Leveraging the domain experts: specializing privacy threat knowledge

Laurens Sion[0000−0002−8126−4491], Dimitri Van Landuyt[0000−0001−6597−2271], and Wouter Joosen[0000−0002−7710−5092]

DistriNet, KU Leuven, 3001 Leuven, Belgium
firstname.lastname@kuleuven.be

Abstract. The design and development of privacy-preserving software systems remains a challenging endeavor, especially with the wide-spread adoption of potentially privacy-harmful technologies such as ML/AI, LLMs, telemetry, etc. Current privacy threat knowledge consolidation efforts mainly focus on the ontological generalization of threat knowledge. The generic encoding of privacy threat knowledge is useful for increasing overall awareness of the diversity and scope of privacy threats and promoting broader application of privacy threat analysis. However, it also inhibits reuse of threat knowledge that is more tailored to the organization context or application domain. There is thus an emerging need to encode, manage, and share specialized privacy threat knowledge that may be more domain-, technology-, or organization-specific.

In this position paper, we outline a vision and roadmap towards improved support for the overall management of privacy threat knowledge, and particularly we envision advanced knowledge modeling support for capturing specialized threat knowledge, supporting evolution, customization, and reuse.

1 Introduction

Engineering privacy-preserving software-intensive systems remains a non-trivial task that requires substantial expertise to assess and strengthen the privacy properties of the system under development. Furthermore, several legislative initiatives (such as the GDPR) stress the importance of considering privacy and data protection in the design and development of systems, and even impose a strong obligation to proactively consider these issues [1].

To assist privacy engineers and developers in this task, diverse approaches such as privacy threat modeling [13], and supporting resources, such as threat trees [38], and other taxonomies [4, 8, 40] have been created. These resources provide support in two complementary dimensions. First, they provide methodological support [17] for the users to perform the actual privacy analysis of the system under design (process dimension). Second, they provide consolidated and refined threat knowledge to assist users in considering diverse threats in the system they are working on (knowledge dimension).

Even with these supporting resources, privacy experts with relevant domain- or application-specific knowledge remain a scarce resource [44]. Existing efforts to capture relevant knowledge are generally focused on abstracting specific threats (or weaknesses) for the purpose of making the resulting information more broadly available and applicable across applications and domains (Section 2). While such a generalization is definitely useful to enable wider re-use, it also tends to abstract away important details that are no longer available to the user [42].

This position paper argues for and envisions new mechanisms for specialization of privacy threat knowledge that can provide compelling benefits on three fronts. (1) Explicitly capturing specialized privacy threat knowledge enables the construction of organization-specific repositories that capture and propagate earlier experiences and best practices; this allows organizations to emphasize or prioritize specific types of threats that have to be explicitly considered across its products. (2) Specializing the threat knowledge can be performed to accommodate the particularities of specific application domains. (3) A consolidated form of privacy threat knowledge can contribute to the overall practice of cyber threat intelligence (CTI) sharing which entails that diverse players active within the same domain or ecosystem (e.g., supply chain) more readily and freely share information about incidents, vulnerabilities, and encountered attack patterns.

The specialization of threat knowledge enables the construction of, for example, specialized threat libraries. However, the threat modeler not only has to interpret this threat knowledge correctly, but also has to translate that to the relevant application domain. This additional translation step further complicates the use of this knowledge. The specialization of privacy threat knowledge to specific domains can already include this translation step and can enrich the threat knowledge with concrete domain-specific examples to make it much more convenient to use.

2 Related work

We discuss the availability of different types of threat knowledge resources and the extent of their specialization. Sections 2.1 and 2.2 first elaborate on security and privacy threat knowledge resources. Next, Section 2.3 assesses the support in current threat modeling tools. Finally, Section 2.4 covers domain-specific and application-specific resources. Figures 1a and 1b visualize the results.[1]

2.1 Security threat knowledge resources

Security threat modeling approaches offer several threat trees [18,38] that capture generic threat knowledge along the STRIDE threat type mnemonic. These trees capture generic ways in which the STRIDE threats can manifest themselves to help threat modelers to instantiate threats.

[1] Note that the exact coordinates are not important; the figures serve to draw attention to the quadrant in which the resources are located (i.e. abstract and generic, abstract and application specific, generic and concrete, or application-specific and concrete).

In addition to those resources, there are more generic resources that capture security knowledge. Examples of these are the Common Attack Pattern Enumeration and Classification (CAPEC) [22] that provides attack patterns, Common Weaknesses and Exposures (CWE) [24] for identifying the underlying causes of vulnerabilities [23], the Common Architectural Weaknesses and Exposures (CAWE) [32] that abstracts these to common architectural design flaws, and finally the Top 10 Secure Design Flaws [5] again focusing on more high-level flaws.

Overall, these resources focus on abstracting, for example the CWE catalog is based on an extensive effort to abstract and generalize specific vulnerability reports (CVE) to construct a model of the root causes. They do not provide domain-specific threat knowledge and lack support for refinement. Section 2.4 provides more details on these.

2.2 Privacy threat knowledge resources

In the space of privacy threat knowledge, there are fewer resources. We extend Wuyts' earlier overview [48] with resources that have been published since [7, 9, 14, 47]. One of the main resources in this area are the LINDDUN threat trees that have undergone several iterations [13, 14, 46] with more detail and refinements in threat examples. While these enrichments increase the level of concretization in support of applying the knowledge, LINDDUN itself remains generic in the sense that it is application domain-agnostic, in analogy with STRIDE. A variant of the LINDDUN knowledge is LINDDUN GO [47] which makes the examples more concrete, but again is not domain-specific. The LINDDUN GO-inspired Plot4AI cards [7], on the contrary, focuses specifically on the domain of AI. A recent newly-introduced knowledge source is the MITRE's PANOPTIC [36]. PANOPTIC uses two taxonomies: privacy contextual domains and privacy activities, and is constructed from knowledge about FTC/FCC privacy attacks. Finally, CNIL's methodology for privacy risk [11] also includes a generic list of threats. While the methodology focuses on privacy risk, the provided threats are generic and focus on security (confidentiality, integrity and availability). In other work, a technique for making specific domain refinements of LINDDUN threat knowledge [49] mainly involved tagging or annotating specific branches of the threat trees with domain information. This mechanism is predominantly suited to express domain-specific *selections* from the broader knowledge, i.e. to indicate that some threat types or sub-types are (or not) applicable or relevant to the specific application domain. It however does not support true *specialization* in the sense that newer subtypes (or leading examples) can be introduced of the generic threat types encoded in the threat library.

2.3 Threat modeling tools

In additional to the generally-available resources on security and privacy threat knowledge, several threat modeling tools also embed threat knowledge and may allow different forms of customization or specialization of this knowledge. This section outlines the extent of support in existing tools. Shi et al. [37]

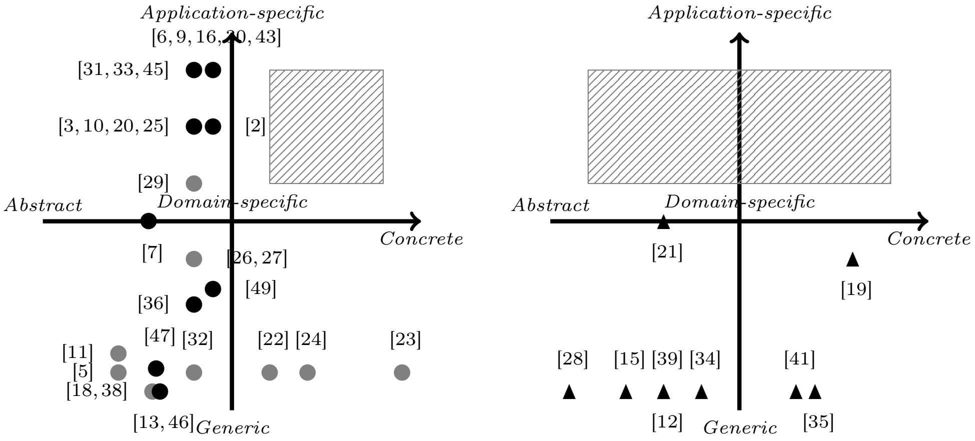

(a) Overview of knowledge resources
This diagram situates security (●) and privacy (●) knowledge resources on a high-level. Most resources provide generic and abstract knowledge (high-level guidance). Furthermore, most domain- or application-specific resources are research results not readily available for practitioners to apply. The marked area highlights the gap.

(b) Overview of tool support
This diagram situates threat modeling tools. Most provide generic and abstract knowledge (high-level guidance). This confirms the observations from Figure 1a in which the domain- and application-specific resources are not available in tools. The marked area highlights the gap.

Fig. 1: Overview of knowledge and tool support

analyzed several open source and freely available commercial tools. They observe tools to either have a self-defined library (based on STRIDE/LINDDUN or fully custom) or leverage existing resources (CVE, CWE, or CAPEC). Microsoft Threat Modeling Tool [21], ThreatDragon [28], pyTM [41], OVVL [34], threagile [35], and IriusRisk [19] all rely on such generic sources. Other tools such as threats manager studio [12] and SPARTA [39] provide similar support for automated elicitation of threats based on similar catalogs of threat types. Finally, CAIRIS [15] also supports DFD creation and documenting threats. IriusRisk and Microsoft's Threat Modeling Tool offer additional support for AWS and Azure, with the Microsoft Threat Modeling Tool even having a threat catalog for medical devices.

2.4 Domain- and application-specific resources

In addition to the generic resources outlined in Sections 2.1 and 2.2, there are several resources that offer more scoped and concrete threat knowledge. OWASP publishes several top 10 security risk rankings for mobile [27], web [26], and LLM applications [29]. While these resources focus on security, taxonomies for website privacy vulnerabilities are also available [4]. Listing all potential domain-specific and application-specific resources is not possible. Instead, a few domains are highlighted to illustrate the where these types of knowledge resources can be situated in Figure 1a. In this case, we highlight the areas of machine learning [3, 10, 20, 25], automotive [6, 9, 16, 30, 43], IoT [2], and social networks [31, 33, 45]. These examples illustrate the range between more generic technologies and highly-specific application domains. However, this specialized knowledge can be mainly found in papers, not in reusable catalogs or tools.

2.5 Observations

The following three observations can be made from the presentation of the different security and privacy threat knowledge resources and threat modeling tool support visualized in, respectively, Figures 1a and 1b.

First, construction of security and privacy threat knowledge happens mainly through abstraction of concrete vulnerabilities, threats, attacks, etc. to create more generally-applicable resources (CWE, top 10s, taxonomies, etc.).

Second, this pattern re-appears in tool support. Several tools [19, 21] make steps towards providing more specific knowledge, for example, for cloud platforms (e.g., AWS or Azure) or specific domains (e.g., medical devices).

Third, a lot of specialized knowledge is available in particular domains (illustrated through the overview in the areas of machine learning, automotive, IoT, and social networks) but this knowledge is not readily available for practitioners.

This is also confirmed by a mapping study of privacy threat analysis assumptions [42] that has pointed out that a number of assumptions can be attributed to a mismatch between the generic information in the resources and the specific domain, resulting in generic and redundant assumptions (e.g., *threat type X does not apply in our domain*), illustrating the issue when specialization falls short.

3 Roadmap towards threat knowledge specialization

To bridge the highlighted gap (Figures 1a and 1b), we argue that more extensive support is needed for domain- and application-specific knowledge. This requires modeling support to capture these domain-specific details in a single threat knowledge representation. In such a representation, domain- or application-specific variants could be realized as filtered views of that threat knowledge [49]. This requires support for encoding information at different abstraction levels (from high-level domain-specific threat types to very concrete example threats). Realizing this requires overcoming a number of key challenges.

Structured and extensible knowledge representation. The first challenge is expanding the modeling support in existing threat libraries to support the representation of domain- and application-specific threat types, characteristics,annotation, and examples. This is necessary for: (1) capturing the application- and domain-specific knowledge at various abstraction levels, and (2) annotating this knowledge with the relevant domain or application to enable filtering and selection support.

Specialization and refinement. Here, the aforementioned mechanisms are leveraged to expand existing threat libraries with more specialized types, characteristics, and examples. To create these resources, existing domain- and application-specific resources (Section 2.4) can be encoded. This provides an opportunity to verify whether existing generic resources are expressive enough to cover them and extends the threat library with more concrete information.

Evolution and versioning. To support continuous evolution and extension (when adding new domain- and application-specific knowledge) requires support for evolution and versioning. Modifications to the existing knowledge representation may be necessary in the future to express new types of knowledge that cannot be accurately captured with previous versions.

Reuse and customization. Specialized privacy threat knowledge bases provide the opportunity for organizations to customize the knowledge for internal reuse.

Community interaction. Finally, the model-driven representation of specialized threat knowledge provides a structured format for exchanging this type of information. This raises opportunities for community interactions and collaborations to pro-actively share their knowledge, examples, and potential improvements.

We presented our vision towards more principled support for encoding domain-, technology-, and organization-specific threat knowledge to address this limitation. The specialization of the threat knowledge offers compelling benefits in: (1) helping threat modelers to apply that knowledge in their applications, (2) encoding and reusing organization-specific threat knowledge in a more systematic fashion, (3) increasing efficiency in eliminating and filtering out irrelevant threat knowledge, and (4) supporting and encouraging inter-organizational sharing of privacy threat knowledge. More structured support for specialized privacy threat knowledge creation will support more efficient encoding and knowledge sharing and help organizations to develop privacy-preserving software systems by comprehensively and systematically assessing the privacy threats.

Acknowledgments. This research is partially funded by the Research Fund KU Leuven and Cybersecurity Research Program Flanders.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016. Official Journal of the European Union **59**(L 119), 1–88 (May 2016)
2. Abdulghani, H.A., Nijdam, N.A., Collen, A., Konstantas, D.: A study on security and privacy guidelines, countermeasures, threats: Iot data at rest perspective. Symmetry **11**(6), 774 (2019)
3. Al-Rubaie, M., Chang, J.M.: Privacy-preserving machine learning: Threats and solutions. IEEE Security & Privacy **17**(2), 49–58 (2019)
4. Antón, A.I., Earp, J.B.: A requirements taxonomy for reducing web site privacy vulnerabilities. Requirements engineering **9**, 169–185 (2004)

5. Arce, I., Daswani, N., Delgrosso, J., Dhillon, D., Kern, C., Kohno, T., Landwehr, C., Mcgraw, G., Schoenfield, B., Seltzer, M., Spinellis, D., Tarandach, I., West, J.: Avoiding the Top 10 Software Security Design Flaws. Tech. rep., IEEE Center for Secure Design (2014)
6. Asuquo, P., Cruickshank, H., Morley, J., Ogah, C.P.A., Lei, A., Hathal, W., Bao, S., Sun, Z.: Security and privacy in location-based services for vehicular and mobile communications: An overview, challenges, and countermeasures. IEEE Internet of Things Journal **5**(6), 4778–4802 (2018)
7. Barberá, I.: PLOT4ai - Privacy Library Of Threats 4 Artificial Intelligence. `https://plot4.ai/` (Feb 2024), `https://plot4.ai/`
8. Barker, K., Askari, M., Banerjee, M., Ghazinour, K., Mackas, B., Majedi, M., Pun, S., Williams, A.: A data privacy taxonomy. In: Dataspace: The Final Frontier: 26th British National Conference on Databases (2009)
9. Chah, B., Lombard, A., Bkakria, A., Yaich, R., Abbas-Turki, A., Galland, S.: Privacy threat analysis for connected and autonomous vehicles. Procedia Computer Science **210**, 36–44 (2022)
10. Chang, S., Li, C.: Privacy in neural network learning: threats and countermeasures. IEEE Network **32**(4), 61–67 (2018)
11. CNIL: Methodology for Privacy Risk Managment: How to implement the Data Protection Act. Tech. rep. (2012)
12. Curzi, S.: Pytm (2024), `https://threatsmanager.com/`
13. Deng, M., Wuyts, K., Scandariato, R., Preneel, B., Joosen, W.: A privacy threat analysis framework: Supporting the elicitation and fulfillment of privacy requirements. Requirements Engineering **16**(1), 3–32 (2011)
14. DistriNet: LINDDUN Website. `https://linddun.org` (2024)
15. Faily, S.: Designing usable and secure software with IRIS and CAIRIS. Springer (2018)
16. den Hartog, J., Zannone, N., et al.: Security and privacy for innovative automotive applications: A survey. Computer Communications **132**, 17–41 (2018)
17. Hernan, S., Lambert, S., Ostwald, T., Shostack, A.: Threat modeling: Uncover security design flaws using the STRIDE approach. MSDN Magazine **6** (Nov 2006), `https://msdn.microsoft.com/en-us/magazine/cc163519.aspx`
18. Howard, M., Lipner, S.: The Security Development Lifecycle. Microsoft Press (2006)
19. IriusRisk: IriusRisk (2024), `https://www.iriusrisk.com/`
20. Lyu, L., Yu, H., Zhao, J., Yang, Q.: Threats to federated learning. Federated Learning: Privacy and Incentive pp. 3–16 (2020)
21. Microsoft Corporation: Microsoft Threat Modeling Tool 7. http://aka.ms/tmt (2023)
22. MITRE: Common Attack Pattern Enumeration and Classification (2024), `https://capec.mitre.org`
23. MITRE: Common Vulnerability Enumeration (2024), `https://cve.mitre.org`
24. MITRE: Common Weakness Enumeration (2024), `https://cwe.mitre.org`
25. Mo, K., Ye, P., Ren, X., Wang, S., Li, W., Li, J.: Security and privacy issues in deep reinforcement learning: Threats and countermeasures. ACM Computing Surveys (2024)
26. OWASP: Top 10 Web Application Security Risks (2021), `https://owasp.org/www-project-top-ten/`
27. OWASP: Mobile Top 10 (2024), `https://owasp.org/www-project-mobile-top-10/`
28. OWASP: Threat Dragon. https://owasp.org/www-project-threat-dragon/ (2024)

29. OWASP: Top 10 for Large Language Model Applications version 1.1 (2024), `https://owasp.org/www-project-top-10-for-large-language-model-applications/`
30. Raciti, M., Bella, G.: How to model privacy threats in the automotive domain. arXiv preprint arXiv:2303.10370 (2023)
31. Rodrigues, A., Villela, M.L.B., Feitosa, E.L.: Privacy threat modeling language. IEEE Access **11**, 24448–24471 (2023)
32. Santos, J.C., Tarrit, K., Mirakhorli, M.: A catalog of security architecture weaknesses. In: 2017 IEEE International Conference on Software Architecture Workshops (ICSAW). pp. 220–223. IEEE (2017)
33. Sanz, B., Laorden, C., Alvarez, G., Bringas, P.G.: A threat model approach to attacks and countermeasures in on-line social networks. In Proceedings of the 11th Reunion Espanola de Criptografıa y Seguridad de la Información (RECSI) (2010)
34. Schaad, A., Reski, T.: "Open Weakness and Vulnerability Modeler" (OVVL): An Updated Approach to Threat Modeling. In: Proceedings of the 16th International Joint Conference on e-Business and Telecommunications. SciTePress (2019)
35. Schneider, C.: Threagile (2024), `https://threagile.io/`
36. Shapiro, S., Bloom, C., Ballard, B., Slotter, S., Paes, M., McEwen, J., Xu, R., Katcher, S.: The PANOPTIC™ Privacy Threat Model. Tech. Rep. V1 (Dec 2023)
37. Shi, Z., Graffi, K., Starobinski, D., Matyunin, N.: Threat modeling tools: A taxonomy. IEEE Security & Privacy **20**(4), 29–39 (2022)
38. Shostack, A.: Threat Modeling: Designing for Security. John Wiley & Sons, Indianapolis, Indiana (2014)
39. Sion, L., Van Landuyt, D., Yskout, K., Joosen, W.: Sparta: Security & privacy architecture through risk-driven threat assessment. In: 2018 IEEE International Conference on Software Architecture Companion (ICSA-C). pp. 89–92 (2018)
40. Solove, D.J.: A taxonomy of privacy. U. Pa. l. Rev. **154**, 477 (2005)
41. Tarandach, I.: Pytm (2024), `https://github.com/izar/pytm`
42. Van Landuyt, D., Joosen, W.: A descriptive study of assumptions made in linddun privacy threat elicitation. In: Proceedings of the 35th Annual ACM Symposium on Applied Computing. pp. 1280–1287 (2020)
43. Vasenev, A., Stahl, F., Hamazaryan, H., Ma, Z., Shan, L., Kemmerich, J., Loiseaux, C.: Practical security and privacy threat analysis in the automotive domain: Long term support scenario for over-the-air updates. In: VEHITS (2019)
44. Verreydt, S., Yskout, K., Sion, L., Joosen, W.: Threat modeling state of practice in dutch organizations. In: SOUPS'24: Proceedings of the Twentieth USENIX Conference on Usable Privacy and Security (2024)
45. Wang, Y., Nepali, R.K.: Privacy threat modeling framework for online social networks. In: 2015 International Conference on Collaboration Technologies and Systems (CTS). pp. 358–363. IEEE (2015)
46. Wuyts, K.: Privacy Threats in Software Architectures. Ph.D. thesis, KU Leuven (Jan 2015)
47. Wuyts, K., Sion, L., Joosen, W.: LINDDUN GO: A Lightweight Approach to Privacy Threat Modeling. In: 2020 IEEE Security and Privacy Workshops (2020)
48. Wuyts, K., Sion, L., Van Landuyt, D., Joosen, W.: Knowledge is power: Systematic reuse of privacy knowledge for threat elicitation. In: 2019 IEEE Security and Privacy Workshops (SPW). pp. 80–83 (2019). `https://doi.org/10.1109/SPW.2019.00025`
49. Wuyts, K., Van Landuyt, D., Hovsepyan, A., Joosen, W.: Effective and efficient privacy threat modeling through domain refinements. In: Proceedings of the 33rd Annual ACM Symposium on Applied Computing. pp. 1175–1178 (2018)