wants to protect a lexical domain representing information or a causal

Table 2: Security Element Elicitation Table for Persistent Meter Data

Asset-Candidate	is asset	Authorized Entity Candidate	is Authorized Entity	rights
Persistent-Meter-Data	<input checked="" type="checkbox"/>	WAN	<input type="checkbox"/>	<input type="checkbox"/> read <input type="checkbox"/> write
		Consumer	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/> read <input type="checkbox"/> write
		HAN	<input type="checkbox"/>	<input type="checkbox"/> read <input type="checkbox"/> write
		AuthorizedExternalEntity	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/> read <input type="checkbox"/> write
		SmartMeter	<input checked="" type="checkbox"/>	<input type="checkbox"/> read <input checked="" type="checkbox"/> write
		LMN	<input type="checkbox"/>	<input type="checkbox"/> read <input type="checkbox"/> write

: yes, : must have, : might have

domain. *For our example, we only select the persistent meter data as an asset, which contains information about the electricity consumption of the consumer. This information has to be protected for privacy reasons, as it, for example, allows to monitor the consumer. The full case study contains 13 further assets (see Section 6).*

Authorized Entities A authorized entity to an asset is every domain, which has an eligible interest in knowing the state / reading, or controlling / writing the asset. *Eligible entities of the meter data are the smart meters, which produce the meter data, the external entities, who need the consumption information for billing, and the consumer, who wants to check his/her electricity consumption.*

Rights Authorized entities have different rights to access the asset. In case of a lexical domain, the rights are to read or write the information in the domain. In case of the causal domain, the rights are to control or know the state of the causal domain. For each right and authorized entity, one has to state if the entity is allowed to have the right or if the entity must have the right. *The smart meters must have the right to write the information, while the consumer and external entities must have the right to read the information. The smart meters do not need to read the stored consumption data, and the external entities and the consumer are not allowed to modify the consumption data.*

The elicited information has to be added to the model. For this purpose, we use *domain knowledge diagrams*. In domain knowledge diagrams additional knowledge about domains and relations between domains can be modeled. To support modeling security-related domain knowledge we developed UML profiles. Figure 4 shows a domain knowledge diagram modeled using the profiles. It




Figure 4: Asset Knowledge for the Rights of the Consumer regarding the PersistentMeterData

depicts that there is domain knowledge (stereotype «domainKnowledge» at class *ConsumerPersistentMeterDataAccess*) about the persistent meter data and the consumer. The persistent meter data is an asset (stereotype «asset» at class *PersistentMeterData*), the consumer is a authorized entity of this asset (stereotype «authorizedEntity» at class *Consumer*), and the consumer is allowed to read («read») from the persistent meter data («refersTo»). The diagrams are generated in the background while the requirements engineer completes a wizard which is similar to the security element elicitation table. The result of the step are asset knowledge diagrams.

Attacker(s) Elicitation In this step, the requirements engineer and a security expert have to collaborate to define those *attackers* who might attack our system-to-be. While the requirements engineer has a deeper understanding of the system-to-be and its domain, the security expert adds his/her vital knowledge about attackers, attacker abilities, possible attack vectors, and so forth. Hence, it is not mandatory that the requirements engineer has a security background.

Beckers et al. (Beckers et al., 2013c) enumerate different *types of attackers*: *physical attacker*, *software attacker*, *network attacker*, and *social attacker*. Regarding their abilities, we have chosen the abilities as described by Dolev and Yao (Dolev and Yao, 1983): *read* (*read message / get state of domain*), *write* (*write message / change state of domain*), *interfere* (*intercept message / prevent the change of state*). Again, describing attackers with their basic abilities is in line with the state of the art (ISO/IEC, 2009a; Howard and Lipner, 2006; ISO/IEC, 2009b; Jürjens, 2005). For the purpose of eliciting the information about attackers, we use the generated attacker elicitation tables (see Table 3).

Attacker First, we have to reason for each domain and type of attacker about the question if this type of attacker might exist for the domain at hand. *For simplicity's sake, we assume for the running example that we only have to defend against network attackers. We distinguish between two network attackers: The inter-*

Table 3: Attacker Elicitation Table for HAN

Attack-Candidate	is attack-able	Attacker	Name	Abilities
HAN	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/> network attacker	internal net. attacker	<input checked="" type="checkbox"/> read <input checked="" type="checkbox"/> write <input checked="" type="checkbox"/> interfere
		<input type="checkbox"/> physical attacker		<input type="checkbox"/> read <input type="checkbox"/> write <input type="checkbox"/> interfere
		<input type="checkbox"/> social attacker		<input type="checkbox"/> read <input type="checkbox"/> write <input type="checkbox"/> interfere
		<input type="checkbox"/> software attacker		<input type="checkbox"/> read <input type="checkbox"/> write <input type="checkbox"/> interfere

: yes, : has the ability




Figure 5: Attacker Knowledge regarding the internal network attacker’s abilities regarding the HAN

nal network attacker, who has access to the HAN and LMN, where the smart meters reside, and the external network attacker, who can attack via the WAN. Note that for the full case study we found and modeled 7 attackers in total, including all kinds of attackers.

Abilities For each attacker and each domain the attacker has access to, we have to state which abilities the attacker has. Whenever there is no detailed information about the attackers and their abilities regarding a domain they have access to, one should assume the strongest attacker. This might lead to an overestimation of the threats afterward. But adding an unnecessary security requirement is not so much of an issue, while missing one is critical. *After an assessment of all attackers of our example and their abilities, we could not exclude any of the basic abilities. Hence, our attackers have all abilities regarding the domains they have access to.*

The elicited information has to be added to the model to be available for our analysis, too. Again, the modeling can be done semi-automatically using the wizards our tool provides. The result are attacker knowledge diagrams. Figure 5 shows the domain knowledge about the internal network attacker and his/her abilities regarding the HAN. It depicts that there is domain knowledge (stereotype «domainKnowledge» at class *InternalNetworkAttackerExists*) about the internal network attacker and the HAN. The internal network attacker is a network attacker (stereotype «networkAttacker» at class *InternalNetworkAttacker*). This attacker is able to read («read») from the HAN («refersTo»), and the attacker («refersTo») is able to write to and interfere with the HAN («write, interfere»).

5.3 Graph Generation

The automated part of the security analysis relies on graphs, which visualize information flows and access flows. The attacker asset access graphs, which contains the potential security threats towards the functional requirements, are generated stepwise. The steps and intermediate graphs are explained in the following.

Global Access Graph All graphs $(\mathcal{V}, \mathcal{E})$ that we use for our security analysis in the PresSuRE method are labeled and directed. The set of vertices is a subset of the domains occurring in the model, formally $\mathcal{V} \subseteq \text{Domain}$. An edge is annotated with a diagram and a type. The diagram can be a problem diagram or a domain knowledge diagram. The type can be required (*req*), implicit (*imp*) or attack (*att*) ($\text{Type} ::= \text{req}|\text{imp}|\text{att}$). The type indicates if the edge is *required* or *implicitly* given by the problem diagram or if it shows a possible *attack* relationship defined in a domain knowledge diagram. The edges point from one domain to another, formally $\mathcal{E} \subseteq \text{Domain} \times \text{Diagram} \times \text{Type} \times \text{Domain}$. For the rest of the paper we will regard such an edge as an access flow. In the following, we describe a graph $(\mathcal{V}, \mathcal{E})$ only by its edges \mathcal{E} .

For the analysis of the threats towards an asset we will use the *global access graph*. This graph contains the information about access flows between domains, and which problem diagrams are the source of these flows. For the flows, we distinguish between required flows as stated by the requirement and implicit ones which are modeled due to the given environment. To set up the global access graph we use the problem diagrams as an input. The predicates *constrains*, *refersTo* : $(\text{Domain} \times \text{Diagram})$ and *controls* : $\mathbb{P}(\text{Domain} \times \text{Domain} \times \text{Diagram})$ can be derived from the problem frame model and are used to generate the global access graph. We have $(d, p) \in \text{constrains}$ and $(d, p) \in \text{refersTo}$ iff a requirement or domain knowledge in diagram p constrains the domain d or refers to it, respectively. $(d_1, d_2, p) \in \text{controls}$ is true iff the domain d_1 controls an interface that d_2 observes in the diagram p .

Using these predicates, we create the global access graph \mathcal{G} , which is an overapproximation of the access flows occurring in the system-to-be. An edge $(d_1, p, \text{req}, d_2)$ is in \mathcal{G} iff the domains d_1 and d_2 are not equal, and the domain d_1 is referred to and the domain d_2 is constrained in p . For example, the problem diagram for RQ 1 (see Figure 2) contains the smart meter and the temporary storage. The smart meter is referred by RQ 1 and the temporary storage




Figure 6: Global Access Graph (also Asset Access Graph for Persistent Meter Data)

is constrained by RQ 1. Hence, we add a required access flow edge (solid arrow) between smart meter (node with name SmartMeter) and temporary meter data (node with name TemporaryMeterData) annotated with RQ 1 (see graph shown in Figure 6).

Additionally, an edge $(d_1, p, \text{imp}, d_2)$ is in \mathcal{G} iff the $(d_1, p, \text{req}, d_2)$ is not already in \mathcal{G} , the domains d_1 and d_2 are not equal, and d_2 observes an interface controlled by d_1 in p . Note that machines are treated as transitive forwarders in this case. This means that whenever a machine m observes an interface controlled by d_1 and d_2 observes an interface controlled by m , we assume that d_2 observes an interface of d_1 . For example, the domain LMN controls a phenomenon forwardMeterData which is observed by the machine (see Figure 2). The domain temporary meter data observes a phenomenon writeTemporaryData from the machine. Hence, an implicit access flow edge (dotted edge) is added between the LMN and the temporal meter data annotated with RQ 1 (see Figure 6). We define the graphs as follows.

$$\begin{aligned} \mathcal{G}_{req} = \{ & (d_1, p, t, d_2) : \\ & \text{Domain} \times \text{ProblemDiagram} \times \{\text{req}\} \\ & \times \text{Domain} \mid d_1 \neq d_2 \wedge \\ & (d_2, p) \in \text{constrains} \wedge (d_1, p) \in \text{refersTo} \} \\ \mathcal{G}_{imp} = \{ & (d_1, p, t, d_2) : \text{Domain} \times \\ & \text{ProblemDiagram} \times \{\text{imp}\} \times \text{Domain} \mid \\ & d_1 \neq d_2 \wedge (d_1, p, \text{req}, d_2) \notin \mathcal{G}_{req} \\ & \wedge ((d_1, d_2, p) \in \text{controls} \vee \exists m : \text{Machine} \bullet \\ & (d_1, m, p) \in \text{controls} \wedge (m, d_2, p) \in \text{controls}) \} \\ \mathcal{G} = & \mathcal{G}_{req} \cup \mathcal{G}_{imp} \end{aligned}$$

Because of the annotation of the edges we keep the information which problem diagram causes the access flow. Thus, our global access graph contains traceability links that are used in our further analysis. The semantics of an edge $(d_1, p, t, d_2) \in \mathcal{G}$ is that in problem diagram p there is possibly a required or implicit (depending on t) access flow from domain d_1 to domain d_2 .

Asset Access Graph As the global access graph can be huge for a complex system-to-be, we introduce an asset access graph which focuses the view on one asset only. It only contains access flows given by the requirements directly or indirectly concerning the asset. Thus, we get one asset access graph per asset. The asset access graph makes the information for the requirements engineer easier to comprehend. Hence, it improves the scalability of our method. An edge (d_1, p, t, d_2) is in \mathcal{G}_{asset} iff p is in \mathcal{P}_{access} . A problem diagram p is in \mathcal{P}_{access} iff

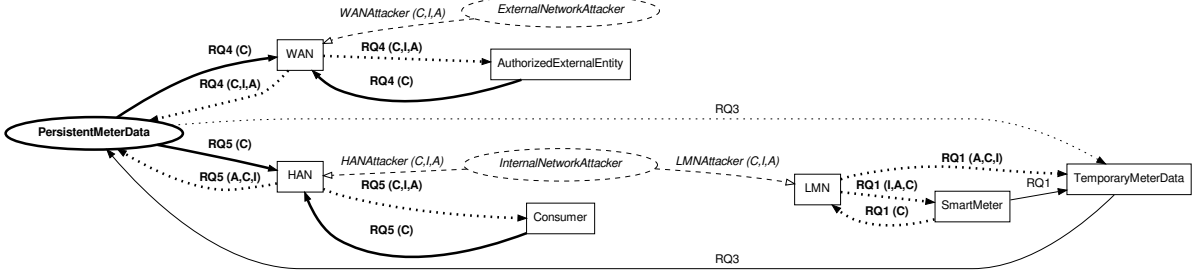


Figure 7: Attacker Asset Access Graph for Persistent Meter Data

there is an edge (d_1, p, t, d_2) which is required and d_1 and d_2 are both in \mathcal{D}_{access} . \mathcal{D}_{access} is a union of \mathcal{D}_{active} and $\mathcal{D}_{passive}$. A domain d_1 is in \mathcal{D}_{active} iff there is a required access flow which starts at d_1 and the target domain d_2 is already in \mathcal{D}_{active} . Initially, only the *asset* is in \mathcal{D}_{active} . Hence, \mathcal{D}_{active} contains all domains which have a required direct or indirect (via another domain) access flow towards the asset. A domain d_2 is in $\mathcal{D}_{passive}$ iff there is a required access flow which ends at d_2 and the source domain d_1 is already in $\mathcal{D}_{passive}$. Initially, only the *asset* is in $\mathcal{D}_{passive}$. Hence, $\mathcal{D}_{passive}$ contains all domains which are the target of a required direct or indirect (via another domain) access flow from the asset. Formally, we define \mathcal{D}_{active} , $\mathcal{D}_{passive}$, \mathcal{D}_{access} , \mathcal{P}_{access} , and \mathcal{G}_{asset} as follows.

$$\begin{aligned}
& \text{required } asset \in Domain \\
& asset \in \mathcal{D}_{active}, \quad asset \in \mathcal{D}_{passive}, \\
& (d_1, p, req, d_2) \in \mathcal{G} \wedge d_2 \in \mathcal{D}_{active} \Rightarrow d_1 \in \mathcal{D}_{active} \\
& (d_1, p, req, d_2) \in \mathcal{G} \wedge d_1 \in \mathcal{D}_{passive} \Rightarrow d_2 \in \mathcal{D}_{passive} \\
& \mathcal{D}_{access} = \mathcal{D}_{active} \cup \mathcal{D}_{passive} \\
& \mathcal{P}_{access} = \{p : ProblemDiagram \mid \\
& \quad \exists (d_1, p, req, d_2) : \mathcal{G} \bullet d_1, d_2 \in \mathcal{D}_{access}\} \\
& \mathcal{G}_{asset} = \{(d_1, p, t, d_2) : \mathcal{G} \mid p \in \mathcal{P}_{access}\}
\end{aligned}$$

The resulting asset access graph for the persistent meter data is shown in Figure 6, as for our small example the global and the asset access graph do not differ. For a complex scenario the asset access graph is significantly smaller than the global access graph. The asset access graph can be used to check if a stakeholder can gain more rights than he/she should. For reasons of space, we do not go into detail on this matter.

Attacker Asset Access Graph For each asset, we generate the attacker asset access graph, which visualizes the information and control flows from attackers to the asset and from the asset to the attackers. At this point, we focus on the basic information security goals confidentiality, integrity, and avail-

ability (short CIA), which are suggested by the Common Criteria (ISO/IEC, 2009a) and ISO 27000 family of standards (ISO/IEC, 2009b). The problematic access flows are annotated with the information which CIA property(ies) are threatened ($CIA ::= C|I|A|\epsilon$). First, the domains which are directly connected to attackers are identified. Note that for this purpose we use the information given in domain knowledge diagrams created during the step *Identify assets, authorized entities and rights* described in Section 5.2. From these diagrams, we can derive the predicates *read*, *write*, *interfere* : $\mathbb{P}(Domain \times Diagram)$. We have $(d, dk) \in read$, $(d, dk) \in write$, and $(d, dk) \in interfere$ iff domain knowledge in diagram dk has a read, write, or interfere dependency, respectively, to the domain d .

A domain can d be object to be attacked if it is in \mathcal{D}_{access} for the asset at hand. That is, an attacker can access or influence information on the asset through the domain d . We define the sets \mathcal{D}_w , \mathcal{D}_i , and \mathcal{D}_r as the sets of all domains for which an attacker has the ability to *write*, *interfere*, or *read* it, respectively. A domain d is in \mathcal{D}_w iff there exists an attacker a and a domain knowledge diagram dk , in which d is written and a is referred to by the domain knowledge. The domain d is in \mathcal{D}_i iff there exists an attacker a and a domain knowledge diagram dk in which d is interfered and a is referred as sources of the interference. The domain d is in \mathcal{D}_r iff there exists an attacker a and a domain knowledge diagram dk in which the information in d is referred to and a reads this information. Based on the three sets of domains which might be attacked, the asset threat graph \mathcal{G}_{threat} can be set up. \mathcal{D}_w , \mathcal{D}_i , and \mathcal{D}_r are formally defined as follows.

$$\begin{aligned}
\mathcal{D}_w &= \{d : \mathcal{D}_{access} \mid \exists a : Attacker; dk : Diagram \bullet \\
& \quad (d, dk) \in write \wedge (a, dk) \in refersTo\} \\
\mathcal{D}_i &= \{d : \mathcal{D}_{access} \mid \exists a : Attacker; dk : Diagram \bullet \\
& \quad (d, dk) \in interfere \wedge (a, dk) \in refersTo\} \\
\mathcal{D}_r &= \{d : \mathcal{D}_{access} \mid \exists a : Attacker; dk : Diagram \bullet \\
& \quad (a, dk) \in read \wedge (d, dk) \in refersTo\}
\end{aligned}$$

\mathcal{G}_{threat} contains all edges, and therefore problem diagrams, of the corresponding asset access graph which

might allow an attacker to successfully attack the asset at hand. An access flow $(d_1, p, t, d_2) \in \mathcal{G}_{asset}$ represents that information is transferred from d_1 to d_2 that possibly comes from the asset or that possibly will be stored in the asset. Hence, such an access flow is a possible threat to the confidentiality of an asset if an attacker has the ability to read one of the domains d_1 or d_2 ($d_1 \in \mathcal{D}_r \vee d_2 \in \mathcal{D}_r$). In this case, we add the edge (d_1, p, t, C, d_2) to \mathcal{G}_{threat} . An access flow $(d_1, p, t, d_2) \in \mathcal{G}_{asset}$ is a possible threat to the integrity of an asset if an attacker has the ability to write the source d_1 of the access flow ($d_1 \in \mathcal{D}_w$), because an attacker could change the information of the asset or the information sent to the asset at domain d_1 , which forwards it to domain d_2 . In this case, we add the edge (d_1, p, t, I, d_2) to \mathcal{G}_{threat} . We have to consider an access flow $(d_1, p, t, d_2) \in \mathcal{G}_{asset}$ as a possible threat to the availability of an asset if an attacker has the ability to interfere one of the domains d_1 or d_2 ($d_1 \in \mathcal{D}_i \vee d_2 \in \mathcal{D}_i$), because an attacker is then able to threaten the availability of information flowing from or to the asset through the domains d_1 and d_2 . In this case, we add the edge (d_1, p, t, A, d_2) to \mathcal{G}_{threat} . \mathcal{G}_{threat} is defined as follows.

$$\begin{aligned} \mathcal{G}_{threat} = \{ & (d_1, p, t, cia, d_2) : Domain \times Diagram \times \\ & Type \times CIA \times Domain \mid (d_1, p, t, d_2) \in \mathcal{G}_{asset} \wedge \\ & [(d_1 \in \mathcal{D}_r \vee d_2 \in \mathcal{D}_r) \wedge cia = C \\ & \vee d_1 \in \mathcal{D}_w \wedge cia = I \\ & \vee (d_1 \in \mathcal{D}_i \vee d_2 \in \mathcal{D}_i) \wedge cia = A] \} \end{aligned}$$

The full attacker asset access graph \mathcal{G}_{attack} is an extension of $\mathcal{G}_{threat} \subset \mathcal{G}_{attack}$. We add an edge $(d_1, p, t, \varepsilon, d_2)$ to \mathcal{G}_{attack} iff (d_1, dk, t, d_2) is in \mathcal{G}_{asset} but not in \mathcal{G}_{threat} . These edges visualize how the attacks on the access flows in \mathcal{G}_{threat} are may propagated over the system due to the functional requirements. Additionally, the attackers are added to the attacker asset access graph. \mathcal{G}_{attack} contains an edge (a, dk, att, cia, d) if a is an attacker and a domain knowledge diagram dk exists, in which d is referred to and d is written ($cia = I$) or interfered ($cia = A$). Additionally, \mathcal{G}_{attack} contains an edge (a, dk, att, C, d) if a is an attacker and a domain knowledge diagram dk exists in which d is referred to and a is read. Formally, we define \mathcal{G}_{attack} as follows.

$$\begin{aligned} \mathcal{G}_{attack} = \{ & (d_1, p, t, \varepsilon, d_2) : Domain \times Diagram \times \\ & Type \times CIA \times Domain \mid (d_1, p, t, d_2) \in \mathcal{G}_{asset} \wedge \\ & \forall st : CIA \bullet (d_1, p, t, st, d_2) \notin \mathcal{G}_{threat} \} \cup \\ & \{ (a, dk, att, cia, d) : Attacker \times Diagram \times Type \times \\ & CIA \times Domain \mid (d, dk) \in refersTo \wedge \\ & [(a, dk) \in read \wedge cia = C \vee (a, dk) \in write \wedge cia = I \\ & \vee (a, dk) \in interfere \wedge cia = A] \} \cup \mathcal{G}_{threat} \end{aligned}$$

The generated attacker asset access graph for the persistent meter data is shown in Figure 7. Note that for reasons of readability, the PresSuRE tool merges edges and their annotation if they have the same source and target, and are of the same type. The *asset* is now visualized as *ellipse with bold border* and the

asset name (*PersistentMeterData*) is written in bold. The *attackers* internal and external network attacker are also added as *ellipses with dashed borders* and in italic font. Their *attack flow edges* are shown as dashed edges, which are annotated with the domain knowledge diagram they are described in and the security goals they may threaten. A bold (both, edge and annotation) access flow indicates a flow for which a security property might be threatened by an attacker. The threatened security property is annotated in brackets. For example, the implicit access flow edge between the nodes *LMN* and *TemporaryMeterData* is annotated with *RQ1 (A,C,I)*. Hence, it might be possible that for RQ1 the confidentiality, integrity and availability of persistent meter data is threatened.

5.4 Analyze Attacker Asset Access Graph

Last, a requirements engineer and a security expert have to analyze the problematic (potentially threatened) access flows. The input to this step are the attacker asset access graphs. The attacker asset access graph contains all information regarding access flows to and from the asset. And it contains the information where the asset might be threatened by an attacker. For each previously identified asset, it has to be checked if the original requirements related to the asset have to be augmented with security requirements. At this point we skip the detailed discussion how to do this analysis for reasons of space.

6 Validation

We validated PresSuRE using two real-life case studies, the already introduced smart meter and a voting system. The voting system requirements were obtained from a Common Criteria profile for voting systems (Volkamer and Vogt, 2008). For more details on the voting system case study, see (Faßbender and Heisel, 2013). The results for applying PresSuRE are reported in the following in detail for the Smart

Table 5: Results of the assessment for the smart meter case study

	Confidentiality	Integrity	Availability
Precision per attack asset access edges	36.14%	66.35%	62.52%
Recall per attack asset access edges	100.00%	100.00%	100.00%
Aggregated precision per attack asset access edges	55.00%		
Aggregated recall per attack asset access edges	100.00%		
Precision per requirement	92.59%	96.30%	96.30%
Recall per requirement	100.00%	100.00%	100.00%
Aggregated precision per requirement	95.06%		
Aggregated recall per requirement	100.00%		

Table 4: Effort spent for conducting the method

		Method Step								
		Model Functional Requirements		Security Knowledge Elicitation			Graph Generation			Analyze Attacker Access Graphs
		Model Problem Diagrams	Adjust Problem Diagrams	Prepare Knowledge Elicitation	Security Element Elicitation	Attacker(s) Elicitation	Global Access Graph	Asset Access Graphs	Attacker Asset Access Graphs	
Effort	∅ Per Item	28 min. per Problem Diagram	5.5 min. per Domain	< 1 sec. per Domain	10 min. per Domain	12 min. per Domain	3 sec. per Problem Diagram	9 sec. per Asset	40 sec per Attacker and Asset	15 min. per Problem Diagram
	Number of Items	27 Problem Diagrams	17 Domains	20 Domains	20 Domains	20 Domains	27 Problem Diagrams	14 Assets	14 Assets 7 Attacker	27 Problem Diagrams
	Total	12.5 person hours	1.6 person hours	0 computer hour	3.3 person hours	4 person hours	0 computer hour	0 computer hour	1 computer hour	6.75 person hours
	# Involved Persons	1	2	0	2	2	0	0	0	2
	Accumulated Total	12.5 person hours	3.2 person hours	0 computer hour	6.6 person hours	8 person hours	0 computer hour	0 computer hour	1 computer hour	13.5 person hours

Grid. The original functional requirements were obtained from (ope, 2009) and the NESSoS case studies provided by the industrial partners of the project. For conducting our method, we selected 13 minimum uses cases, which embody 27 requirements in total. For these requirements, 14 assets and 7 attackers of all kinds, as described in Section 5.2, were identified. Based on this information the graphs were generated, and the initial security requirements elicited.

We analyzed each attacker asset access graph for assessing the tool support, and we also analyzed the initial security requirements found for assessing the overall method. For the graph, we checked for each edge in the attacker asset access graph at hand if the annotated threats are existing according to the threats and security requirements of the original documents (e.g. (ope, 2009; Kreutzmann et al., 2011) for smart meter). We also looked for threats and security requirements which are defined in such documents, but which were not identified using PresSuRE. In this way we were able to measure the precision and recall of our method. Unfortunately, we do not know which security analysis was used for eliciting the security requirements reported in those documents. But we assume that security experts were involved in writing the documents and the documents were reviewed thoroughly. Hence, these documents are a good benchmark.

Next, we aggregated the results of the edges of the attacker asset access graph for each requirement. Thus, we derived for each requirement the information if the requirement has to be complemented by security requirements according to PresSuRE. Again, we also checked if the found security requirements are compliant with the original documents. Last, we measured the precision and recall of PresSuRE on the requirements level.

The results of this analysis for the smart meter is shown in Table 5. Speaking of the precision on the level of edges of the attacker asset access graph, we have many false positives, especially for confidentiality. This is because the original documents do not demand a high level of confidentiality. Additionally,

PresSuRE discovered potential indirect information flows between assets which will not occur in the system later on. Thus, PresSuRE is very strict and defensive, which is not appropriate in every case. Note that even though the indirect flows often turned out to be irrelevant, but they have to be checked anyway. Often attacks use such indirect relations to tamper with a system. Overall, the precision on the level of edges of the attacker asset access graph is acceptable (55%), but should be improved. The recall is perfect (100%) as we did not find any false negatives. On the requirements level, our results are satisfying. Whenever PresSuRE suggested to add a complementing security requirement for a functional requirement, this suggestion was correct with a precision of 95%, and no security requirement was missed (recall 100%).

Similar results were obtained for the voting system case study. The precision on the level of edges of the attacker asset graph is slightly higher, as the voting system documents are very strict regarding confidentiality. This fact also shows on the requirements level, but the difference to the smart meter case study is not significant. For the attacker asset access edge and the requirements level the recall was 100% again.

Speaking of the effort (see Table 4), the 43 person hours spent are a significant effort, but seem to be reasonable. The main effort is spent on the domain knowledge elicitation and the security analysis. Both are vital for security requirements elicitation and have to be conducted even when not using PresSuRE. But PresSuRE speeds up these activities and lowers the error rate, as it provides guidance and means for a detailed analysis and discussion. Additionally, from our experience, PresSuRE scales well, as most steps rely on the complexity of the environment. And the environment does not grow significantly when adding further requirements. When discussing the requirements itself, PresSuRE enables to focus on local and small problems, which also supports the scalability in means of comprehensibility and an only linear increase time for the security analysis.

7 Related Work

Schmidt and Jürjens (Schmidt and Jürjens, 2011) propose to integrate the SEPP method, which is based on problem frames, and UMLSec (Jürjens, 2005), which is based on a UML profile and allows tool-based reasoning about security properties. In this way, they can express and refine security requirements and transfer the security requirements to subsequent design artifacts. A similar method is described by Haley et al. (Haley et al., 2008), which also relies on problem frames for security requirements analysis. The first method (Schmidt and Jürjens, 2011) starts after the initial security requirements are already known, while the latter one already embodies a step for security requirements elicitation. But this particular step is described very sparsely and informally. Hence, our work can complement and improve these works.

There are many publications concerning goal-oriented security requirements analysis (e.g. (Liu et al., 2003; Mouratidis and Giorgini, 2007; Salehie et al., 2012; Van Lamsweerde, 2004)). But goal models are of a higher level of abstraction than problem frames. Goal models are stakeholder-centric, while problem frames are system-centric. Therefore, refining functional requirements taking into account more detail of the system-to-be and analyzing the system-to-be described by the functional requirements is reported to be difficult for goal-oriented methods (Alrajeh et al., 2009). Alrajeh et al. try to tackle this problem by introducing refinement steps which rely on heavy weight formalizations. We offer an alternative way of bridging the this gap. Thus, even though the goals of an attacker and their implication for the goals of stakeholders are already known, one might benefit from using our method. Recently, there were papers which reported a successful integration of goal-oriented and problem-oriented methods (Beckers et al., 2013b; Mohammadi et al., 2013). Use case-oriented security requirements methods such as misuse cases (Sindre and Opdahl, 2005) or abuse cases (McDermott and Fox, 1999) are one possible step before executing our method.

8 Conclusion

In this paper, we introduced a methodology for Problem-based Security Requirements Elicitation (PresSuRE). PresSuRE is a method for identifying security needs during the requirements analysis of software systems using a problem frame model. In summary, the PresSuRE method has following ad-

vantages: It introduces attacker asset access graphs and the graphs they are based on, which 1) visualize the access flows given by the functional requirements, 2) show entry points for attackers and subsequently threatened assets, 3) directly relate functional requirements and threats they are exposed to, and 4) can be generated fully automatically. And it is a re-usable requirements security analysis method which 1) relies on functional requirements only, as all security aspects are added while conducting the method in a structured way, 2) ensures that crucial domain knowledge is elicited, 3) enforces and supports the collaboration with system stakeholders and security experts, 4) is applicable to different domains, and 5) is tool supported to ease the elicitation, modeling, and analysis necessary for the method.

We validated our method and tool with two real-life case studies in the fields of smart grid and voting systems. The results show the suitability of our method to detect initial security requirements. For the future, we will investigate in which way the false positive rate can be lowered in order to decrease the effort spent on analyzing attack asset access graphs. Moreover, we are working on making the graphs to be analyzed even smaller by focusing a graph only on one attacker, and one asset property at a time. Furthermore, we plan to investigate how the basic CIA properties can be refined further systematically into more fine-grained security requirements such as authentication and authorization.

Acknowledgment

Part of this work is funded by the German Research Foundation (DFG) under grant number HE3322/4-2 and the EU project Network of Excellence on Engineering Secure Future Internet Software Services and Systems (NESSoS, ICT-2009.1.4 Trustworthy ICT, Grant No. 256980).

REFERENCES

- (2009). Requirements of AMI. Technical report, OPEN meter project.
- Alrajeh, D., Kramer, J., Russo, A., and Uchitel, S. (2009). Learning operational requirements from goal models. In *ICSE '09*, pages 265–275.
- Beckers, K., Faßbender, S., Heisel, M., and Meis, R. (2013a). A problem-based approach for computer aided privacy threat identification. In *APF '12*, pages 1–16. Springer.

- Beckers, K., Faßbender, S., Heisel, M., and Paci, F. (2013b). Combining goal-oriented and problem-oriented requirements engineering methods. In *CD-ARES '13*, pages 278–294.
- Beckers, K., Hatebur, D., and Heisel, M. (2013c). A problem-based threat analysis in compliance with common criteria. In *ARES '13*. IEEE Computer Society.
- Boehm, B. W. and Papaccio, P. N. (1988). Understanding and controlling software costs. *IEEE Transactions on Software Engineering*, 14(10):1462–1477.
- Cavusoglu, H., Mishra, B., and Raghunathan, S. (2004). The effect of internet security breach announcements on market value: Capital market reactions for breached firms and internet security developers. *Int. J. Electron. Commerce*, 9(1):70–104.
- Dolev, D. and Yao, A. C. (1983). On the security of public key protocols. *IEEE Transactions on Information Theory*, 29(2):198–207.
- Faßbender, S. and Heisel, M. (2013). From problems to laws in requirements engineering using model-transformation. In *ICSOFT '13*, pages 447–458. SciTePress.
- Firesmith, D. (2003). Specifying good requirements. *Journal of Object Technology*, 2(4).
- Haley, C. B., Laney, R., Moffett, J. D., and Nuseibeh, B. (2008). Security requirements engineering: A framework for representation and analysis. *IEEE Transactions on Software Engineering*, 34(1):133–153.
- Hatebur, D. and Heisel, M. (2010). Making pattern- and model-based software development more rigorous. In *ICFEM '10*, pages 253–269. Springer.
- Howard, M. and Lipner, S. (2006). *The Security Development Lifecycle : SDL : A Process for Developing Demonstrably More Secure Software*. Microsoft Press.
- ISO/IEC (2009a). Common Criteria for Information Technology Security Evaluation. ISO/IEC 15408, International Organization for Standardization (ISO) and International Electrotechnical Commission (IEC), Geneva ,Switzerland.
- ISO/IEC (2009b). Information technology - Security techniques - Information security management systems - Overview and Vocabulary. ISO/IEC 27000, International Organization for Standardization (ISO) and International Electrotechnical Commission (IEC), Geneva ,Switzerland.
- Jackson, M. (2001). *Problem Frames. Analyzing and structuring software development problems*. Addison-Wesley.
- Jürjens, J. (2005). *Secure Systems Development with UML*. Springer.
- Khansa, L., Cook, D. F., James, T., and Bruyaka, O. (2012). Impact of HIPAA provisions on the stock market value of healthcare institutions, and information security and other information technology firms. *Computers & Security*, 31(6):750 – 770.
- Kreutzmann, H., Vollmer, S., Tekampe, N., and Abromeit, A. (2011). Protection profile for the gateway of a smart metering system. Technical report, BSI.
- Liu, L., Yu, E., and Mylopoulos, J. (2003). Security and privacy requirements analysis within a social setting. In *RE '03*, pages 151–161.
- McDermott, J. and Fox, C. (1999). Using abuse case models for security requirements analysis. In *ACSAC '99*, pages 55–64.
- Mohammadi, N. G., Alebrahim, A., Weyer, T., Heisel, M., and Pohl, K. (2013). A framework for combining problem frames and goal models to support context analysis during requirements engineering. In *CD-ARES '13*, pages 272–288.
- Mouratidis, H. and Giorgini, P. (2007). Secure Tropos: a security-oriented extension of the tropos methodology. *International Journal of Software Engineering and Knowledge Engineering*, 17(2):285–309.
- Salehie, M., Pasquale, L., Omoronyia, I., Ali, R., and Nuseibeh, B. (2012). Requirements-driven adaptive security: Protecting variable assets at runtime. In *RE '12*, pages 111–120.
- Schmidt, H. and Jürjens, J. (2011). Connecting security requirements analysis and secure design using patterns and UMLsec. In *CAiSE '11*, pages 367–382. Springer.
- Sindre, G. and Opdahl, A. L. (2005). Eliciting security requirements with misuse cases. *Requir. Eng.*, 10(1):34–44.
- Van Lamsweerde, A. (2004). Elaborating security requirements by construction of intentional anti-models. In *ICSE '04*, pages 148–157.
- Volkamer, M. and Vogt, R. (2008). *Common Criteria Protection Profile for Basic set of security requirements for Online Voting Products*. Bundesamt f"ur Sicherheit in der Informationstechnik.
- Willis, R. (1998). *Hughes Aircraft's Widespread Deployment of a Continuously Improving Software Process*. AD-a358 993. Carnegie-mellon university.