

Model-driven Privacy Assessment in the Smart Grid

Fabian Knirsch*, Dominik Engel*, Christian Neureiter*, Marc Frincu[†] and Viktor Prasanna[†]

*Josef Ressel Center for User-Centric Smart Grid Privacy, Security and Control

Salzburg University of Applied Sciences

Urstein Sued 1, A-5412 Puch/Salzburg, Austria

Email: {fabian.knirsch, dominik.engel, christian.neureiter}@en-trust.at

[†]Ming-Hsieh Department of Electrical Engineering

University of Southern California

Los Angeles, USA

Email: {frincu, prasanna}@usc.edu

Abstract—In a smart grid, data and information are transported, transmitted, stored, and processed with various stakeholders having to cooperate effectively. Furthermore, personal data is the key to many smart grid applications and therefore privacy impacts have to be taken into account. For an effective smart grid, well integrated solutions are crucial and for achieving a high degree of customer acceptance, privacy should already be considered at design time of the system. To assist system engineers in early design phase, frameworks for the automated privacy evaluation of use cases are important. For evaluation, use cases for services and software architectures need to be formally captured in a standardized and commonly understood manner. In order to ensure this common understanding for all kinds of stakeholders, reference models have recently been developed. In this paper we present a model-driven approach for the automated assessment of such services and software architectures in the smart grid that builds on the standardized reference models. The focus of qualitative and quantitative evaluation is on privacy. For evaluation, the framework draws on use cases from the University of Southern California microgrid.

I. INTRODUCTION

In a smart grid a number of stakeholders (actors) have to cooperate effectively. Interoperability has to be assured on many layers, ranging from high level business cases to low level network communication. Data and information is sent from one actor to another in order to ensure effective communication. Furthermore, the exchange of vast amounts of data is crucial for many smart grid applications, such as demand response (DR) or electric vehicle charging [1], [2]. However, this data is also related to individuals and privacy issues are an upcoming concern [3], [4]. Especially the combination of data, e.g., meter values and preferences for DR can exploit serious privacy threats such as the prediction of personal habits. In system engineering, privacy is a cross-cutting concern that has to be taken into account throughout the entire development life-cycle, which is also referred to as *privacy by design* [1].

Model-driven privacy assessment is especially useful when applied in software engineering. Boehm et al. [5] thoroughly investigated the phases in software engineering and the expected costs for error correction and change requests. Costs double with every phase and once an application or a service is delivered, the additional adding of crosscutting concerns such as privacy is tied to enormous costs. As a result, design time privacy assessment is preferred in early phases of the software

engineering process. Therefore, a framework is needed to (i) model the system, including high-level use cases and concrete components and communication flows; and (ii) to assess the system's privacy impact using expert knowledge from the domain. Related work in the domain of automated assessments in the smart grid mainly focuses on security aspects and is not primarily concerned with privacy and the modeling in adherence to reference architectures.

In this paper we address these issues and present an approach for the model-driven assessment of privacy for smart grid applications. The framework proposed in this paper is designed to assist system engineers to evaluate use cases in the smart grid in an early design phase. For evaluation only meta-information is used and no concrete data is needed. We use Data Flow Graphs (DFG) to formally define use cases according to a standardized smart grid reference architecture. The assessment is based on an ontology driven approach taking into account expert knowledge from various domains, including customer views on privacy as well as system engineering concerns. The output is a set of threats and a quantitative analysis of risks, i.e., a number indicating the strength of that threat. To evaluate the system we draw on insights from the University of Southern California microgrid. The primary contributions of this paper are (i) the use of Data Flow Graphs to model use cases in the smart grid; (ii) the usage of data flow graphs for a quantitative privacy assessment; and (iii) the use of an ontology driven approach to capture domain knowledge.

The remainder of this paper is structured as follows: In Section II related work in the area of smart grid reference architectures, privacy evaluation and automated assessment tools is presented. In Section III the architecture of the proposed framework and its components are described. This includes the concept of DFGs for modeling use cases in the smart grid, the principal design of the ontology and the mapping of data flow graphs to the ontology, the methodology for defining threat patterns and finally, how these patterns are matched to use cases. The framework is evaluated with a set of representative use cases in Section IV. Section V summarizes this paper and gives an outlook to further work in this area.

II. RELATED WORK

In this section related work in the field of smart grid reference architectures, privacy evaluation and assessment as

well as automated assessment tools are presented. Often, privacy and security are used interchangeably. For the purpose of this paper we refer to privacy as legally accessing data but not using it for the intended purpose. Security, by contrast, would involve the illegal acquisition of data. In both cases, the well established and widely understood terminology from security assessment is used, i.e., *threat*, *attacker*, *vulnerability* and *countermeasure*.

A. Reference Models

Stakeholders in the smart grid come from historically different areas, including electrical engineering, computer science and economics. To ensure interoperability and to foster a common understanding, standardization organizations are rolling out reference models and road maps. In the US the NIST Framework and Roadmap for Smart Grid Interoperability Standards [6] and in the EU the Smart Grid Reference Architecture [7] were published. The European Smart Grid Architecture Model (SGAM) is based on the NIST Framework, but extends the model to better meet European requirements, such as distributed energy resources. In this paper we investigate use cases from the US. Investigations have, however, shown that for the purpose of this project all use cases from the US can be directly mapped to the European SGAM without the loss of information. Therefore we propose the utilization of the SGAM for two reasons: (i) the SGAM builds on the NIST model and allows to capture both, use cases from the US and the EU; and (ii) with the SGAM Toolbox Dänekas et al. [8] present a framework for modeling use cases based on the SGAM; in that way formally modeled use cases are the input for the evaluation.

B. Privacy

Privacy (and security) issues in the smart grid are addressed by standards in the US [9] and the EU [10]. Privacy, in specific, has no clear definition. According to a thorough analysis by Wicker and Schrader [11], privacy can be defined as the right of an individual's control over personal information. More formally this is defined by Barker et al. [12] in a four dimensional privacy taxonomy. The dimensions are *purpose*, *visibility*, *granularity* and *retention*. The *purpose* dimension refers to the intended use of data, i.e., what personal information is released for. The purpose ranges from single, a specific use only, to any. *Visibility* refers to who has permitted access. The range is from owner to all/world. *Granularity* describes to what extent information is detailed. The *retention* dimension finally is the period for storage of data. In any case, privacy is assured if all these dimensions are communicated clearly and fully disclosed to data owners and the compliance to the principles is governed. Hence, data is collected and processed for the intended purpose only, and the degree of visibility, granularity and retention is at the necessary minimum.

C. Assessment Tools

To measure the degree to which systems adhere to privacy requirements, approaches for automated qualitative assessments (resulting in statements of possible privacy impacts due to privacy critical actions or relationships) and quantitative

assessments (resulting in a numeric value that determines the risk of privacy impacts) exist.

In [13], Ahmed et al. present an approach towards ontology based risk assessment. The authors propose three ontologies, the *user environment ontology* capturing where users are working, i.e., software and hardware, the *project ontology* capturing concepts of project management, i.e., work packages and tasks and the *attack ontology* capturing possible attacks, e.g., non-authorized data access, virus distribution or spam emails. For a risk assessment, attacks (defined in the attack ontology) are matched with information available from the other ontologies. For a quantitative assessment, the annual loss expectancy is calculated by combining a set of harmful outcomes and the expected impact of such an outcome with the frequency of that outcome. The approach presented by Ahmed et al. is designed for security issues and does not explicitly cover privacy assessments.

In [14] and [15] an ontology driven approach for privacy evaluation is presented. The aim of these papers is to integrate privacy in the design process. High-level privacy statements are matched to system specifications and implementation details. The proposed *privacy by design* process includes the following phases: identification of high-level privacy requirements, translation of abstract privacy requirements to formal privacy descriptions, realization of the requirements and modeling of the system and analyzing the system by matching formal privacy requirements to the formal system model. Contrary to our work this approach is not focused on use cases in the smart grid and therefore does not model systems based on a standardized reference architecture.

A workflow oriented security assessment is presented by Chen et al. in [16]. This approach is not based on ontologies but on argument graphs. The presented framework uses *security goal*, *workflow and system description*, *attacker model* and *evidence* as an input. This information is aggregated in a discriminative set of argument graphs, each taking into account additional input. Nodes in the graph are aggregated using boolean expressions and the output is a quantitative assessment of the system. Instead of focusing on workflow analysis using graphs, we model systems as a whole in adherence to the standardized reference architecture using an ontology driven approach to integrate expert knowledge.

III. ARCHITECTURE

This section is dedicated to an architectural overview as well as a detailed discussion of the components. Figure 1 shows the principal components of the proposed architecture, including input and output. For a privacy assessment, the framework accepts two inputs, a use case *UC* modeled as a DFG in adherence to the SGAM and a set of threat patterns *T*. In order to qualitatively analyze this input the use case is mapped to individuals – i.e., instances of classes – of an ontology (sometimes referred to as the *assertion box*, *ABox* [17]). The corresponding class model (sometimes referred to as the *terminological box*, *TBox* [17]) is based on the SGAM. This qualitative analysis provides explicit and implicit information about the elements from the DFG: actors, components, information objects and their interrelation. The results

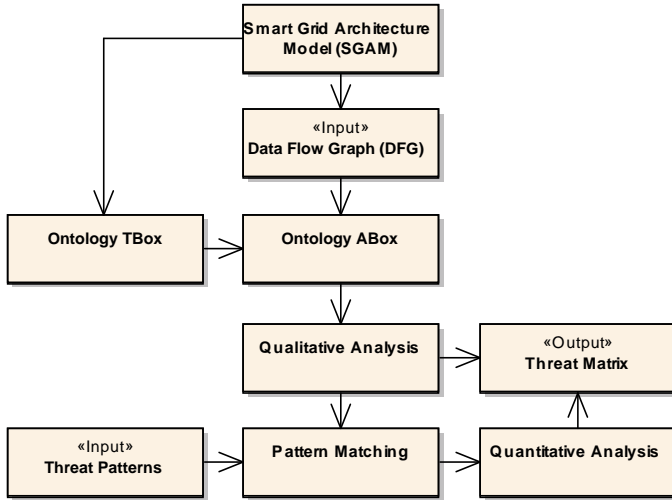


Fig. 1. Architecture overview showing input, output, components and principal information flows of the framework.

of the qualitative assessment are the input for the subsequent quantitative analysis. The output of that analysis is finally a class c from a set of classes C that the use case is assigned to. Threat patterns are used to describe potential threats $t \in T$ and a class represents a subset of threats T^* . A class c describes how threat patterns and the qualitative results are combined, which is presented as a threat matrix as an output. Formally, the classifier is defined as Assign UC to c_i if $t \in T_i^*$, $\forall t \in T, 1 \leq i \leq \{C\}$. A threat exploits a set of vulnerabilities and is mitigated by a set of countermeasures. Each threat pattern can be evaluated for itself or multiple patterns are combined to classes of threats. A vulnerability is any kind of privacy impact for any kind of stakeholder or actor. Threats are evaluated using the attack vector model which is adapted from security analysis and defined in detail later in this paper. In general, an attack is feasible, if given (i) an attacker; (ii) a privacy asset; and (iii) the resources to perform the attack. Hence, a receiver or collector of privacy critical data items is potentially able to access these assets and to use them in a way not corresponding to the original purpose. This is formally represented as \langle data access, privacy asset, attack resources \rangle .

A. Data Flow Graphs

In order to qualitatively and quantitatively assess the privacy impact of a use case a formalization is crucial. In this section we introduce the concept of Data Flow Graphs (DFG) for the smart grid based on a model-driven design approach originally presented in [8], [18]. DFGs formally capture all aspects of use cases in the smart grid in adherence to the SGAM. They contain high-level business cases as well as detailed views of a system's characteristics such as encryption and protocols. DFGs are a powerful tool as they allow both, easy modeling and full adherence to the reference architecture. Furthermore, in the graph relationships between actors, as well as the transported information objects (IO) are modeled. Nodes in a graph represent business actors, system actors or components and edges represent data flows annotated with

IOs. In accordance to the standard [7], DFGs consist of the following five layers:

- 1) Business Layer. In a DFG this layer is a high level description of the business case. Business actors, their common business goal and their business requirements are modeled.
- 2) Function Layer. The function layer details the business case by mapping business actors to system actors and by dividing the high level business goals in use cases and steps.
- 3) Information Layer. This layer describes information flows in detail. System actors communicate to each other through IOs. IOs are characterized by describing information attributes on a meta-level. An IO is one of the key data used for classification and is discussed in greater detail below.
- 4) Communication Layer. The communication layer is a more detailed view on communication taking into account network and protocol specifications.
- 5) Component Layer. In a DFG this layer contains concrete components. Therefore system actors are mapped to components and devices.

Each layer is a directed graph. Both, nodes and edges can have attributes. The semantics, however, are varying. For instance, where attributed edges in the business layer describe a business case, in the information layer concrete meta-data of communication flows are captured. Even though implicitly covered in the model presented above, for automated evaluation we introduce two additional layers: Between business and function layer we include the *Business Actor to System Actor Mapping* and between communication and component layer the *System Actor to Component Mapping*. This allows to capture the complexity of use cases on different levels while still maintaining the cross-layer relationship between high-level business actors and their representation as components. These layers are directed graphs as well, with edges indicating the mapping. The mapping defines a one to many relationship from business actors to system actors and from system actors to components. In the European Smart Grid Reference Architecture with the SGAM Methodology an approach for mapping use cases to the reference model is suggested. DFGs build on this methodology focusing on actors and their interrelation. An implementation for modeling DFGs in UML is available as the *SGAM Toolbox*¹. Data Flow Graphs contain explicit information (what is modeled) and implicit information (what can be concluded). Conclusions are drawn using ontology reasoning. Furthermore the framework allows to be extended with an arbitrary number of attributes to meet individual requirements.

B. Ontology Design

The ontology driven approach for classification has been chosen for two main reasons: (i) ontologies are powerful for capturing domain knowledge explicitly; and (ii) through logic reasoning [17] ontologies are a source for implicit knowledge. The power of ontologies to formally capture knowledge and

¹<http://www.en-trust.at/downloads/sgam-toolbox/>

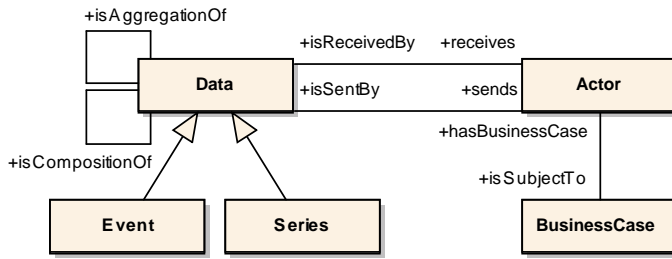


Fig. 2. Principal components of the ontology, showing a subset of the relationships between actor and data.

how to draw conclusions is discussed by Guarino et al. in [19]. The power of reasoning for gaining additional, implicit knowledge can easily be outlined with two examples: In a DFG, information objects may be sent from an actor A to an actor B and from there to another actor C . This is explicitly modeled in the DFG. A reasoner in an appropriate ontology, however, may conclude directly the transitivity, hence that actor A in fact sends information to actor C . Another example is concerned with compositions of data. An information object I_1 may contain sensitive data and it may be used by an actor D to compose another information object I_2 that is sent to a collecting actor E . It is not explicitly modeled in the DFG, but it can be concluded by the reasoner, that E receives an information object which is of type sensitive data since I_2 is a composition of I_1 . The ontology we propose here is designed to capture all aspects of a DFG. The ontology is modeled in OWL² and class expressions are stated in Manchester Syntax³. Therefore, all components available for modeling DFGs are represented either directly or as an abstraction in the ontology (referred to as the *TBox*). The DFG is represented in the ontology as a set of individuals (referred to as the *ABox*). Figure 2 depicts the principal classes and relationships of the ontology and therefore the most relevant concepts for mapping a DFG to the ontology. This view shows the main classes and relationships for illustration purposes only; our current ontology comprises more than 60 classes, data properties and object properties. Crucial concepts represented immediately, include which actor sends or receives which data and IO and how these IOs are composed. Furthermore, a set of pre-classifiers is defined to determine implicit knowledge. These classifiers are OWL classes using an equivalent class expression in Manchester Syntax. For instance, to determine if some aggregation consists of direct personal data, the following expression is used: `Data and isAggregationOf some DirectPersonalData`. To determine the multiplicity of the sending actor and if the data is a composition sent by many of such actors, more elaborate expressions can be phrased: `Data and isSentBy some Actor and Multiplicity value "n" and isCompositionOfMany some Data`.

²<http://www.w3.org/TR/owl-features/>

³<http://www.w3.org/TR/owl2-manchester-syntax/>

C. Threat Patterns

In this paper we evaluate the privacy impact on customers, thus we identified the following list of typical high-level threats based on literature reviews [1], [2], [4]. These threats have been modified in order to be more representative for the use cases from the University of Southern California microgrid that are investigated in this paper. Subsequently, IOs that may cause these threats are determined.

Customer presence at home. This privacy concern is discussed by Cavoukian et al. in [1]. To potentially determine a person's presence at home, some device in the customer premises is needed. This device collects data at a certain frequency, high enough to have a resolution that allows to draw conclusions on the energy usage of specific devices. Furthermore, data collected from that device needs to be sent to another actor (i.e., a utility). At the utility an individual or a system needs to have access to the data in an appropriate resolution. Since we always assume that data is accessed legally, we do not focus on unallowed data access. Additionally, the total delay of the data transmission is of relevance. If data is collected and transmitted in almost real time the presence at home can be determined immediately. If data is available with a delay only, the analysis of past events and predictions might be possible. If this information is published, an attacker might exploit this vulnerability in order to break in the house.

Tracking customer position. This threat is especially interesting for electric vehicle charging. Assuming the customer has some identification towards the charging station, at least the location, a timestamp and the amount of energy consumed will be recorded for billing. Depending on the design of the infrastructure only little information will be sent to the operator or a very detailed profile of the customer is maintained. Here, the multiplicity of the actors is crucial and the fact that different actors have access to the same data. Attacks for this threat are described by Langer et al. in [2], e.g., using information for targeted ads, for tracking movements to certain places or to infer the income based on recharges.

D. Pattern Matching

Actual classification is done in the pattern matching process. For each actor in the DFG and the ontology, respectively, the attack vector is determined, i.e., to which resources does an actor have access and what is the effort. If that shows feasible matching this is seen as a threat. It can be retrieved immediately from the ontology if an actor has access to a certain IO. This is done by evaluating actor and data object properties and by incorporating information from the pre-classifiers. Furthermore, relationships on the business layer and data properties such as encryption are taken into account. The following, discriminative set of classifiers is used to determine potential threats: first, for each information object the data provider and the data collector are determined (according to the terminology defined in [12]) and it is assessed who has access to the data. This yields a list of three-tuples in the form $\langle \text{information object (IO)}, \text{data provider (DP)}, \text{data collector (DC)} \rangle$. Then it is determined if an information object either contains sensitive

or direct personal data (according to the terminology defined in [20]). This yields another three-tuple in the form $\langle \text{information object (IO), sensitive (S), direct personal (DP)} \rangle$. Finally it is determined if the attacker has actual data access, yielding one more three-tuples in the form $\langle \text{information object (IO), data collector (DC), access (A)} \rangle$. Data access depends on the relationship of actors, on data resolution, retention and encryption. Matching these tuples to each other results in the components of the attack vector, recalling $\langle \text{data access, privacy asset, attack resources} \rangle$ yields $\langle \langle \text{IO, DP, DC} \rangle, \langle \text{IO, S, DP} \rangle, \langle \text{IO, DC, A} \rangle \rangle$. An exemplary attack vector for a DR use case where DR preferences are sent to the utility is $\langle \langle \text{DR preferences, customer, utility} \rangle, \langle \text{DR preferences, false, false} \rangle, \langle \text{DR preferences, utility, true} \rangle \rangle$. This already provides thorough qualitative analysis. It is possible to determine which actor can potentially threaten the privacy of another actor. It is even possible to conclude how and where this might happen. However, for a quantitative assessment the risk for a particular threat is calculated. While a qualitative assessment is useful in supporting detailed system design decisions and evaluation, for a very first outline of the overall system characteristics, a quantitative value is much more expressive. Further, providing a numeric value for the system's privacy impact helps to easily compare and contrast proposed designs. For the set T^* a number of patterns is defined. A pattern contains a set of conditions for vulnerabilities and countermeasures. Conditions are SPARQL ASK queries⁴ that return either *true* or *false* if the pattern applies or not. t'_v denotes the number of vulnerabilities that apply, t'_c the number of countermeasures that apply and t_v and t_c denote the total number of vulnerabilities and countermeasures, respectively. In this paper we propose the following approach for determining values for $P(t'_v, t'_c)$ and $L(t'_v, t'_c)$: $P(t'_v, t'_c)$ is determined by defining a plane that satisfies the following conditions: $P(t'_v = t_v, t'_c = 0) = 1$, $P(t'_v = 0, t'_c = t_c) = 0$ and $P(t'_v = 0, t'_c = 0) = \frac{1}{2}$. This yields $P(t'_v, t'_c) = \frac{1}{2}(\frac{t'_v}{t_v} - \frac{t'_c}{t_c} + 1)$. A condition that is of type *vulnerability* increases $L(t'_v, t'_c)$, a condition of type *countermeasure* decreases $L(t'_v, t'_c)$. The value of $L(t'_v, t'_c)$ is defined in the pattern. Risk R is finally defined by $R = \sum_{t \in T^*} P(t'_v, t'_c) L(t'_v, t'_c)$. To feed in the results gained from the qualitative analysis, certain variables in the query can be bound to instances. For example, given the following fraction of a query (where `usc` denotes the namespace prefix for actors and IOs in the University of Southern California microgrid) `$io usc:isSentBy ?systemactor . ?systemactor usc:isRealizationOf ?businessactor . ?businessactor a usc:BusinessActor` to determine if *some* information object is sent by *some* business actor. It is now possible to bind the variable `$io` to a concrete value as determined in the qualitative assessment, e.g., `$io ← InformationObject.CustomerName`. This allows to assess a particular impact on a particular information object or component/actor based on the previously calculated attack

vectors. We developed generic patterns for *typical* threats, i.e., such as the ones mentioned above. The framework is, however, not limited to this set of patterns and allows the definition of an arbitrary number of additional patterns to meet the individual needs of the application scenario. The output of the framework is a threat matrix contrasting the results from the qualitative analysis and from the quantitative risk assessment. For a *UC*, a threat matrix contains the attack vector and the assigned risk for the determined class c .

IV. EVALUATION

For evaluating the framework new, previously unused use cases are applied and the result is compared to state of the art assessments discussed in literature. We are therefore using a representative set of use cases describing typical applications in the smart grid. This includes, but is not limited to, smart metering, electric vehicle charging and DR. In this section a real-life use case from the University of Southern California microgrid is evaluated as an example. This use case has been chosen as it is (i) simple enough to verify results based on literature reviews; and (ii) complex enough to have an interesting combination of actors and information flows. We are focusing on a DR scenario similar to the one described in [21]. A customer interested in DR creates an online profile stating on which DR actions the customer is interested to participate (e.g., turning down air condition). When the utilities want to curtail load with DR, a customer whose profile fits the current requirements is sent a text message to, e.g., turn down the air condition. This message is acknowledged by the customer and the utility further reads the meter values to track actual power reduction. Besides the data flows mentioned, this further involves the storing of the profile and the past behavior of the customer for a more accurate prediction. For modeling this use case as a DFG, the following actors and IOs are identified.

Actors. Business actors are the *user* and the *utility*. The user is mapped to the system actors *smart meter, device and portal*. DR requests are sent to the user device (e.g., a cell phone) and the user's DR preferences are set in the portal (e.g., a web service). The smart meter is used to measure actual curtailment. The utility is mapped to a *DR repository*, containing preferences for each user and past behavior, to a *prediction unit* predicting DR requests based on the preferences and a *control unit* to meter user feedback and actual curtailment.

Information Objects. Cross-domain/zone information flows include user preferences sent to the utilities, DR requests sent to the user from the utility and both, the user acknowledge/decline and the meter values sent back to the utility. Information flows within the utilities' premises are from the DR repository to the prediction unit and from the control unit to the DR repository. Given the threat patterns introduced in Section III, we use our framework to determine the privacy impact of this use case which provides the following results.

Customer presence at home. The qualitative analysis shows that in the DR repository of the utility information about both, past customer behavior and customer data is brought together, i.e., direct personal data is composed with a detailed history of a person's actions. Furthermore, the customer's acknowledge/decline and the measured curtailment reveal if

⁴<http://www.w3.org/TR/sparql11-query/>

a customer (i) responded to the DR request; and (ii) actually participated in DR; both is a indication for the presence at home. For this threat we identified four vulnerabilities and one countermeasure, resulting in a *PO* of 0.9, an *EL* of 11.5 and a risk value of 10.35.

Tracking customer position. In our case, this threat might apply in two different scenarios: First, this threat is immediate if the acknowledge/decline response to DR requests contains the customer position (e.g., if sent by a cell phone or other mobile device). This does not only show the customers past and present position, but also if the customer is able to remotely control devices in his premises. Second, when the customer is represented by an additional component *electric vehicle charging station*. Assuming that DR requests are also sent with respect to the charging behavior. Based on the amount of energy the customer is willing to DR it might be possible to estimate the consumption of the electric vehicle and subsequently the traveled distance. For this threat we identified two vulnerabilities and one countermeasure, resulting in a *PO* of 0.66, an *EL* of 5 and a risk value of 3.3.

V. CONCLUSION AND FUTURE WORK

In this paper we introduced a framework for the model-driven privacy assessment in the smart grid. The framework builds on an ontology driven approach matching threat patterns to use cases that are modeled in adherence to standardized reference architectures. The approach presented here builds on meta-information and high-level data flows. It has been shown how to utilize this framework to successfully assess the privacy impact on use cases in early design time. Exemplary threats and exemplary use cases draw on insights from the University of Southern California microgrid. Future work will include an evaluation of the systems ability to generalize to arbitrary kinds of threats in the smart grid. Furthermore the system will be extended to serve as a policy decision point for system developers and customers in a smart grid IT infrastructure.

ACKNOWLEDGMENT

The financial support of the Josef Ressel Center by the Austrian Federal Ministry of Economy, Family and Youth and the Austrian National Foundation for Research, Technology and Development is gratefully acknowledged. Funding by the Austrian Marshall Plan Foundation is gratefully acknowledged. The authors would like to thank Norbert Egger for his contribution to the prototypical implementation.

This material is based upon work supported by the United States Department of Energy under Award Number number DE-OE0000192, and the Los Angeles Department of Water and Power (LA DWP). The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof, the LA DWP, nor any of their employees.

REFERENCES

[1] A. Cavoukian, J. Polonetsky, and C. Wolf, "Smartprivacy for the smart grid: embedding privacy into the design of electricity conservation," *Identity in the Information Society*, vol. 3, no. 2, pp. 275–294, 2010.

[2] L. Langer, F. Skopik, G. Kienesberger, and Q. Li, "Privacy issues of smart e-mobility," in *Industrial Electronics Society, IECON 2013 - 39th Annual Conference of the IEEE*, Nov 2013, pp. 6682–6687.

[3] P. McDaniel and S. McLaughlin, "Security and privacy challenges in the smart grid," *Security Privacy, IEEE*, vol. 7, no. 3, pp. 75–77, May 2009.

[4] Y. Simmhan, A. Kumbhare, B. Cao, and V. Prasanna, "An analysis of security and privacy issues in smart grid software architectures on clouds," in *IEEE International Conference on Cloud Computing (CLOUD)*, 2011. IEEE, 2011, pp. 582–589.

[5] B. Boehm, "A view of 20th and 21st century software engineering," in *Proceedings of the 28th International Conference on Software Engineering*, ser. ICSE 2006. New York, NY, USA: ACM, 2006, pp. 12–29.

[6] National Institute of Standards and Technology, "NIST Framework and Roadmap for Smart Grid Interoperability Standards, Release 2.0," National Institute of Standards and Technology, Tech. Rep. NIST Special Publication 1108R2, February 2012.

[7] CEN, Cenelec, and ETSI, "Smart Grid Reference Architecture," CEN/Cenelec/ETSI Smart Grid Coordination Group Std., Tech. Rep., November 2012.

[8] C. Dänekas, C. Neureiter, S. Rohjans, M. Uslar, and D. Engel, "Towards a model-driven-architecture process for smart grid projects," in *Digital Enterprise Design & Management*, ser. Advances in Intelligent Systems and Computing, P.-J. Benghozi, D. Krob, A. Lonjon, and H. Panetto, Eds. Springer International Publishing, 2014, vol. 261, pp. 47–58.

[9] National Institute of Standards and Technology, "Guidelines for smart grid cyber security: Vol. 2, privacy and the smart grid," The Smart Grid Interoperability Panel – Cyber Security Working Group, Tech. Rep., August 2010.

[10] CEN, Cenelec, and ETSI, "Smart Grid Information Security," CEN/Cenelec/ETSI Smart Grid Coordination Group Std., Tech. Rep., November 2012.

[11] S. Wicker and D. Schrader, "Privacy-aware design principles for information networks," *Proceedings of the IEEE*, vol. 99, no. 2, pp. 330–350, Feb 2011.

[12] K. Barker, M. Askari, M. Banerjee, K. Ghazinour, B. Mackas, M. Majedi, S. Pun, and A. Williams, "A data privacy taxonomy," in *Proceedings of the 26th British National Conference on Databases: Dataspace: The Final Frontier*, ser. BNCOD 26. Berlin, Heidelberg: Springer, 2009, pp. 42–54.

[13] M. Ahmed, A. Anjomshoaa, T. Nguyen, and A. Tjoa, "Towards an ontology-based risk assessment in collaborative environment using the semantictlife," in *Proceedings of the The Second International Conference on Availability, Reliability and Security*, ser. ARES 07. Washington, DC, USA: IEEE Computer Society, 2007, pp. 400–407.

[14] M. Kost, J.-C. Freytag, F. Kargl, and A. Kung, "Privacy verification using ontologies," in *Proceedings of the 2011 Sixth International Conference on Availability, Reliability and Security*, ser. ARES '11. Washington, DC, USA: IEEE Computer Society, 2011, pp. 627–632.

[15] M. Kost and J.-C. Freytag, "Privacy analysis using ontologies," in *CODASPY '12 Proceedings of the second ACM conference on Data and Application Security and Privacy*. San Antonio, Texas, USA: ACM, February 2012, pp. 205–2016.

[16] B. Chen, Z. Kalbarczyk, D. Nicol, W. Sanders, R. Tan, W. Temple, N. Tippenhauer, A. Vu, and D. Yau, "Go with the flow: Toward workflow-oriented security assessment," in *Proceedings of New Security Paradigm Workshop (NSPW)*, Banff, Canada, September 2013.

[17] R. Shearer, B. Motik, and I. Horrocks, "Hermit: A highly-efficient owl reasoner." in *OWLED*, ser. CEUR Workshop Proceedings, C. Dolbear, A. Ruttenberg, and U. Sattler, Eds., vol. 432. CEUR-WS.org, 2008.

[18] C. Neureiter, G. Eibl, A. Veichtlbauer, and D. Engel, "Towards a framework for engineering smart-grid-specific privacy requirements," in *Proc. IEEE IECON 2013, Special Session on Energy Informatics*. Vienna, Austria: IEEE, November 2013.

[19] N. Guarino, D. Oberle, and S. Staab, *What Is an Ontology?*, 2nd ed., ser. Handbook on Ontologies – International Handbooks on Information Systems, S. Staab and R. Studer, Eds. Springer, 2009.

[20] The European Parliament and the Council, "Official Journal L 281, 23/11/1995 P. 0031 - 0050 – Directive 95/46/EC of the European Parliament and of the Council of 24 October 1995," Online, November 1995.

[21] Y. Simmhan, Q. Zhou, and V. Prasanna, "Semantic information integration for smart grid applications," in *Green IT: Technologies and Applications*, J. H. Kim and M. J. Lee, Eds. Berlin Heidelberg, Germany: Springer, 2011, pp. 361–380.