Evaluating the Efficacy of LINDDUN GO for Privacy Threat Modeling for Local Renewable Energy Communities

Oliver Langthaler[a], Günther Eibl[b], Lars-Kevin Klüver and Andreas Unterweger[c]
Salzburg University of Applied Sciences, Urstein Süd 1, 5412 Puch/Hallein, Austria
{*oliver.langthaler, guenther.eibl, lars-kevin.kluever, andreas.unterweger*}*@fh-salzburg.ac.at*

Keywords: Threat Modeling, Privacy, Local Energy Communities, LINDDUN

Abstract: While security is considered an essential aspect of the design and implementation of many systems, privacy is often overlooked, especially in early planning phases. Although methodologies for the identification of privacy threats have been proposed, the number of studies outlining their practical application is limited. As a consequence, practical experience with these methods is sparse. This raises questions about their practicality and applicability for the energy domain. As a first step towards the assessment of the practical properties, we apply a lightweight version of the most prominent methodology, LINDDUN GO, to an intelligent charging use case for local renewable energy communities that is based on load forecasting. We find that one of the main advantages of LINDDUN GO is the completeness of the analysis, which was able to identify not only a built-in privacy deficiency but also unforeseen privacy threats for the considered use case. However, we also found that LINDDUN GO is not applicable for all privacy categories: Detectability was not assessable since it required detailed information that was not contained in our data flow graph in the design phase. In contrast, non-compliance was treated too generically, its intention is more to complete the list of important topics.

1 INTRODUCTION

The European Union aims to minimize carbon emissions through a variety of measures. Two of them are the introduction of smart metering and the possibility of allowing distributed generation and distribution of renewable energy.

Local, renewable energy communities (LECs) are a concept in the field of smart grid and sustainable energy, signifying a shift toward more distributed and localized forms of energy generation, distribution, and consumption. Broadly defined, energy communities are groups of individuals or organizations collaborating to produce, consume and manage energy resources, typically with an emphasis on renewable sources. These communities can take various forms, from small-scale, neighborhood-focused projects to larger collaborative initiatives involving multiple stakeholders (Bauwens, 2016; Walker, 2008). As part of the Clean Energy Package, the European Commission has established a legal framework for Energy Communities (Directorate-General for Energy, EC, 2019). The package defined two types of energy communities (1) Citizen Energy Communities and (2) Renewable Energy Communities, with different regulations and limitations (European Parliament and Council of the EU, 2018).

In Austria, local renewable energy communities are registered associations that can freely trade electricity between their members. Financial benefits arise from the avoidance of network tariffs as well as, typically, a gap between feed-in and consumption tariffs. For billing purposes, the smart meter data are forwarded by the Distribution System Operator (DSO) with a time resolution of 15 minutes. While these data can only be obtained at the end of the day, smart meter data of members can also be obtained through a legally mandated customer interface in (i) a timely manner and (ii) usually at a higher resolution (e.g. 5 second intervals[1]), depending on the respective DSO and smart meter model. These two properties of energy communities might enable additional use cases.

However, this might come at the cost of privacy which should, similarly to security, be considered already at the design phase. While many security threat

[a] https://orcid.org/0009-0001-4748-8272
[b] https://orcid.org/0000-0001-9570-5246
[c] https://orcid.org/0000-0002-3374-1636

[1] https://www.salzburgnetz.at/content/dam/salzburgnetz/dokumente/stromnetz/Technische-Beschreibung-Kundenschnittstelle.pdf

modeling methods are available (Azam et al., 2023), this is not the case for privacy, where LINDDUN (Deng et al., 2011) is the commonly considered standard. Although LINDDUN is widely known, work that demonstrates the practical application is rare: LINDDUN is applied to a scenario in the automotive domain in (Chah et al., 2023) and to a use case that deploys the OSIA standard for a national identity management architecture in (Nweke et al., 2022). This small number of examples of practical analysis might be attributed to the high startup cost that requires extensive privacy expertise and threat modeling expertise. The LINDDUN developers are aware of this issue, so they developed a more lightweight approach called LINDDUN GO (Wuyts et al., 2020). The goal of this paper is to perform a first case study that explores, whether this lightweight methodology is beneficial, especially for the energy domain, thereby adding to the empirical studies called for in (Wuyts et al., 2020).

More precisely, the goals are to determine (i) the extent of information needed for such an analysis and what knowledge is required from the analyst; (ii) if the methodology is suitable to detect both previously considered and unconsidered privacy deficiencies; and finally (iii) if the methodology is suitable to propose countermeasures. These questions are tackled by a case study, in which the additional capabilities of LECs are used to support the use case of an intelligent charging point.

This paper is organized as follows: Section 2 provides basic background information about threat modeling and analysis using LINDDUN and the LINDDUN GO card deck. Section 3 describes the process of its practical application using the concrete example of an electric vehicle (EV) charging use case within an LEC. This includes the required modeling of the use case, the identification of threats and the derived countermeasures. Section 4 discusses the fulfillment of the goals and gives an overview of possible future activities.

2 Background

Threat modeling is a process to identify and prioritize threats and vulnerabilities to an analyzed system, originating from software development (Shostack, 2014a). The most well-known methodology is STRIDE (Shostack, 2014a), see (Azam et al., 2023) for a survey. Inspired by STRIDE, LINDDUN was developed to model privacy threats (Deng et al., 2011). Both techniques are model-driven and based on data flow diagram (DFD) modeling of the system-under-investigation. A DFD is a graph constructed with the following elements: processes, data stores, entities, data flows and trust boundaries, see Figure 1. Trust boundaries divide the system into areas with varying levels of trust. Data flows that cross trust boundaries require special attention, as each of them represents a potential attack surface or entry point for security and possible transfer of sensible information for privacy. Based on the DFD, both STRIDE and LINDDUN first perform a threat analysis and then try to find countermeasures.

Figure 1: The elements used in DFDs

2.1 Privacy Threat Analysis with LINDDUN

LINDDUN utilizes DFDs as input and employs a systematic approach in which the individual components of a system, as well as the entirety of the system, are analyzed. The acronym LINDDUN represents the seven different types of privacy threats considered[2]: Linking (L): associating data items or user actions to learn more about an individual or group. Identifying (I): learning the identity of an individual. So LINDDUN distinguishes between threats in the context of identified data (where there is an explicit link), and identifiable data (where the link can for example be derived based on a pseudonym). Non-repudiation (Nr): being able to attribute a claim to an individual. As examples of this evidence include log files or digital signatures, non-repudiation is incompatible with common security means. Detecting (De): deducing the involvement of an individual through observation. Data Disclosure (Dd): excessively collecting, storing, processing or sharing personal data. Unawareness & Unintervenability (U): insufficiently informing, involving or empowering individuals in the processing of personal data. Non-compliance (Nc): deviating from security and data management best practices, standards and legislation.

(Wuyts et al., 2014) carry out an empirical examination of the LINDDUN methodology. Through three empirical user studies, the results of different groups of students are compared with a reference solution. They found, that correctness rate is satisfactory while its completeness rate could use some improvements. This underlines the importance of showing use cases in which the method is demonstrated and described in detail.

[2]https://linddun.org/threats/

2.2 The LINDDUN GO Card Deck

While the original version of LINDDUN provides a systematic and extensive assessment of a system's design, it has two main drawbacks: the analyst needs extensive privacy expertise and also experience with the threat modeling process itself. In order to enable a more wide-spread adoption of system privacy assessments in practice, a lightweight LINDDUN GO version has been created with the goal of reducing the initial effort to start privacy threat modeling (Wuyts et al., 2020). This was achieved by replacing the 100 leaf node threat tree of LINDDUN with 35 threat type cards, at a slight reduction of thoroughness. In addition, hotspots (as areas of the system from which a threat can originate) were introduced.

LINDDUN GO's modeling approach is akin to a card game, but it should not be mistaken for an actual game. It is intended to be performed by a group of participants, ideally comprising domain experts, system architects, developers, legal experts and data protection and information security officers. They use a basic representation of the system under analysis, as well as the card deck as a basis (example card see figure 2). Each card describes the hotspot, some elicitation questions to determine if the threat applies, an overall description, examples illustrating the threat, consequences demonstrating the importance of the threat and additional information. The participants in turn pick threat cards from a randomized deck and openly put them on the table. Subsequently, the applicability of the illustrated privacy threat with regard to the system or its users is assessed. For each listed hotspot, the given elicitation questions are carefully considered. If the threat is applicable for a given hotspot, it is documented accordingly. All players actively participate in this task; the higher their number, the smaller the odds of threats being overlooked. When no new threats are discovered, the next player draws a new card and the routine is repeated, until all cards (and thereby all threats) have been accounted for. This process typically takes several hours.

As work on LINDDUN GO is ongoing, the versions available at linddun.org are updated occasionally. The version applied in this paper is the latest GO-version which is available in digital form on the homepage (version v240118 at the time of writing). There is also a set of **p**hysical cards available from Agile Stationery[3]. While no version information for a purchase of the physical set made in 2023 was provided, the authors assume that this represents the original but now outdated version of LINDDUN GO, as it exhibits notable differences compared to the digital version. Despite having performed the analysis with both versions, only the results from the digital version are described. The currently available physical card deck is titled "2024 version", so it is assumed that the physical and digital versions are now identical.

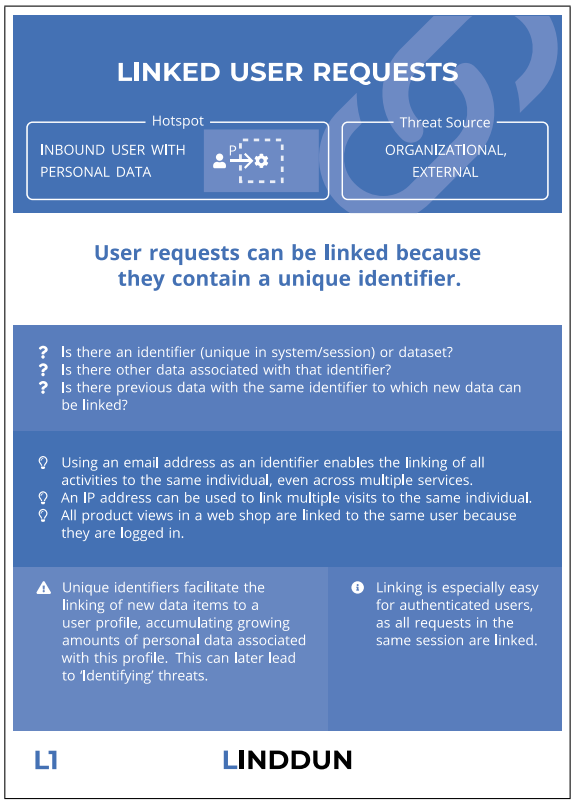

Figure 2: LINDDUN GO Card L1: Linked User Data.

3 Practical Application of LINDDUN GO

This section describes the practical application of LINDDUN GO. It is structured according to the three main steps: first the system is modeled as a DFD, then the card game is applied, and finally countermeasures are derived.

The result might be subjective and depend on the knowledge of the analysts. Therefore, a brief description of the analysis team is added: No team member has previous experience with LINDDUN or STRIDE, which is not ideal but may be representative of other analysis teams. The team consists of three people who share the following expertise among them: two computer scientists with common knowledge about systems engineering, two with domain knowledge

[3]https://agilestationery.com/collections/threat-modeling/products/linddun-go-privacy-threat-modeling-cards

about energy systems and one with knowledge about privacy-enhancing technologies. One of the members also has a background in software development as well as in business management. The preparation of the analysis resembles a practical scenario: it consisted of reading the instructions of the game and the included information about the definitions of the categories.

3.1 Description of Use Case Intelligent Charging Point

Due to the a rising number of EVs and high loads during charging, charging will have an increasing impact on the power grid. Consequently, self-consumption maximization within LECs using intelligent charging points will play an important role. Other performance criteria like Fairness, Quality of Service and Quality of Experience as discussed in (Danner and de Meer, 2021) need to be taken into account as well. In this work, we would further like to add the privacy perspective.

In this context, the use case of an intelligent shared charging station should serve as a basic test case, to see how well LINDDUN is able to identify associated privacy threats. The intention of the use case is to provide a forecast of future net production and consumption to the intelligent charging point which enables it to create an optimal charging schedule. For example, based on this information a household can identify times where less energy will be consumed within the LEC and therefore can be bought cheaper from the LEC than from the normal energy provider. Note that the forecast needs fine-grained metering data, i.e., net production and consumption, provided in a timely manner which is not possible with the data obtained from the DSO which is only available at the end of a day.

As the suitability of the methodology to detect existing privacy deficiencies should be investigated, two scenarios of this use case are created. In the first, external forecasting scenario, the LEC does not have any forecasting capabilities, see Figure 3 for a first, intuitive description. Therefore, forecasting is done by an external forecasting provider to whom the data are forwarded. This scenario is considered less private because the data needed for learning the prediction is sent to an external service provider. It has also been chosen such that privacy could be improved by data minimization, i.e., by only providing spatially aggregated metering data to the external forecasting service. In the second scenario (LEC forecasting), the forecasting can be the done within the LEC (Figure 4). As the LEC is considered more trustworthy than the external party, this scenario is expected to be more private. Although this latter scenario is more private, it is less likely, as in Austria a LEC is only a registered, non-profit association that likely lacks the tools and competence to produce accurate forecasts.

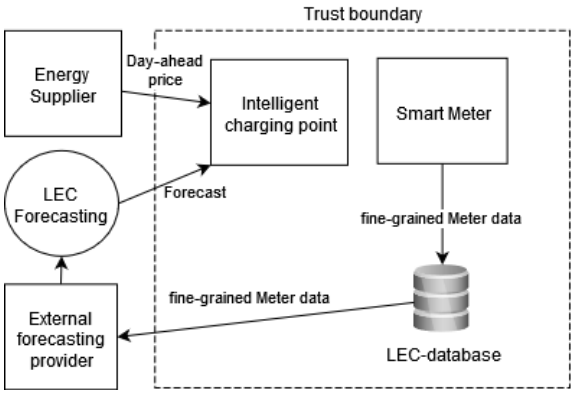

Figure 3: Scenario1: Intelligent EV-charging with external forecasting.

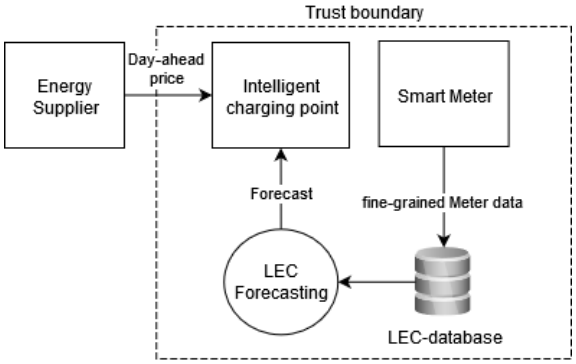

Figure 4: Scenario2: Intelligent EV-charging with internal forecasting.

3.2 Creation of a DFD

LINDDUN provides no instructions on how to create a suitable DFD. While we expected a DFD containing Data Flows between entities, we did not anticipate – before the study of the method – that we also need a description of the associated processes. However, despite that lack of instructions, we found that the coarse modeling of the DFD is rather straightforward: for example, it was clear where to introduce and locate processes. As an example, the DFD for the external forecasting scenario is shown in black and white in Figure 5. Note that for brevity, Figure 5 also contains some components, marked green, that are introduced as countermeasures in Section 3.4. The red data flows are the ones that were identified as most important for

the privacy analysis. Further note that in the internal scenario, no external party is needed (compare Figure 4). Thus, also the red data flow between the Community Energy Management System (CEMS) and the External forecasting provider does not exist.

Some aspects of this modeling were less obvious. The concept of a trust boundary stems from the security domain where it is used to separate different kinds of privileges like for example "read" or "write". In the case at hand, we decided to introduce three levels of trust which are separated by two trust boundaries. The highest trust level is clearly within a household. The next level consists of the LEC which is more trustworthy than external parties. This can be justified by an additional consideration of the billing use case. As in the billing use case the LEC gets meter data with a certain time granularity anyway, we deduce that the LEC is assumed to be trustworthy enough to get, store and process such meter data. By contrast, external parties have the lowest trust level and are not allowed to get these data. In addition, the exact kind of data transmitted is not exactly clear. We decided that in the external scenario, the CEMS just forwards the data without changes. This includes the identifiers of the smart meters which is especially important for the Linking threat.

3.3 Application of LINDDUN GO

Assumptions: A1: As part of a critical infrastructure, the smart metering system is assumed to have security measures in place: All users are authenticated. **A2**: As described in Section 3.2 the LEC is trustworthy enough to get and store meter data in the same time granularity as for billing, which is not true for the external forecasting provider. The outer trust boundaries of the DFD in Figure 5 indicate this trust. **A3**: As the information collected also needs to be usable for billing, it is assumed that some kind of identifier is contained in data sent between Member Energy Management System (EMS) and Community EMS. It is assumed, that this data is simply forwarded to the external forecasting provider.

Linking: In the internal forecasting scenario, the following threats were found: Threat L1 (Linked User Requests) is found due to the link between the Community EMS and the Member EMS (marked in red in Figure 5); L4 (Linkable Dataset) as the LEC Database stores linkable user data; L5 (Profiling Users) since consumption data has been shown to reveal private user patterns. However, as long as the time granularity is not finer than the one for the billing use case, this is not considered a problem. For the external scenario, these threats originate from the link between the Community EMS and the external forecasting provider as well as the data storage there. Because the forecasting provider is not considered trustworthy enough to have this kind of smart meter data (outer trust boundaries in the DFD), these threats are considered a problem for the external scenario and marked bold in Table 1. So, as expected at least one category was found where the external scenario is worse than the internal scenario.

Table 1: Summary of threats found using LINDDUN GO, named by card number for the scenarios where forcasting is done by the LEC and an external service provider, respectively. Threats that are a problem are printed in bold. The question mark "?" stands for unclear.

Category	scenario internal	scenario external
L	L1,L4,L5	**L1,L4,L5**
I	I1,I4,I5	**I1,I4,I5**
Nr	Nr1,...,Nr4	Nr1,...,Nr4
De	unassessable	unassessable
Dd	**DD1, DD4**	**DD1, DD4**
U	**U1,U3,U4?,U5**	**U1,U3,U4?,U5**
Nc	**Nc1,...,Nc4**	**Nc1,...,Nc4**

Identification: Table 1 shows that the found Identification threats have the same card numbers as the Linking threats. This does not happen by accident. The five Identifiability cards are structured the same way as the Linkability cards, so I1 corresponds to L1 and so on. This makes sense, as these two concepts are closely related. Since the sent data contains the identifiers (assumption A3), the analysis for Linkability also applies to identifiability, leading to the corresponding threats.

Non-repudiation has been denoted by the LINDDUN inventors themselves as a "category that only applies to some niche privacy applications (like whistle-blowing or e-voting systems), and not to the smart grid system" (Wuyts et al., 2014). The reasoning for the system at hand is the following: first, for billing purposes any form of plausible deniability should be avoided. In addition, for critical infrastructure security and its repudiation requirement must have a higher priority. Therefore, non-repudiation of service usage (Nr1), of sending (Nr2), of receipt (Nr3) and of storage (Nr4) cannot be fulfilled, but this is not a problem for both scenarios. Since no documents with hidden data or metadata are sent, Nr5 is not a threat.

Detectability: The questions in this threat category require detailed knowledge about the system: status messages (informational, warnings, errors) for De1 (detectable users) and De4 (detectable records), observability of communication for De2 (detectable

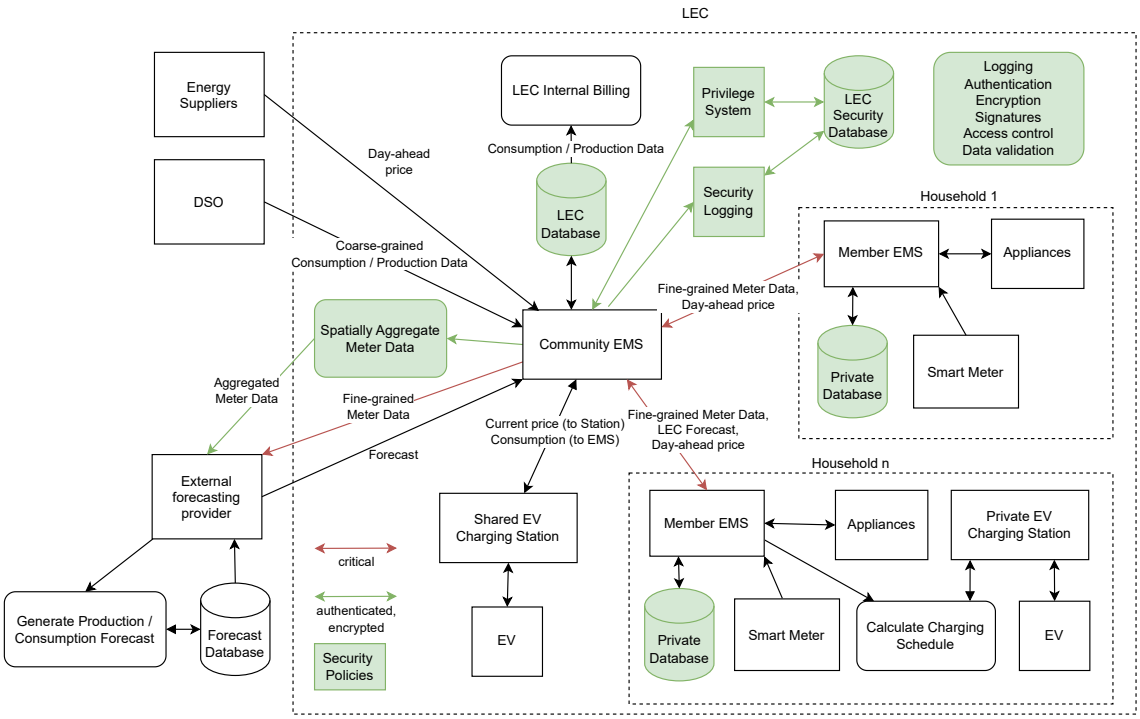

Figure 5: DFD for the external EV charging use case scenario. Red arrows are critical for the threat analysis, green parts denote privacy countermeasures.

service usage) and side effects of the system like traces left behind by deleted applications for De3 (detectable events). Therefore, this part of the threat analysis is considered as unassessable in the design phase based on a DFD such as the one shown here.

Data disclosure: We found threat DD1 (excessive sensitive is collected) as smart meter data have been shown to be sensitive with higher time resolution and aggregated data would suffice. As no system exists that deletes stored data, threat DD4 (data is stored longer than necessary) is also found for both scenarios: The system acquires more data than strictly needed for its functionality. Threat DD5 (Personal data is shared with more services or external parties than necessary) is not considered a threat. In the internal scenario this is clear. In contrast, in the external scenario, an external party gets personal data, but this data is needed to fulfill the forecasting task.

Unawareness: Since the current design of both systems does not provide a means to ensure awareness, it is tempting to consider all of these as threats. However, a strict analysis shows only U1 (lacking information about data processing and collection purpose), U5 (the ability to rectify or erase personal data) and U3 (insufficient privacy controls) as relevant for both scenarios, noting that no control of the time granularity is provided as modeled in the DFD. Threat U2 is ruled out, since users do have data that includes personal information of other users. For U2, the examples provided were especially helpful in clarifying the nature of the threat. The evaluation of threat U4 (access to personal data) is not completely clear in this context: one might consider it as a non-threat because the user has access to their data in its database with highest resolution and can therefore use it; in addition, rectification and erasure is treated by card U3. However, one might also consider it as a threat because a user can not directly access their data stored in the community EMS or at the forecasting provider.

Non-compliance is not in any way tackled by a simple DFD as modeled in Figure 5. Therefore, we considered all occurring threats because needed processes are missing: Nc1 (non-compliance with privacy regulations like GDPR for processing of personal data), Nc2 (non-compliance with privacy standards and best practices), Nc3 (improper data lifecycle management). While Nc4 is also considered a threat, it is fundamentally different because it covers best practices and standards of data *security* measures and processes. The LINDDUN creators are aware of this and write in the additional information section of card Nc4 that one should “consider complementary methods like security threat modeling”.

3.4 Derived Countermeasures

Linking and identification threats were identified for both scenarios but they are only a problem for the external scenario (first two lines in Table 1). These problematic threats are caused by sending the fine-grained meter-data to the external forecasting service. They could be mitigated by performing the forecasting inside the energy community, which results in the internal scenario. However, this might not be applicable since the energy community might not have the capabilities to create accurate forecasts. The second mitigation strategy introduces a process that spatially aggregates the meter data before sending them to the external forecasting service (marked in green in the DFD in Figure 5). This works because (i) the forecast should also work with aggregated data and (ii) data aggregation considerably enhances privacy. Note that there are many other options for mitigation such as federated learning and reducing the temporal resolution of the data, but a detailed evaluation and comparison of all of them is beyond the scope of this document.

Detection is unassessable in the current stage of modeling of the system. Therefore no countermeasures can directly be stated. However, the analysis process is still useful, since LINDDUN GO emphasizes to have an eye on details about handling of warnings or errors in a later stage.

Data disclosure: The two found threats DD1 and DD4 would require spatial aggregation directly during collection of measurements. The corresponding countermeasures are privacy-preserving aggregation protocols. These protocols, which are typically based on masking or the usage of additively homomorphic cryptosystems, are not easy to achieve in practice (Kursawe et al., 2011; Erkin et al., 2013; Li et al., 2010), which underscored the need to consider privacy already in the design phase. Note that spatial aggregation is also beneficial for DD5, although this was not considered a threat.

Unawareness needs to be addressed by a new, dedicated privacy module: the found threats can be used as some sort of requirements documents. For example, inform households about the processing of data and involved parties (U1), offer configuration possibilities such as the choice of time resolution (U3) or possibilities to correct/delete their stored data (U5).

Non-compliance cards do not lead to such direct countermeasures since they can not be evaluated based on the DFD. We consider the Non-compliance cards more as a list of topics that must be tackled in addition to privacy. While this does not solve the problem, it ensures completeness of the analysis. For security, one could consider utilizing the elevation of privilege card game (Shostack, 2014b; Tøndel et al., 2018) or using tools like ThreatGet (El Sadany et al., 2019). As an illustrative example for security countermeasures, the addition of a security module is roughly demonstrated in Figure 5.

4 Conclusion and Outlook

In this paper, the practical applicability of the LINDDUN GO Card deck for privacy threat modeling for LECs has been studied. The evaluation focused on the identification of the information needed to conduct the privacy analysis, LINDDUN's ability to detect privacy threats and the derivation of countermeasures.

Information needed and required knowledge of the analysts: Creating the DFD was relatively straightforward, as a simple diagram of the use case was already available and the team carrying out the threat modeling had a profound understanding of the use case. In contrast, in the first empirical study of LINDDUN in (Wuyts et al., 2014), creating the DFD is rated at medium difficulty by the participating students, researchers and practitioners.

It turned out that for the practical analysis of the card deck, a common and precise understanding of the threat is important. The examples provided on the cards are essential for this understanding. For example, based solely on the general description, we may have rated U2 as a privacy threat. However, the examples illustrate the concept more clearly, which led to it being classified as not being a threat. The practical analysis of the card deck was possible based on this DFD for all but two categories: detectability required more detailed information than contained in the DFD and Non-compliance was also unassessable. However, in the latter case we consider this category more as a way to reach completeness of important topics.

Detection of threats: We tested LINDDUN GO by considering an internal and external, less private scenario of the use case. LINDDUN GO did not find new, but the same threats in the external scenario. However, the questions were formulated in such a way, that the linkage and identification threats were not considered as a problem in the internal scenario. Therefore, we consider it as successful in detecting that the external scenario is less private. The detection of unconsidered privacy deficiencies is one of the main advantages of LINDDUN, as it aims at completeness. The methodology yielded many unintended threats especially in the Unawareness category, which likely stems from our

more technology-oriented approach. This is a consequence of the technology-oriented composition of the analysis team. While LINDDUN GO's methodology in itself is clearly beneficial, our experience confirms that diverse analysis teams are indeed desirable.

Derivation of countermeasures can not be done using LINDDUN GO alone as it requires some deeper privacy knowledge, for example about privacy enhancing technologies.

The present work is considered as a starting point in a detailed evaluation of LINDDUN GO. In this first step, experience is gained by a detailed analysis of a single use case. In future work, we would like to study the applicability on a broader variety of use cases and also for more complex use cases. One such example would consider flexibility provision, where the system operates in a distributed way. This happens, for example, when reinforcement learning is used for the determination of charging and discharging actions. Methodologically, we see the greatest potential for improvement in the development of more concrete Non-compliance cards (for example regarding compatibility with GDPR), as those were rather generic. We also intend to compare LINDDUN GO to the more complex, regular version LINDDUN.

Acknowledgements

Funding from the Federal State of Salzburg (project SEEEG) and the Austrian Research Promotion Agency (FFG project number 881165) is gratefully acknowledged.

REFERENCES

Azam, N., Michala, L., Ansari, S., and Truong, N. B. (2023). Data privacy threat modelling for autonomous systems: A survey from the gdpr's perspective. *IEEE Transactions on Big Data*, 9:388–414.

Bauwens, T. (2016). Explaining the diversity of motivations behind community renewable energy. *Energy Policy*, 93:278–290.

Chah, B., Lombard, A., Bkakria, A., Yaich, R., and Belfort, F. (2023). Exploring Privacy Threats in Connected and Autonomous Vehicles : An Analysis. *Journal of Ubiquitous Systems & Pervasive Networks*, 19(1):25–32.

Danner, D. and de Meer, H. (2021). Quality of service and fairness for electric vehicle charging as a service. *Energy Informatics*, 4:1–20.

Deng, M., Wuyts, K., Scandariato, R., Preneel, B., and Joosen, W. (2011). A privacy threat analysis framework: supporting the elicitation and fulfillment of privacy requirements. *Requirements Engineering*, 16:3–32.

Directorate-General for Energy, EC (2019). *Clean energy for all Europeans*. Publications Office of the European Union.

El Sadany, M., Schmittner, C., and Kastner, W. (2019). Assuring compliance with protection profiles with threatget. In *International Conference on Computer Safety, Reliability, and Security (SAFECOMP) 2019*, pages 62–73.

Erkin, Z., Troncoso-Pastoriza, J. R., Lagendijk, R. L., and Perez-Gonzalez, F. (2013). Privacy-preserving data aggregation in smart metering systems: An overview. *IEEE Signal Processing Magazine*, 30:75–86.

European Parliament and Council of the EU (2018). Directive (eu) 2018/2001 of the european parliament and of the council of 11 december 2018 on the promotion of the use of energy from renewable sources (recast). *Official Journal of the European Union*, 61:82–209.

Kursawe, K., Danezis, G., and Kohlweiss, M. (2011). Privacy-friendly aggregation for the smart grid. In *Privacy Enhanced Technology Symposium*, pages 175–191.

Li, F., Luo, B., and Liu, P. (2010). Secure information aggregation for smart grids using homomorphic encryption. In *Proceedings of First IEEE International Conference on Smart Grid Communications*, pages 327–332. IEEE.

Nweke, L. O., Abomhara, M., Yayilgan, S. Y., Comparin, D., Heurtier, O., and Bunney, C. (2022). A linddun-based privacy threat modelling for national identification systems. In *Proceedings of the 2022 IEEE Nigeria 4th International Conference on Disruptive Technologies for Sustainable Development, NIGERCON 2022*. Institute of Electrical and Electronics Engineers Inc.

Shostack, A. (2014a). Elevation of privilege: Drawing developers into threat modeling. In *2014 USENIX Summit on Gaming, Games, and Gamification in Security Education (3GSE 14)*.

Shostack, A. (2014b). *Threat Modeling: Designing for Security*. Wiley Publishing, 1st edition.

Tøndel, I. A., Oyetoyan, T. D., Jaatun, M. G., and Cruzes, D. (2018). Understanding challenges to adoption of the Microsoft Elevation of Privilege game. In *Proceedings of the 5th Annual Symposium and Bootcamp on Hot Topics in the Science of Security*, pages 1–10.

Walker, G. (2008). What are the barriers and incentives for community-owned means of energy production and use? *Energy Policy*, 36:4401–4405.

Wuyts, K., Scandariato, R., and Joosen, W. (2014). Empirical evaluation of a privacy-focused threat modeling methodology. *Journal of Systems and Software*, 96:122–138.

Wuyts, K., Sion, L., and Joosen, W. (2020). LINDDUN GO: A Lightweight Approach to Privacy Threat Modeling. In *IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*, pages 302–309.