LINDDUN GO: A Lightweight Approach to Privacy Threat Modeling

Kim Wuyts
imec-DistriNet, KU Leuven
3001 Leuven, Belgium
kim.wuyts@cs.kuleuven.be

Laurens Sion
imec-DistriNet, KU Leuven
3001 Leuven, Belgium
laurens.sion@cs.kuleuven.be

Wouter Joosen
imec-DistriNet, KU Leuven
3001 Leuven, Belgium
wouter.joosen@cs.kuleuven.be

***Abstract*—Realizing privacy-preserving software requires the application of principles such as Privacy by Design (PbD) which require the consideration of privacy early on in the software development lifecycle. While privacy threat modeling approaches, such as LINDDUN, provide such a systematic and extensive assessment of a system's design, their application requires the analyst performing the assessment to have (i) extensive privacy expertise and (ii) sufficient experience with the threat modeling process itself. Hence, there is a high startup cost to apply these techniques. To reduce this initial threshold, more lightweight privacy analysis approaches are necessary.**

In this paper, we (i) discuss the requirements for early lightweight privacy analysis approaches; (ii) present LINDDUN GO, a toolkit that supports lightweight privacy threat modeling; (iii) describe the pilot studies that were conducted for the preliminary evaluation with industry professionals.

The availability of lightweight privacy analysis approaches reduces the initial effort to start privacy threat modeling and can therefore enable a more wide-spread adoption of system privacy assessments in practice.

***Index Terms*—threat modeling, privacy by design, privacy engineering**

1. Introduction

Creating secure and privacy-preserving software by design requires the assessment of privacy problems early on in the software development lifecycle. Threat modeling supports such an approach, as it is a method to systematically elicit and mitigate privacy and security threats in a software architecture. It therefore requires both extensive knowledge of the system's domain and architecture, as well as expert knowledge on privacy and/or security, and experience with the threat modeling process itself. Fortunately, threat modeling frameworks, such as LINDDUN [1], [2] and STRIDE [3], [4], provide both process and knowledge support. The threat modeling process systematically guides the analyst through the different steps, which are strengthened with privacy (or security) knowledge. To guarantee a full-fledged analysis, each component (or interaction) of the system should be systematically examined for potential privacy threats. This is labor-intensive and still requires quite some expertise of privacy concepts and sufficient experience with the threat modeling process. Therefore, the threshold is relatively high and people tend to only use fragments of the approach, or even just the mnemonic for brainstorming rather than apply the envisioned full threat modeling approach.

By providing more lightweight hands-on support for privacy threat modeling, this burden could be alleviated. We therefore propose LINDDUN GO, which aims to achieve this by simplifying both the LINDDUN method and the provided knowledge, while still adhering to the core process and privacy knowledge. This lightweight approach can thus lower the threshold to get started with threat modeling.

The contributions in this paper are threefold: (i) we describe the challenges with current approach and propose requirements for a lightweight approach; (ii) we propose LINDDUN GO, a lightweight privacy threat modeling approach; and (iii) we present preliminary evaluation results, among others obtained from industry professionals.

This paper is structured as follows. Section 2 introduces the background on LINDDUN privacy threat modeling and its current state of practice. Section 3 suggests requirements for a lightweight approach to fill the gap between a full-fledged privacy threat modeling approach and an unstructured brainstorming exercise. The section concludes with a mapping of these requirements to known privacy threat modeling approaches to highlight the current gap and potential value and opportunity for a lightweight approach. Section 4 introduces LINDDUN GO. The preliminary evaluation is presented in Section 5 and discussion and future work are described in Section 6. Finally, Section 7 concludes the paper.

2. Background

This section summarizes threat modeling approaches, with a focus on LINDDUN, which will be used as basis to create a lightweight privacy threat modeling approach, and highlights its current state of practice.

2.1. LINDDUN privacy threat modeling

LINDDUN [1], [2], [5] is a privacy threat modeling framework that provides support to systematically elicit and mitigate privacy threats in software architectures. It was inspired by STRIDE [3], [4], [6], Microsoft's threat modeling approach for security. STRIDE was developed more than 20 years ago at Microsoft as part of their security development lifecycle. LINDDUN is currently one of the most mature privacy threat modeling approaches. While

This research is partially funded by the Research Fund KU Leuven, the Vlaio TETRA-GDPR project, and the Cybersecurity Initiative Flanders.

also other (security) threat modeling approaches exist [7], [8], all follow the same four high-level steps described as four questions by Shostack [4]: (1) What are you building? (2) What can go wrong? (3) What are you going to do about it? (4) Did you do an acceptable job?

LINDDUN provides both process and knowledge support, i.e. each of the steps of the method are complemented with privacy-specific knowledge to aid the analysts. The knowledge is structured according to the threat categories encompassed in the LINDDUN acronym: linkability, identifiability, non-repudiation, detectability, disclosure of information, unawareness, non-compliance. LINDDUN consists of the following three steps:

1) *Model the system.* Typically a data flow diagram (DFD) [9] is used, which is a very simple representation of the system consisting of only 5 building blocks: processes (i.e. computational units), data stores (i.e. passive containers of information), external entities (i.e. users or third party services), data flows (i.e. communication between the different components), trust boundaries (i.e. logical or physical borders).
2) *Elicit threats.* To elicit threats, you need to systematically go over each of the DFD elements or interactions and determine which threats are applicable. LINDDUN provides knowledge support in the form of privacy threat trees, which describe the most common attack paths for each combination of LINDDUN threat category and DFD element type. Each identified threat needs to be documented.
3) *Manage threats.* In this solution-oriented step, the threats are prioritized according to their risk, and then suitable mitigation strategies and corresponding privacy enhancing solutions can be selected to resolve each elicited threat.

2.2. LINDDUN in practice

LINDDUN has received growing attention from both academia and industry. LINDDUN has been examined and applied in several academic projects [10], [11] and has been recommended by several authorities in the field of privacy engineering [12]–[14].

Similar to observations in security threat modeling [15], a full-fledged application of LINDDUN can be considered complex and time-consuming. LINDDUN has frequently been used as a mnemonic for a brainstorming-style exercise rather than a means to systematically elicit privacy threats. The LINDDUN threat trees are sometimes considered complex to use. While they provide a valuable overview of potential threat types, they may lack some semantics and only have minimal selection criteria to support the applicability assessment of potential threats [16]. The content itself requires a sufficient level of privacy expertise (similar to the STRIDE knowledge for security [15]). In addition, the assessment of each potential threat currently needs to be executed manually. The rather extensive set of threat types (i.e. leaf nodes in the LINDDUN trees that describe potential issues in the system) that need to be systematically examined for each applicable DFD element or interaction quickly leads to a large number of individual threats that need to be analyzed. Each step also needs to be fully documented. An element-based execution of LINDDUN, for instance, implies that a large mapping table with a row for each DFD element and column for each threat category needs to be created, and systematically updated for each covered cell. Each identified threat needs to be properly documented, and also assumptions should be made explicit. So, in addition to the already quite labor-intensive exercise of manually eliciting threats, creating the expected documentation also requires a lot of time.

Feedback from industry highlights that (i) the fairly complex content of the LINDDUN threat trees requires considerable privacy expertise, and (ii) the overhead of the labor-intensive systematic elicitation and documentation are the main concerns that prevent the use of a full privacy threat modeling exercise. The rather big gap between its prescribed use and its application in practice provided inspiration to create a more lightweight version of LINDDUN which aims to lower the threshold of required expertise and reduce the overall effort.

3. Requirements for a lightweight approach

As the state of practice shows (Section 2.2), a sufficient level or privacy expertise is required. In addition, the threat modeling method, though reasonably simple, still requires a lot of manual effort [15]. Tool support could help reduce some of the effort, but currently available tools [17]–[20] require more extensive models as input (and thereby put a heavier burden on the modeling step) and often still require a significant amount of manual assessment. This paper focuses on scoping the method and supporting knowledge, rather than creating tooling for existing approaches.

Based on the feedback received throughout the years from industry professionals who have put LINDDUN to the test, and based on the empirical studies on privacy and security threat modeling we have executed in the past years [21]–[24] and the (limited set of) experience reports from industry on threat modeling [15], [25], we distilled requirements for a lightweight approach. We have categorized these into the two main threat modeling building blocks: the *method* and the provided *knowledge* (also known as technique and repertoire [4]). These requirements focus on lightening the elicitation step of the threat modeling process. The final part of this section maps the requirements to existing privacy threat modeling approaches.

3.1. Method requirements

REQ 1 - Simple. While the essence of the threat modeling process is fairly easy (i.e. systematically iterate over each system element), the execution is sometimes still considered too complex, especially when analyzing less critical systems. For each system element, the applicability of potential threats needs to be assessed, which is far from trivial. In order to get people eager to get started with an approach, the actions to be executed should be as clear and easy as possible. Therefore, simplifying the process itself and reducing the expected documentation would lighten the method as a whole, while, ideally, also the time required to complete a first and complete threat modeling exercise will be decreased.

REQ 2 - Comprehensive. Although the process itself should be fairly simple to apply, it should still remain comprehensive. As one of the strengths of threat modeling is its systematic approach, a lightweight variant should still be able to guarantee a sufficient degree of coverage by taking into account an as complete as possible set of potential threat types in combination with all relevant system elements.

REQ 3 - Collaboration. Typical threat modeling approaches are mainly shaped to aid individual threat modelers (in isolation). This implies that the threat modeler solely carries the burden and responsibility of creating a correct representation of the system. A threat modeling exercise in practice should however aim to gather a diverse group of participants who can cover all aspects of the system (ranging from the threat modeler and privacy and/or security expert, to the business stakeholder, developer, architect and project manager). Ideally, these stakeholders are not only around the table to fill in the gaps in the system model that is jointly being assembled, but they can also provide valuable insights in the threat elicitation process. Having at least some guidance on how the method could be applied by a team would be thus useful.

3.2. Knowledge requirements

REQ 4 - Understandable privacy knowledge description. The description of potential threats (i.e. threat trees) is typically very limited, both in LINDDUN and STRIDE. This can hinder adoption by novice privacy analysts as they typically require more information than a brief (often one-sentence) summary of potential threats. In order for non-experts to be able fully grasp the actual threat, a more extensive documentation of threat types is required. This documentation could be extended with, for instance, specific examples to illustrate the potential threats, and the threat type's consequences to highlight relevance and importance of the threat.

REQ 5 - Applicability criteria. In addition to comprehensible content in general, it would be useful if the provided knowledge would be able to actually guide the user in the selection and assessment process. By explicitly including applicability criteria, the user can quickly assess the relevance of the threat description with respect to specific components (or context) of the system-under-analysis.

3.3. Comments on existing privacy threat modeling approaches

To illustrate the current gap in lightweight privacy threat modeling support, we map known approaches to our requirements.

Given the heavyweight nature of threat modeling, many analysts tend to come up with their own more ad-hoc method. Some more lightweight approaches have however been formalized. They mainly focus on security, and are aimed at lowering the threshold to get started with threat modeling by gamifying the process. The Elevation of Privilege game (EoP) [26] is probably the best known example. It is a card game that introduces the players to threat modeling. The card deck has two extensions for privacy, TRIM [27] and STRIPED [28]. STRIPED introduces 13 additional privacy threat cards. TRIM introduces 35 privacy threat cards, divided in 4 categories: transfer, retention/removal, inference, and minimization. Other gamified approaches[1] have mainly an educational purpose [29] or focus on security [30], [31].

Evaluation. In the remainder of this section, we map STRIPED and TRIM to the requirements we position in this paper and compare them to the original LINDDUN approach.

LINDDUN is not considered as a *simple* process to execute. STRIPED and TRIM, on the other hand, are very simple to get started with. The LINDDUN process systematically iterates over each DFD element, making it a very *comprehensive* method. STRIPED and TRIM simplified their process by limiting the iteration over all DFD elements to just iterating until 1 threat has been identified. Obviously the comprehensiveness of the process is lower. Concerning their content, STRIPED condensed privacy to merely 12 threat cards, while TRIM describes 31 distinct threats. TRIM can thus, content-wise, be considered more elaborate than STRIPED. As extensions of EoP, STRIPED and TRIM are also *collaborative*. LINDDUN itself does not provide any explicit support more collaboration. Both LINDDUN, STRIPED and TRIM extensions only have a very limited *description* of each potential threat (i.e. a one-sentence summary). This thus requires sufficient privacy expertise to use. Concerning *applicability criteria*, LINDDUN only scopes each threat tree to the applicable DFD element type. STRIPED and TRIM do not provide guidance with respect to applicability. While STRIPED and TRIM are overall great approaches to get started with privacy threat modeling, our requirements highlighted some opportunities of improvement.

4. LINDDUN GO

LINDDUN GO is a toolkit, in the sense that it offers both privacy threat information to use and guidelines on how to apply it in a systematic way. It consisting of a set of threat type cards (of which an example is shown in Figure 1) that describe the most common privacy threats for each threat category. LINDDUN GO is primarily targeted at industry professionals who want to get started with privacy threat modeling, but can also be used by more experienced privacy analysts who are looking for a more lightweight approach, yet still want to adhere to the systematic threat modeling method and privacy knowledge.

In this section, we first describe how LINDDUN GO was created, then introduce LINDDUN GO's threat type cards and overall process.

4.1. Creating LINDDUN GO

Looking for a way to lower the threshold when learning and applying LINDDUN, we found inspiration in the Elevation of Privilege (EoP) card game and its privacy extension, STRIPED and TRIM. We particularly valued the card deck representation as it is very accessible and inviting. Experience however showed that people who

[1] An overview of existing information security cards and games is available on https://adam.shostack.org/games.html

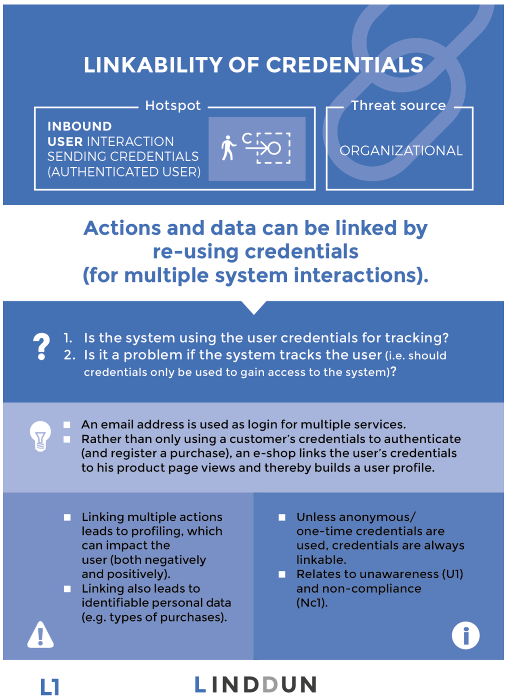

LINKABILITY OF CREDENTIALS

Hotspot

INBOUND
USER INTERACTION
SENDING CREDENTIALS
(AUTHENTICATED USER)

Threat source

ORGANIZATIONAL

Actions and data can be linked by re-using credentials (for multiple system interactions).

1. Is the system using the user credentials for tracking?
2. Is it a problem if the system tracks the user (i.e. should credentials only be used to gain access to the system)?

- An email address is used as login for multiple services.
- Rather than only using a customer's credentials to authenticate (and register a purchase), an e-shop links the user's credentials to his product page views and thereby builds a user profile.

- Linking multiple actions leads to profiling, which can impact the user (both negatively and positively).
- Linking also leads to identifiable personal data (e.g. types of purchases).

- Unless anonymous/ one-time credentials are used, credentials are always linkable.
- Relates to unawareness (U1) and non-compliance (Nc1).

L1

LINDDUN

Figure 1. Example LINDDUN GO threat type card - Linkability of credentials (L1)

are new to the security or privacy domain require more information to properly grasp each threat. We therefore decided to extend the description of each threat type. The collaborative approach of EoP was also an element we wanted to integrate in our lightweight privacy approach. The competition aspect has however been made optional, as an initial empirical study with industry professionals (of which more details are provided in Section 5.2) showed that there was no interest in a competitive approach.

We also decided to deviate from the EoP overall process to more closely resemble the LINDDUN process, as this would lead to a higher level of thoroughness. Rather than quickly jumping from one card to the next, LINDDUN GO nudges towards a systematic analysis of each applicable architectural component.

Reduced scope. The full-fledged LINDDUN requires to systematically iterate over all threat trees. With more than a 100 leaf nodes to consider, this can be quite time-consuming. LINDDUN GO condensed this knowledge into 35 threat type cards to consider. Part of this reduction is obtained by combining related threat types (e.g. the lack of data portability closely relates to the lack of data access). Also some less important threats are discarded in LINDDUN GO. For instance, the hard privacy categories (i.e. Linkability, Identifiability, Non-repudiation, and Detectability) do not take into account any process-specific threats as experience shows that they have a lower likelihood and impact. Cases where these types of threats would still have a high impact would typically benefit more from the heavyweight LINDDUN approach. For LINDDUN GO the usability aspects (including effort and ease of use) were considered a priority and therefore thoroughness was slightly reduced.

Content updates. Not only was each threat type's documentation extended and polished, we also wanted to update the content itself. Since the latest release of the LINDDUN threat trees, the GDPR has entered into force. An alignment with the data protection principles and data subject rights was therefore the focus. The unawareness category captures threats against data subject rights (e.g. transparency, erasure, access, portability). The non-compliance category captures data protection principles violations (e.g. disproportionality, unlawful processing). Although these concepts are inspired by GDPR, they are also applicable independent of legislation, as they encompass general privacy principles. Nonetheless, this alignment can create a leverage to futher adoption and integration of privacy in the software development lifecycle.

LINDDUN GO release. After several internal iterations of the structure of the cards and the content itself, internal runs of the method, and a final trial run with industry professionals (Section 5.2), LINDDUN GO was released to the public with a request for feedback [32]. Initial feedback results are discussed in Section 5.

4.2. LINDDUN GO cards

The main contribution of LINDDUN GO is its collection of threat type cards that describe potential privacy threats. This section will provide more insights in these cards, their properties and their creation.

LINDDUN GO extends and structures the threat type description of LINDDUN's threat trees [2] and documents them as threat type cards. Each of these cards will need to be analyzed with the system that is being examined in mind.

Each threat type card follows the same template, of which an example is shown in Figure 1:

- *Title.* Title of the threat type.
- *Hotspots.* The area where the threat occurs in the system. (More details on hotspots are discussed in the remainder of this section)
- *Threat source.* The origin type of the threat (i.e. organizational, external to the system, or the receiving party of the interaction).
- *Summary.* Short description of the threat type.
- *Elicitation questions.* Two questions to help determine the applicability of the threat type. The first question mainly determines whether the prerequisites are fulfilled, while the second question aids in assessing the applicability itself.
- *Examples.* Illustration(s) of the threat type.
- *Consequences/impact.* Rationale about why the threat is important.
- *Additional information.* Extends the description with supplementary information about the threat type.
- *Card identifier.* Identifier for easy reference.
- *LINDDUN category.* Highlight of the corresponding threat category in the LINDDUN acronym.

Hotspots represent affected system components. There are roughly two approaches to map threat knowledge to a system model: element-based [1] and interaction-based [5]. In the element-based approach each DFD element is considered individually, while the interaction-based approach rather considers for each flow the origin, destination of communication. As we see value in an interaction-based approach for LINDDUN [5], we wanted to find an understandable way to represent this, preferably without relying on a specific model notation.

The LINDDUN GO threat types therefore refer to hotspots rather than DFD elements or interactions. While hotspots still refer to typical DFD interactions, they can be easily mapped to other model representations as well. In additional, hotspots contain more information than a typical DFD element (e.g. data types), similar to Dhillon's dataflow diagram patterns for security threat modeling [15].

Figure 2 summarizes the hotspots that need to be considered when applying LINDDUN GO. The model used to represent the system that is being analyzed should therefore also contain notions of these building blocks. The hotspots include inbound and outbound communication flows, storage and retrieval actions to data storage, and processing operations. Each of these can be further scoped, by, for instance, focusing on communication with direct user interaction, or interactions that consist of a specific type of personal data.

These hotspots can thus already quickly indicate whether the threat type is applicable to the system-under-analysis, and, if so, to which parts particularly. This thus facilitates the elicitation process. In addition, they are a step towards *model-agnostic* threat modeling. Any model that supports the concepts represented by the hotspots can be used (including client-server view, BPMN, a white-board sketch, or even a list of processing activities).

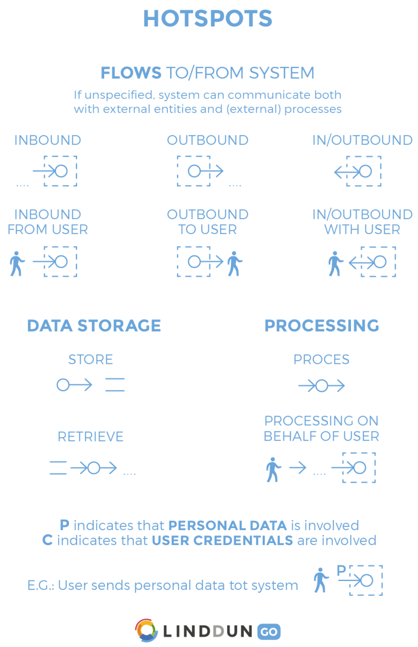
HOTSPOTS
FLOWS TO/FROM SYSTEM
If unspecified, system can communicate both with external entities and (external) processes
INBOUND
OUTBOUND
IN/OUTBOUND
INBOUND FROM USER
OUTBOUND TO USER
IN/OUTBOUND WITH USER
DATA STORAGE
STORE
RETRIEVE
PROCESSING
PROCES
PROCESSING ON BEHALF OF USER
P indicates that PERSONAL DATA is involved
C indicates that USER CREDENTIALS are involved
E.G.: User sends personal data tot system
P
LINDDUN GO

Figure 2. Summary of the hotspots used in LINDDUN GO.

Additional cards. The card deck also contains additional information. Threat category cards, for instance, summarize each LIND(D)UN category and highlight applicable hotspots. They can be used as background information, but can also be instrumental to an initial brainstorming exercise (i.e. 'freestyle' variant, as discussed in Section 4.3). Instructions on how to apply LINDDUN GO are provided, as well as information on hotspots and threat sources. Also recurring privacy terminology is summarized in a glossary, and references to related work are provided.

4.3. Using LINDDUN GO

The execution of LINDDUN GO follows a very simple and collaborative approach, as outlined in Figure 3. Every participant takes turns to draw a random threat type card from the pile and tries to elicit a corresponding threat. Afterwards all other participants fill in any missing threats that correspond to the card. When no threats can be found anymore, the next participant draws a card and the process recommences.

To elicit threats, the participants read the drawn threat type card. For each of the system components that correspond with the hotspot described on the card, they assess whether the threat is applicable. To aid this assessment, they can answer the two applicability questions on the card. When, for instance, the 'linkability of credentials' card (Figure 1) is drawn, the interactions with the system that involve users who are sending their credentials (for authentication purposes) should be examined (thereby heavily scoping the system elements that need to be considered). For each corresponding element, the applicability questions help determine whether the threat actually poses and is considered a problem.

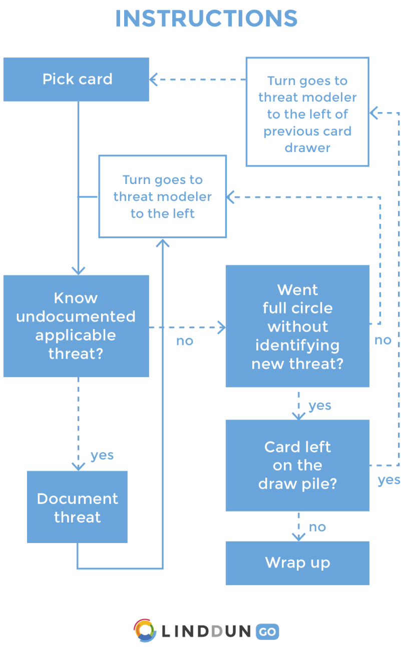
INSTRUCTIONS
Pick card
Turn goes to threat modeler to the left of previous card drawer
Turn goes to threat modeler to the left
Know undocumented applicable threat?
no
Went full circle without identifying new threat?
no
yes
yes
Document threat
Card left on the draw pile?
yes
no
Wrap up
LINDDUN GO

Figure 3. Flowchart of the LINDDUN GO process

Variants. Different analysts prefer different execution styles. Some, for instance, prefer a more playful approach to lighten the exercise. We therefore also provide some variants in addition to the main LINDDUN GO process (which is summarized above). In total, LINDDUN GO suggests five different variants to meet the broad range of stakeholder requirements: *quick* (i.e. no group iteration over each card), *time-boxed* (i.e. limit time or number of cards), *fun* (i.e. make it a competition), *solitary* (i.e. single threat modeler who uses cards as catalog), *freestyle* (i.e. only use generic category cards to brainstorm threats).

While the process thus has several variants, the cards will stay the same in each variant. Only the 'freestyle' variant, which only uses the generic threat category cards to guide a less structured brainstorming exercise, does not include the LINDDUN GO threat type cards.

5. Evaluation

In this section we first assess the requirements we set based on industry input (as described in Section 3). The remainder of this section will describe (i) the intermediary dry-runs and try-outs we executed to come to the current version of LINDDUN GO, (ii) the feedback we received from students who applied LINDDUN GO as part of a course assignment, and (iii) the feedback received from industry professionals.

5.1. Requirements assessment

As a first evaluation step, we map LINDDUN GO to the presented requirements from Section 3.

Understandable description (REQ 4). The structured description of the LINDDUN GO threat type cards provides additional insights in each potential threat, including guiding questions, examples, and potential consequences to illustrate the impact. This information increases the comprehensibility and allows people with limited privacy expertise to easily grasp the concepts.

In addition, the extensive documentation on each card also allows people with various background (e.g., analysts, developers, data protection experts, managers, etc.) to collaborate. This appears particularly useful as these stakeholders have different, yet complementary insights in the system.

Applicability criteria (REQ 5). LINDDUN GO cards contain several elements that guide the selection of applicable threats: The applicability assessment questions aid the analyst to quickly decide whether the threat type is applicable in the system-under-analysis. The hotspots also help reduce the system components the analyst needs to consider. Even the examples and consequences information can be instrumental to further determine the applicability of the threat.

Simple process (REQ 1). The overall process stripped from documentation steps and the content itself has been reduced to simplify the approach. In addition, the applicability criteria also make the process as a whole easier to execute. STRIPED and TRIM will likely be considered slightly more simple than LINDDUN GO though, as their process does not require a full iteration of system elements for each cards. LINDDUN GO however also provides a 'quick' variant for those who prefer this more simplified approach.

Comprehensive process (REQ 2). LINDDUN GO reduced its process to a minimal, while still including the threat modeling foundation, i.e. systematically iterating over *all* potential threats while taking into account all applicable system components. Compared to the full-fledged LINDDUN, some simplifications have been made (e.g. by defining hotspots for each threat type, only a subset of all system elements needs to be examined), however the LINDDUN GO process can be considered comprehensive as it forces iteration over the complete set of threat type cards while taking into account all applicable system elements and interactions. STRIPED and TRIM, on the other hand, can be considered less comprehensive than LINDDUN GO, given their process which does not require full iteration over all system elements.

Collaborative (REQ 3). The main LINDDUN GO method (as summarized in Section 4.3) is collaborative and involves multiple team members to jointly elicit privacy threats. LINDDUN GO includes all participants in the threat elicitation process and nudges also the less vocal people to speak up by explicitly taking turns in picking threat type cards and discussing them.

Overall. Both the overall simplified process and the structured, extensive documentation of threat cards (including applicability criteria and hotspots) are instrumental to lowering the required privacy expertise and reduce the effort needed to execute LINDDUN GO. STRIPED and TRIM also score very high on this overall low-effort requirement. In fact, when compared to LINDDUN GO, they are likely to outrank it, as their process is slightly simpler. LINDDUN GO can however be considered more comprehensive and better suited for non-experts given its understandable description and applicability criteria.

5.2. Trial run

Before releasing the card deck, we executed an internal trial run as well as a try-out that involved industry partners to assess and streamline both the content and process.

The external trial run was integrated in a knowledge sharing event with industry partners. In total, 23 people joined the session. Their background was quite diverse and included software architects, developers, data protection officers, legal researchers, and computer science researchers. After a 1-hour tutorial on the LINDDUN principles and LINDDUN GO, they split up in 5 groups and collaboratively played LINDDUN GO for 1h30min.

Afterwards, we held a round table to gather feedback and also send out a link to an online questionnaire (on which we received 8 responses). Overall, participants indicated that the cards contain the right amount of information (62.5%) and right level of detail (75%). One of the participants also pointed out that one of the biggest advantages of LINDDUN GO is its ability to invoke discussions between interdisciplinary teams. We also received suggestions for improvement in two main categories.

Content improvement suggestions. Some participants indicated that there were too many cards that needed to be evaluated. At that time, each threat category still had 8 to 12 threat type cards. This feedback triggered us to further reduced the deck. The participants also indicated that they would appreciate some guidance concerning the impact and likelihood of each threat. We therefore extended the description and included some remarks on this aspect (although it is difficult to give generic indications).

Collaboration preferred over competition. The card deck was presented to the participants as a card game (the 'fun' variant of the current LINDDUN GO). Also the documentation template included a field to keep track of the score. Although all of the groups had very lively discussions based on the cards, none of the groups felt the need to actually 'play' and compete. This observation made us decide to move away from the gamified competitive version and only include it as a variant of the method.

5.3. Feedback from students with hands-on experience

In addition to industry input, we also received student feedback via their course instructors. The students, who were already familiar with LINDDUN, were asked to apply LINDDUN GO and provide feedback afterwards. LINDDUN GO was applied by 4 groups of 3 students each. The four groups unanimously agreed that LINDDUN GO was easier to apply than the LINDDUN trees. The method was considered simple to apply and the cards were perceived as easy to understand. The students particularly valued the questions and examples in the threat type cards as they aided the selection of applicable threats. Also the collaborative aspect of LINDDUN GO was appreciated. As input for improvement, one group indicated that the hotspots concepts can be a bit confusing. This might be due to previous familiarity with DFD modeling concepts. Another group proposed to also include some solution suggestions for each threat type card.

5.4. Feedback from industry professionals

With the public release of LINDDUN GO [32], we made a questionnaire available to collect feedback. We also had a limited set of printed card decks which we distributed among interested privacy professionals with an explicit request for feedback (including a link to the questionnaire). The questionnaire consists of the ten standard System Usability Scale (SUS) questions) [33] and some additional questions to get more content-specific details (including open questions to describe advantages and disadvantages) as well as get some background information on the participants (e.g., whether they actually applied LINDDUN GO, read through it thoroughly or only briefly skimmed through it, and what their experience with privacy, security and threat modeling is).

As LINDDUN was only recently released, we have so far only received detailed feedback of 10 industry professionals (via questionnaire or personal communication). They all have experience in privacy engineering or threat modeling and most were already familiar with the LINDDUN framework, making their input highly valuable. The professionals who had prior experience with LINDDUN all indicated that LINDDUN GO simplifies the process with respect to LINDDUN and that the cards make the content less heavy and easier to understand; making it therefore a useful aid for people who are new to the field of privacy threat modeling.

Advantages of LINDDUN GO. Detailed comments highlighted, for instance, that the card structure was considered clear and well-explained, and the extensive description of threats is very much appreciated. The use of hotspots is considered a great idea by some, as they state that it resembles the STRIDE approach. Overall, it is considered a good initiative with easy-to-use cards.

Suggestions for LINDDUN GO. As this is only the first publicly released version of LINDDUN GO, we are mostly interested in feedback to further improve the toolkit. We already received the following suggestions: (i) More support to properly documentation the identified threats would be useful. (ii) There are also some concerns regarding the symbols used to describe the hotspots, as well as the use of hotspots in general. (iii) It would also be appreciated if even more information about the risk of each threat type could be provided (i.e. a sense of criticality or urgency). (iv) It seems to be less suited for analysis of complex microsystems.

Further evaluation. A more comprehensive evaluation in a more controlled environment (as part of a student course, or, preferably, with a group of industry professionals) could provide more detailed insights on usability, productivity gain, target audience, etc. This will be part of future work.

6. Discussion

This section discusses additional LINDDUN GO properties and describes future work to further integrate LINDDUN GO in the LINDDUN framework.

Education. While the main purpose of the LINDDUN GO card deck is supporting the threat modeling process, it can also be instrumental as educational tool. Especially as the provided knowledge is understandable by an audience with limited or no privacy background, it can also be used as aid to raise privacy awareness and learn about potential threats.

Gamification. A possible way to make a method more lightweight is to gamify it. Our experience shows that opinions about gamification seem to strongly differ however. While some see great value in a playful, competitive approach, others seem to dislike the concept of serious games in a work-context. In addition, by gamifying threat modeling, part of the method itself can get lost, which can potentially reduce the systematicity and completeness associated with it. Although these opposing opinions make an interesting future research track, they are considered out of scope for this paper. We therefore only included a gamified variant of the method as 'fun' variant, rather than imposing this as the main method. If demand is high, future work can extend the gaming aspects of the method.

LINDDUN integration. LINDDUN GO largely follows the same process as LINDDUN and the threat type knowledge reflects the content of the original threat trees.

Future work will therefore integrate the updated LINDDUN GO content and the received feedback from industry in the LINDDUN framework. For instance, an overarching LINDDUN knowledge base [16] will allow the extraction of outputs scoped specifically for different target audiences, domain contexts and activities. Also further empirical studies are required to capture feedback on LINDDUN GO when applied in practice. This will provide useful input to further improve and extend LINDDUN and LINDDUN GO to facilitate further adoption by industry.

7. Conclusion

Threat modeling approaches, such as LINDDUN, provide systematic support for privacy threat assessment. They however require extensive privacy expertise and threat modeling experience and therefore demand for a relatively high startup cost. There is thus an industry demand for more lightweight approaches.

In this paper, we (i) proposed a set of requirements for lightweight threat modeling based on industry input and empirical studies on threat modeling; (ii) presented LINDDUN GO, a toolkit that supports lightweight privacy threat modeling, based on the LINDDUN privacy threat modeling framework; (iii) describe the initial feedback and preliminary qualitative evaluation of LINDDUN GO with respect to the proposed requirements.

Initial feedback from industry experts is positive and describes LINDDUN GO as a useful tool to get started with privacy threat modeling. Light-weight approaches, such as LINDDUN GO, that allow non-experts to execute a privacy assessment exercise with limited effort can be highly instrumental to the industry adoption of the privacy-by-design paradigm.

Acknowledgments

We would like to thank the industry experts who have (anonymously) provided feedback on the LINDDUN GO card deck. We are especially grateful to Ala'a Al-Momani and Christoph Bosch (ULM University) for having their students apply LINDDUN GO in one of their courses.

References

[1] M. Deng, K. Wuyts, R. Scandariato, B. Preneel, and W. Joosen, "A privacy threat analysis framework: supporting the elicitation and fulfillment of privacy requirements," *Requirements Engineering*, 2011.

[2] K. Wuyts, "Privacy Threats in Software Architectures," Ph.D. dissertation, KU Leuven, jan 2015.

[3] M. Howard and S. Lipner, *The Security Development Lifecycle*. Microsoft, 2006.

[4] A. Shostack, *Threat Modeling: Designing for Security*. Wiley, 2014.

[5] L. Sion, K. Wuyts, K. Yskout, D. Van Landuyt, and W. Joosen, "Interaction-based privacy threat elicitation," in *International Workshop on Privacy Engineering*, 2018.

[6] L. Kohnfelder and P. Garg, "The threats to our products," 1999.

[7] T. UcedaVelez and M. M. Morana, *Risk Centric Threat Modeling: Process for Attack Simulation and Threat Analysis*. Wiley, 2015.

[8] TRIKE tools - octotrike. [Online]. Available: http://www.octotrike.org/

[9] T. DeMarco, *Structured Analysis and System Specification*. Yourdon Press, 1979.

[10] Y.-S. Martin and A. Kung, "Methods and tools for gdpr compliance through privacy and data protection engineering," in *2018 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*. IEEE, 2018, pp. 108–111.

[11] A. Crespo Garcia, N. N. McDonnell, C. Troncoso, D. Le Metayer, I. Kroener, D. Wright, J. M. del Alamo, and Y. S. Martin, "PRIPARE Privacy- and Security-by-Design Methodology Handbook." [Online]. Available: http://www.trialog.com/wp-content/uploads/2018/02/PRIPARE-Methodology-Handbook-Final-Feb-24-2016.pdf

[12] European Data Protection Supervisor, "Preliminary opinion on privacy by design (opinion 5/2018)," Tech. Rep., 2018.

[13] G. Danezis, J. Domingo-Ferrer, M. Hansen, J.-H. Hoepman, D. Le Metayer, R. Tirtea, and S. Schiffner, *Privacy and Data Protection by Design - from policy to engineering*. ENISA, 2015.

[14] "ISO/IEC TR 27550: Information technology — Security techniques — Privacy engineering for system life cycle processes," Tech. Rep., 2019.

[15] D. Dhillon, "Developer-driven threat modeling: Lessons learned in the trenches," *IEEE Security & Privacy*, vol. 9, no. 4, 2011.

[16] K. Wuyts, L. Sion, D. Van Landuyt, and W. Joosen, "Knowledge is power: Systematic reuse of privacy knowledge for threat elicitation," in *International Workshop on Privacy Engineering*, 2019, pp. 80–83.

[17] Microsoft Corporation, "Microsoft threat modeling tool 2016," 2016. [Online]. Available: https://www.microsoft.com/en-us/download/details.aspx?id=49168

[18] "IriusRisk Threat modeling tool." [Online]. Available: https://iriusrisk.com/threat-modeling-tool/

[19] "OWASP threat dragon." [Online]. Available: https://owasp.org/www-project-threat-dragon/

[20] L. Sion, D. Van Landuyt, K. Yskout, and W. Joosen, "SPARTA: Security & Privacy Architecture Through Risk-Driven Threat Assessment," in *International Conference on Software Architecture*. IEEE, 8 2018, pp. 89–92.

[21] R. Scandariato, K. Wuyts, and W. Joosen, "A descriptive study of microsoft's threat modeling technique," *Requirements Engineering*, vol. 20, no. 2, pp. 1–18, 12 2013.

[22] K. Yskout, K. Wuyts, D. Van Landuyt, R. Scandariato, and W. Joosen, "Empirical research on security and privacy by design: What (not) to expect as a researcher or a reviewer." CRC Press, 2017, pp. 1–46.

[23] K. Wuyts, R. Scandariato, and W. Joosen, "Empirical evaluation of a privacy-focused threat modeling methodology," *The Journal of Systems and Software*, vol. 96, pp. 122–138, 10 2014.

[24] D. Van Landuyt and W. Joosen, "A descriptive study of assumptions made in linddun privacy threat elicitation," in *Proceedings of the 35rd Annual ACM Symposium on Applied Computing*. ACM, 2020.

[25] A. Shostack, "Experiences threat modeling at Microsoft," in *Modeling Security Workshop. Dept. of Computing, Lancaster University, UK*, 2008.

[26] ——, "Elevation of privilege: Drawing developers into threat modeling," in *2014 {USENIX} Summit on Gaming, Games, and Gamification in Security Education (3GSE 14)*, 2014.

[27] F-Secure, "Elevation of Privacy - Elevation of Privilege privacy extension (TRIM)." [Online]. Available: https://github.com/F-Secure/elevation-of-privacy

[28] Logmein, "STRIPED – Elevation of Privilege privacy extension." [Online]. Available: http://bit.ly/2COzntD

[29] J. Cleland-Huang, T. Denning, T. Kohno, F. Shull, and S. Weber, "Keeping ahead of our adversaries," *IEEE Software*, 2016.

[30] "OWASP Cornucopia." [Online]. Available: https://owasp.org/www-project-cornucopia/

[31] L. Williams, M. Gegick, and A. Meneely, "Protection poker: Structuring software security risk assessment and knowledge transfer," in *International Symposium on Engineering Secure Software and Systems*. Springer, 2009, pp. 122–134.

[32] "LINDDUN website." [Online]. Available: www.linddun.org

[33] J. Brooke, "SUS-A quick and dirty usability scale," *Usability evaluation in industry*, vol. 189, no. 194, pp. 4–7, 1996.