Themes in Privacy Threat Modeling

BEN BALLARD, MARK PAES, and SHELBY SLOTTER, MITRE
RYAN XU, MITRE
JULIE MCEWEN, MITRE
STUART SHAPIRO, MITRE
CARA BLOOM*, MITRE

Privacy threat modeling is an extension of cyber threat modeling, and threat modeling in general, that seeks to model the actions and adversaries that lead to privacy harms. When integrated with vulnerability and consequence models, it enables threat-informed risk modeling. On its own, a privacy threat model can be used for threat assessments, red teaming, and other activities. A sufficiently granular threat model could provide a common lexicon to accelerate conversations in the privacy community. At the first Privacy Threat Modeling Workshop, held at the Symposium on Usable Privacy and Security, a group of academics and practitioners gathered to discuss the concept of privacy threat models and two existing models: LINDDUN and MITRE PANOPTIC™. Themes from this workshop about conceptualizing, designing, and operationalizing privacy threat models are discussed in this paper.

Additional Key Words and Phrases: privacy, privacy threat modeling, privacy risk modeling, threat modeling, risk modeling

1 INTRODUCTION

Privacy threats are actors and actions that can lead to privacy harms. These threats exploit privacy vulnerabilities, which lead to adverse consequences for organizations and individuals. Privacy threat modeling enables the identification of privacy threats so that organizations can target the vulnerabilities that are most likely to be compromised by threat actors. Privacy risk modeling includes not only privacy threat modeling, but vulnerability and consequence modeling as well. While there is a connection between privacy risk modeling and cybersecurity risk modeling, not all privacy risks are cybersecurity threats and vice versa. The overlap between cybersecurity and privacy is generally characterized as confidentiality attacks involving personal data. Because these attacks overlap with cybersecurity, they are included in cybersecurity risk and threat models. However, the non-cybersecurity privacy risks are understudied in comparison.

Privacy risk modeling components are somewhat developed in the vulnerabilities and consequences spaces, with some well-known models available for each. For example, Solove's Taxonomy of Privacy [11] and Calo's subjective/objective privacy harms [4] cover privacy consequences and the NIST Privacy Risk Assessment Methodology Catalog of Problematic Data Actions and Problems [1] categorizes both vulnerabilities and consequences. The privacy threat space, in contrast, is largely underexplored, with only one model currently available - LINDDUN.

A holistic approach to representing privacy threats could inform privacy risk models and provide a common lexicon to accelerate conversations in the privacy community. Organizations can improve their privacy risk modeling by making it threat informed. They can use knowledge of privacy threats gained from privacy threat models to move away from inadequate compliance-focused privacy programs to those that manage the risks and constituent threats to individuals

*Corresponding author

Authors' addresses: Ben Ballard, bballard@mitre.org; Mark Paes, mpaes@mitre.org; Shelby Slotter, sslotter@mitre.org, MITRE; Ryan Xu, rxu@mitre.org, MITRE; Julie McEwen, jmcewen@mitre.org, MITRE; Stuart Shapiro, sshapiro@mitre.org, MITRE; Cara Bloom, cbloom@mitre.org, MITRE.

and organizations. Using privacy threat models can help them address privacy threats when developing or acquiring systems, architectures, and data-handling approaches, leading to reduced privacy attack surfaces.

This paper discusses themes in privacy threat modeling. It reflects discussion from the Workshop on Privacy Threat Modeling held on August 7, 2022 at the Symposium on Usable Privacy and Security (SOUPS). The workshop brought together researchers, practitioners, and industry specialists to collaborate on the topic of privacy threats.

The workshop examined how the community defines a privacy threat, incident, breach, or attack, and the bounds on each term. It explored better ways of creating datasets of privacy threats which can be used to generate threat models and better understand the privacy threat environment. It fostered discussion of methods for categorizing and describing privacy threats using taxonomies and ontological structures, including current examples, as well as research and implementation challenges in the space. More specifically, attendees discussed issues related to the development and operationalization of privacy threat taxonomies.

This paper is organized around two central themes: designing privacy threat models (PTMs) and operationalizing PTMs. Topics that are discussed regarding designing PTMs include addressing the context and scope of a model and performing iterative approaches to privacy threat modeling. Addressing compliance pitfalls and the cultural context of compliance are discussed with respect to operationalizing privacy threat models.

2 THREAT MODELING

From a high-level, the conceptual purpose of a threat model is to audit and assess factors that are most likely to threaten a system/process [17]. Applying this to complex systems that work with data, it is unlikely that one would be able to surmise every possible threat posed, nor efficient to do so. Because of this, it is key to direct focus towards those threats that can be seen as having a higher potential to occur, as well as causing harms of greater concern [10].

Privacy threat models are an extension of the well-established practice of threat modeling. The discipline spans many fields and has been defined differently based upon the context of its use [17]. While no definition is all encompassing, Uzunov and Fernandez offer a widely accepted definition, describing threat modeling as “a process that can be used to analyze potential attacks or threats, and can be supported by threat libraries or attack taxonomies” [14]. Whether this process is manual or automated, formal or informal, threat modeling applies structure to what was previously an abstract array of threats. By doing so, one can begin the process of identifying what threats are relevant and should be prioritized.

2.1 Cyber Threat Modeling

Cyber threat modeling is a subset of threat modeling more generally. It carries with it its own challenges and unique characteristics that distinguish it from other frames of analysis [10]. Ultimately, cyber threat modeling is brainstorming threats [7]. In doing so, an analyst is trying to understand attackers (e.g., hackers): what types of attacks they will use, what assets they are targeting, what their goals are, what systems they’re trying to exploit, and how those systems have been developed. Any one of these questions can serve as a lens for analysis. Regardless of which approach one chooses, doing so begins to apply needed scope to one’s analysis, thereby simplifying the threat modeling process. The ultimate result of this is a foundation for effectively identifying and managing threats [13].

A wide range of tools exist to support those engaging in cyber threat modeling. One of the first cyber threat models is Microsoft’s STRIDE, which stands for spoofing, tampering, repudiation, information disclosure, denial of service, and elevation of privilege [9]. While not comprehensive, it has been used extensively since its creation. In 2011 Lockheed Martin introduced the Cyber Killchain®as an early, linear model of advanced persistent threats. It views cyber attacks

through the military lens of a "kill chain" which identifies the structure of attacks. The Cyber Killchain breaks attacks down into seven phases: reconnaissance, weaponization, delivery, exploitation, installation, command and control, and actions on objectives [5]. Frameworks like CAPEC [2] serve as repositories for threats, allowing software and hardware developers to identify relevant threats and adapt their models as new threats arise. The MITRE ATT&CK®framework [12] serves as an attack library compiling information on adversary behaviors; namely their use of specific tactics and techniques to exploit digital systems. Tools such as these can be leveraged to tie cyber threat actions and indicators to specific threat actors and provide potential mitigations. In this manner, threat modeling tools can not only be used to develop threat models and modeling methodologies, but also to improve mitigation as well [12].

2.2 Privacy Threat Modeling

While cyber threat models cover attacks on confidentiality, integrity, and availability of information systems, privacy threat models should cover attacks on informational and other types of privacy (e.g., bodily, decisional, behavioral). Because the definition of privacy can be a hotly contested issue, the scope of privacy threat models is likewise unclear. Arguments could be made that models such as STRIDE include a privacy component (information disclosure) or that privacy models should not extend beyond confidentiality issues. Koops et al. argue for eight different types of privacy [6] while NIST posits that privacy with respect to information systems has three dimensions: predictability, manageability, and disassociability [3]. Privacy also can be defined by the types of harms caused by its violation [11] or by the contextual appropriateness of information flows [8]. The definition of privacy used will inherently dictate the breadth and scope of the model. The scope of threat models for cybersecurity and privacy overlap where confidentiality of personal information is concerned.

Privacy threat modeling diverges from cyber threat modeling in distinct ways, even though the domains of cybersecurity and privacy are inherently linked. Unlike in cyber attacks, malicious intent may not be required for a privacy attack, as many privacy harms are created by entities with good intentions. Additionally, unlike cyber attacks, inaction as well as action can precipitate a privacy harm. The adversarial role is taken by parties that can cause privacy harms, for example by exploiting personal information or interfering with an individual's decision making. Finally, in cyber threat modeling the threat is almost exclusively external to the system, while privacy threat models should consider threats arising from the system itself (whether social, technical, or socio-technical). This adds further complexity to any privacy threat model.

Complexity is also introduced through the contextual nature of privacy. Nissenbaum argues in her book *Privacy in Context* that privacy is provided by appropriate contextual flows of information [8]. Applying this theory to privacy threat modeling, it is understandable that the context of a potential privacy threat is essential to understanding the threat itself. For instance, the same action taken by two different actors could be a threat or not; similarly, the same actor could perform the same action on two different pieces of information and it could be a threat or not. The context around the threat determines not only the level of the threat, but whether it is threatening at all.

Any modeling activity inevitably raises the question of when and how often modeling should occur and privacy threat modeling is no exception. There was general agreement at the workshop that privacy threat modeling should take place early in the development lifecycle, but after that questions of scalability were raised. In particular, it is unclear whether modeling updates should start from a blank state or from the prior modeling effort. Consistent contexts can enable an existing model and output to be used as a starting point and modified as necessary, with a concomitant reduction in time and effort. In contrast, dynamic threat environments will undermine attempts to reuse models and

their prior outputs, resulting in privacy risk assessments that are out of sync with the threat landscape. A crucial task in modeling is finding a satisfactory middle ground.

2.2.1 LINDDUN. An example of an existing privacy threat model that was discussed at the Workshop is the LINDDUN privacy engineering framework [16]. The framework, developed at KU Leuven, provides a methodology to identify and mitigate threats within software systems. The three main steps in the LINDDUN methodology are to model the system by breaking down the data flow processes, elicit threats through analysis of the inner workings of the system, and subsequently manage these threats through prioritization based on likelihood followed by selection of appropriate mitigation strategies and solutions [15]. The framework includes 7 threat categories (linkability, identifiability, non-repudiation, detectability, disclosure, unawareness, and non-compliance) and provides a template to map system elements to these categories, a catalog of threat trees to provide an overview of common "attack paths", a taxonomy of mitigation strategies, as well as a classification of privacy solutions such as privacy enhancing technologies (PETs) [16].

2.2.2 MITRE PANOPTIC™. The MITRE team has been developing a new PTM: MITRE PANOPTIC, the Pattern and Action Nomenclature Of Privacy Threats In Context. This taxonomy and typology of privacy threats was initially developed through the deconstruction and systematic organization of privacy cases from the US Federal Trade Commission (FTC) and Federal Communications Commission (FCC), iterating through these cases to identify recurring specific actions, leading to a set of classifications of attacks that resulted in one or more adverse privacy consequences based on Solove's taxonomy of privacy problems [11]. The resulting hierarchical taxonomy the team has produced employs a variety of threat activity categories made up of specific threat actions. With this taxonomy, the team has been mapping cases identified from an expanding set of privacy incidents, including a dataset generated from open-source analysis of news sources and social media, and is creating a threat pattern typology that generalizes frequently observed clusters of similar attacks.

3 DESIGNING PRIVACY THREAT MODELS

3.1 Theme 1: Context and Model Scope

Given that privacy is contextual, it stands to reason that privacy threat modeling must also be highly contextual. This complicates the question of what level of detail is appropriate for a privacy threat model, as the answer will vary depending on context—both the context in which a model is scoped and that in which it is employed.

3.1.1 Privacy Threat Model Refinement. While the environment in which the model is applied is likely the primary focus when considering the context of a PTM, it is important for developers of PTMs to maintain a critical awareness of their own positions in the process. The PTM refinement process will be guided by the goals, judgments, and biases of the threat modelers. The privacy threat modelers will bring an understanding of the system, the data subjects at risk, the threat actors, and what constitutes a potential privacy harm. The modelers will determine whether the model should address privacy threats to or by their system—a determination likely to change depending on what role the stakeholders fill in relation to the system. Is the PTM modeling team diverse? Is the group representative of the data subjects they are seeking to protect or the broader stakeholders' interests they represent?

These questions serve a valuable role in shaping the appropriate scope of a PTM by urging modelers to reflect on their definition of and justification for what is appropriate. Especially given that the expectations and norms surrounding privacy are contested, it is important for the stakeholders employing PTMs to critically evaluate their personal assessments of privacy threats throughout the modeling process. Routinely reviewing and validating their

assumptions may provide insights for PTM modelers into perspectives they might have otherwise overlooked or stakeholder groups whose interests need to be represented in some fashion.

3.1.2 Privacy Threat Model Application. There is a wide range of variables relating to the environment or situation a PTM is applied to that might affect the model's scope. These dimensions of threat model scope include:

(1) System: What is the nature of the system and how is that represented? This is particularly important if taking a system-centric perspective on threats in which the modeling process is driven by the functional architecture of the target system.
(2) Subject: Who are the data subjects whose privacy is potentially violated by the privacy threat? This is related to understanding the potential harms and negative outcomes from a privacy incident.
(3) Adversary: Who or what potentially poses a threat to the privacy of the system's data subjects? This requires analysis of threats to the system and—by proxy—its data subjects versus threats to data subjects by the system itself.
(4) Stakeholders: What additional stakeholders may be involved in or impacted by the privacy threat? Complex systems can involve a variety of different stakeholders. Determining which specific types or groups of stakeholders will be considered as part of the threat modeling process will serve to circumscribe the resulting model.
(5) Model boundary: How far should the threat model extend past the system boundary? There is a need to look beyond the boundaries of the system as an attack could have secondary and even tertiary impacts. The Cambridge Analytica privacy incident highlighted how a complex attack chain could involve multiple distinct systems, contexts, and populations.
(6) Other risk model components: What are the vulnerabilities and consequences? These may indirectly constrain the threat model. Threat modelers will be disinclined to incorporate threats that don't align with the types of vulnerabilities or consequences of concern if these have been defined first.
(7) Mitigation: What actionable steps can be taken to mitigate the potential harms posed by a privacy threat? If the ultimate purpose of a PTM is to prevent privacy harms, then a PTM must be granular enough to enable the identification of actionable measures to protect against those harms.

While by no means comprehensive, this list provides insight into some of the contextual elements PTM developers should consider when evaluating the environment which a model will represent. These dimensions are all variable and interrelated and, thus, they are likely to change over time and prompt changes across other dimensions.

Even with careful analysis of these different factors, attack scenarios can involve many unforeseen twists and turns. There is a real question about how to perform threat modeling in a way that allows for such scenarios. One approach, which we term "guided" threat modeling, would use a (possibly tree- or graph-like) step-by-step process for in-depth analysis of a threat's potential impact. At each step threat modelers would consider the immediate context and decide whether to proceed down further paths or to backtrack. A process like this would enable modelers to dynamically determine the model boundary.

3.2 Theme 2: Iterative Privacy Threat Modeling

Just as the scope of a PTM hinges on its context, the value of a PTM varies depending on how and when the modeling is performed throughout a system's lifecycle. It is important to approach privacy threat modeling not just as a static exercise, performed retroactively to assess an existing system, but also as a dynamic process—one that develops in stride with the system it is meant to address. Products change throughout the ideation and development processes, and their

privacy threat profiles change with them. Thus, privacy threat modeling should begin with the first conceptualization of a system architecture and continue into development, deployment, and maintenance.

As discussed before, the level of detail for a PTM will change as the context surrounding the modeling process changes, with granularity of the model likely increasing as the decisions guiding system implementation become more concrete. A system's context may change during development—with the introduction of new stakeholders, policies, modes of processing data—necessitating revisiting and possibly adapting the PTM. Employing and reevaluating PTMs in an iterative manner enables those responsible for a system to account for new threats as they arise, increasing their capability to mitigate against a shifting threat landscape. In some cases, integrating privacy threat modeling into the system planning and development process not only prepares the owners of a system to respond to privacy threats but may even enable them to preempt privacy threats by design.

4 OPERATIONALIZING PRIVACY THREAT MODELS

4.1 Theme 3: Compliance Pitfalls

The group's discussion of incorporating privacy threat modeling into existing compliance regimes revealed the difficulties of regulating privacy practices, particularly in a standardized way. Concerns arose around the idea of performative compliance, where the effectiveness of compliance measures is not as important to some stakeholders as the appearance of providing protection against privacy threats, resulting in routines of rote action akin to simply checking off boxes. Each idea posed for addressing the risk of performativity comes with its own potential benefits and drawbacks:

- *Contract a third party to conduct the threat modeling and assessment.* Outsourcing the threat modeling process to a stakeholder group more removed from an organization's regular business rhythm may decrease incentives that promote a performative approach to compliance. At the same time, doing so may simply create a new line of consulting business that similarly engenders a checklist mentality.
- *Enable users to demand the use of threat models as an exercise of user rights.* This shifts the focus from organizations to their stakeholders most affected by threat modeling practices (or a lack thereof). While encouraging users to exercise autonomy over their data and privacy rights sounds promising in theory, we question how likely users are to put this into practice. Though non-governmental organizations (NGOs) could conceivably do this as well, we may not want to rely solely on them to advocate for users' right to privacy protections.
- *Specify a common consequence model for privacy threat modeling.* To render the privacy threat modeling process more tractable, an appropriate authority could require that some or all organizations use a particular consequence model to establish context for their modeling. In addition to the difficulties inherent in determining the appropriate authority and mode of regulation, this would have the unavoidable side effect of at least implicitly limiting the consequences considered.
- *Require organizations to make their privacy threat models public.* This measure would increase transparency and public insight into what organizations are actually doing to protect user privacy. However, those public models are unlikely to be unvarnished, and the utility of such disclosures might be vitiated if other risk model components are not also available for inspection.
- *Create privacy red teams.* The practice of red teaming may make threat modeling more actionable and provide some assurance that organizations are sincerely evaluating and working to protect against potential privacy threats. Even with informed consent and data de-identification, though, the ethical nuances of privacy red

teaming—given the close ties between the data at risk and the person that data represents—may constrain an organization's ability to effectively employ this practice.

The success of any of these approaches to promoting meaningful compliance depends on the context of not only the organization incorporating privacy threat modeling into its practices but also the broader regulatory environment.

4.2 Theme 4: Cultural Context of Compliance

Due to its contextual nature, shared definitions of and expectations for privacy vary across countries, regions, and cultures. These different conceptions of privacy lead to divergent laws, policies, and practices regulating privacy rights and protections. This is evident when comparing the European Union's General Data Protection Regulation (GDPR) to the patchwork of policies governing privacy in the United States, or even when comparing privacy statutes between U.S. states. The successful integration of privacy threat modeling into compliance regimes requires understanding and accounting for these regulatory differences.

Regulatory context adds another layer to the discussion of operationalizing privacy threat models, impacting each of the compliance options outlined in the previous section. The ability of users to demand that organizations use privacy threat models is determined by how individuals' privacy rights are defined and recognized by their government. Not all governments may possess or exercise the authority to regulate privacy threat modeling, limiting the possibilities of establishing a common consequence model or mandating the publication of threat models.

For environments in which privacy regulation is prevalent, organizations must avoid narrowing the scope of their privacy threat modeling to only address highly regulated threats. The purpose of incorporating privacy threat modeling into compliance regimes is to draw an organization's awareness to the wide array of potential privacy threats, expanding their capacity to reduce privacy risk—not to simply replicate existing compliance measures by other means. Thus, while regulatory context will influence the ways in which organizations shape their compliance programs, comprehensive threat identification and assessment requires that organizations develop specific, enforceable standards of their own to guide privacy threat modeling.

REFERENCES

[1] October 25, 2022. https://www.nist.gov/itl/applied-cybersecurity/privacy-engineering/resources

[2] Sean Barnum. 2008. *Common attack pattern enumeration and classification (CAPEC) schema.* Technical Report. U.S. Department of Homeland Security.

[3] Sean Brooks, Michael Garcia, Naomi Lefkovitz, Suzanne Lightman, and Ellen Nadeau. 2017. *An introduction to privacy engineering and risk management in federal systems, NISTIR 8062.* Technical Report. US Department of Commerce, National Institute of Standards and Technology (NIST).

[4] Ryan Calo. 2011. The boundaries of privacy harm. *Indiana Law Journal* 86, 3 (2011). https://ssrn.com/abstract=1641487

[5] Eric M Hutchins, Michael J Cloppert, Rohan M Amin, et al. 2011. Intelligence-driven computer network defense informed by analysis of adversary campaigns and intrusion kill chains. *Leading Issues in Information Warfare & Security Research* 1, 1 (2011), 80.

[6] Bert-Jaap Koops, Bryce Clayton Newell, Tjerk Timan, Ivan Skorvanek, Tomislav Chokrevski, and Masa Galic. 2016. A typology of privacy. *U. Pa. J. Int'l L.* 38 (2016), 483.

[7] Nancy R Mead, Forrest Shull, Krishnamurthy Vemuru, and Ole Villadsen. 2018. *A hybrid threat modeling method, CMU/SEI-2018-TN-002.* Technical Report. Carnegie Mellon University Software Engineering Institute.

[8] Helen Nissenbaum. 2009. *Privacy in Context.* Stanford University Press.

[9] Adam Shostack. 2008. Experiences threat modeling at Microsoft. *MODSEC@ MoDELS* 2008 (2008), 35.

[10] Adam Shostack. 2014. *Threat Modeling: Designing for Security.* John Wiley & Sons.

[11] Daniel J. Solove. 2006. A taxonomy of privacy. *University of Pennsylvania Law Review* 154, 3 (2006). https://ssrn.com/abstract=667622

[12] Blake E Strom, Andy Applebaum, Doug P Miller, Kathryn C Nickels, Adam G Pennington, and Cody B Thomas. 2018. *MITRE ATT&CK: Design and philosophy.* Technical Report. The MITRE Corporation.

[13] Peter Torr. 2005. Demystifying the threat modeling process. *IEEE Security & Privacy* 3, 5 (2005), 66–70.

[14] Anton V Uzunov and Eduardo B Fernandez. 2014. An extensible pattern-based library and taxonomy of security threats for distributed systems. *Computer Standards & Interfaces* 36, 4 (2014), 734–747.

[15] Kim Wuyts. 2012. *LINDDUN: A privacy threat analysis framework.* Technical Report. CS, KU Leuven.

[16] Kim Wuyts. 2020. LINDDUN GO: A lightweight approach to privacy threat modeling. *IEEE European Symposium on Security and Privacy Workshops* 2020 (2020), 302–309.

[17] Wenjun Xiong and Robert Lagerström. 2019. Threat modeling–A systematic literature review. *Computers & Security* 84 (2019), 53–69.