

Threat modeling in Dutch organizations

State of practice

Laurens SION
Stef VERREYDT
Koen YSKOUT

DistriNet Research Group
Department of Computer Science
KU Leuven

Client: NCSC Nederland

Version: 1.0
November 7, 2023



Executive summary

From smart watches and vehicles to digitalized hospitals and power plants, software is becoming increasingly prevalent throughout society. As software-enabled technology is nearly indispensable, so is software security, with single vulnerabilities potentially leading to significant financial, reputational, and physical harm, and in some cases even threats to national security. With this in mind, security is currently being ‘shifted left’ in industry: rather than implementing all functional requirements before testing for vulnerabilities, security is taken into account throughout the whole development lifecycle, including during requirement analysis and architectural design.

Threat modeling is a widely known technique to implement such ‘security by design’ mindset, allowing to identify potential security and privacy threats based on a software design or architecture. Despite threat modeling being considered best practice, there are limited studies on where, how, and how often threat modeling is applied in contemporary software development projects. The goal of this report is therefore to provide qualitative insights into the state of practice on threat modeling.

These insights were gathered through a set of interviews with practitioners from large Dutch organizations. During these interviews, participants were asked to describe how threat modeling is embedded within their organization, which roles are involved in threat modeling activities, how threat modeling is performed concretely, and their general experiences with threat modeling.

The observations in the interviews lead to the following main findings. (1) While there is no singular definition of threat modeling, it is recognized as an important activity for uncovering threats, but also (and perhaps even more so) for raising security awareness among developers. (2) Developers are most prominently involved in the threat modeling process. The introduction of threat modeling into the organization is frequently triggered by security team members with prior experience. Many emphasize the importance of intrinsic motivation in the development team to perform threat modeling rather than imposing it as a mandatory step. In the same spirit, pragmatic techniques are preferred over rigorous processes. (3) Organizations still face numerous challenges related to threat modeling, such as managing the scope, obtaining relevant documentation about the software system, scaling the activity to multiple teams, and systematically following up on the results.

Organizations can use this report to assess their current threat modeling activities, and help inform decisions to start, extend, or reorient them. Furthermore, threat modeling facilitators and researchers may base future efforts on the challenges identified as part of this study.



Contents

Contents	iii
1 Introduction	1
1.1 Threat modeling	1
1.2 Goal and scope	2
1.3 About the researchers	2
1.4 Structure of this report	3
2 Research methodology	5
2.1 Study design	5
2.2 Recruitment process	5
2.3 Data collection process	6
2.4 Analysis procedure	7
3 Thematic observations	9
3.1 Definition and interpretation of threat modeling	9
3.2 Introducing threat modeling in the organization	10
3.3 Motivation for threat modeling	10
3.4 Triggers for threat modeling	11
3.5 Training	12
3.6 Use of models	13
3.7 Eliciting threats	14
3.8 Output and follow-up	15
3.9 People involved	16
3.10 Challenges	17
3.10.1 Planning	17
3.10.2 Training	18
3.10.3 Timing	19
3.10.4 Process	19
3.10.5 Follow-up	21
3.10.6 Use of tools	22
3.10.7 Involving management	23
3.10.8 Demonstrating effectiveness	23
3.10.9 Intra-organizational differences	24
4 State of practice	25
4.1 Embedding of threat modeling activities	25
4.2 Involved organizational roles	26
4.3 Threat modeling process	27
4.4 Threat modeling experiences	29

5 Discussion	31
5.1 Advice	31
5.2 Limitations of this study	32
5.3 Relationships with other studies	33
6 Conclusion	35
Bibliography	37
A Information Sheet	39
B Informed Consent	41
C Interview Guide	43
C.1 Introduction	43
C.2 Demographics	43
C.3 Threat Modeling	43
C.4 Closing/debriefing	44

Software engineering has become an integral part of our digitalized world. It revolves around the process of specifying, designing, programming, documenting, testing, and bug fixing applications, frameworks, or other software components.

With the increasing pervasiveness of technology and its influence on our daily lives, the focus on software security has become paramount. Software security is a quality that must be paid attention to, in order to prevent potential breaches and cyber threats. This is especially important in a setting where a single security vulnerability can lead to significant financial losses, damage to reputation, and even pose threats to national security. A secure software system can be defined as a system that maintains the integrity, confidentiality, and availability of the assets that it manages (in particular functionality and data) to a sufficient degree. Sufficiency alludes to software not being designed to be perfectly secure, but ‘secure enough’ [Sandhu, 2003], as determined by its application domain, stakeholders, and risk appetite.

Contemporary software development happens in small teams that iteratively refine, extend, and refactor a software system in fast cycles (sprints). Furthermore, the DevOps movement calls for a tighter integration between development and operational activities. In this development process, many security-enhancing activities can be performed, ranging from training and the specification of adequate security requirements over source code analysis to pentesting and incident response handling [Microsoft, 2023]. One of these activities, and the focus of this report, is ‘threat modeling’.

1.1 Threat modeling

Threat modeling (sometimes known as architectural risk analysis [McGraw, 2006]) is widely promoted as a best practice in a (secure) software development lifecycle, implementing ‘security by design’. For example, it takes a prominent role in Microsoft’s Security Development Lifecycle (SDLC) [Microsoft, 2023], OWASP’s Software Assurance Maturity Model (SAMM) [OWASP, 2022], NIST’s Secure Software Development Framework (SSDF) [NIST, 2022], and others. Insecure design, for which threat modeling is considered one of the main important mitigation strategies, appears at the fourth place in the most recent (2021) edition of the OWASP Top 10 [OWASP, 2021].

Threat modeling, in essence, is the practice of identifying the most important security threats to a system, which drives the selection of countermeasures to implement. In the words of the ‘Threat Modeling Manifesto’ [Braiterman et al., 2020], and in alignment with the ‘four questions’ framework of Shostack [Shostack, 2014],

“threat modeling is analyzing representations of a system to highlight concerns about security and privacy characteristics. At the highest levels, when we threat model, we ask four key questions: 1) What are we working on? 2) What can go wrong? 3) What are we going to do about it? and 4) Did we do a good enough job?”.

Despite the strong advocacy of threat modeling as a best practice, not much is known about where, how, and how often threat modeling is applied in practice.

1.2 Goal and scope

The goal of this report is to provide qualitative insights into the state of practice on threat modeling within large Dutch organizations that are part of the target audience of NCSC (the sponsor of this research). The insights in this report originate from a set of interviews with practitioners from these organizations, conducted between August 2022 and February 2023. We focus primarily on organizations that have in-house software development teams, but also include organizations without such teams yet focusing on Information Technology (IT) and Operational Technology (OT) infrastructure, as well as one organization that has an advisory role.

Our assessment of the state of practice addresses four broad research questions:

- RQ1. How is threat modeling embedded in the organization?
- RQ2. Which organizational roles are involved in threat modeling activities?
- RQ3. How is threat modeling performed within the organization?
- RQ4. What are the experiences with threat modeling within the organization?

This report can be used by other organizations to compare and assess their current practice to that of other organizations, gain insights and experiences, and help inform decisions to start, extend, or reorient existing threat modeling programs.

Due to the rather small number of interviewees and organizations, it is explicitly *not* the goal of this research to provide a quantitative characterization of the state of practice. Hence, this report refrains from using precise numbers or percentages when discussing observations, as they would give a false impression of accuracy.

1.3 About the researchers

The authors of this report, which also conducted the underlying research, are academic researchers from KU Leuven, Belgium. They are all connected to the DistriNet research group of the department of Computer Science. Their primary

research focus lies on approaches and techniques to provide automated support for threat modeling.

The authors of this report declare no conflict of interest regarding the research presented herein. The research questions were determined in collaboration with the sponsor (NCSC), ensuring that the study addressed relevant and meaningful aspects of the subject matter. However, the subsequent research was conducted independently, safeguarding the objectivity and impartiality of the findings and conclusions.

1.4 Structure of this report

The report is structured as follows. Chapter 2 explains the followed research methodology, including the study design and recruitment process. Afterwards, Chapter 3 bundles the observations from the different participant interviews along a range of different themes. Chapter 4 subsequently answers the research questions based on those observations. Chapter 5 relates the insights from this study to other studies, and discusses the limitations of this study. Finally, Chapter 6 concludes the report.

2.1 Study design

The study design starts from the four research questions (see Section 1.2) to determine the scope and goal of the interviews. Based on these research questions, the interview guide was constructed. This interview guide lists the different topics to discuss with the interview participants in the form of questions. Each of those questions is linked to the particular research questions it addresses. The full interview guide is included in Appendix C.

The interviews themselves are performed using the technique of responsive interviewing [Rubin and Rubin, 2011], to allow the interviewers to pick up on certain answers and delve into more detail when appropriate. This means that the questions from the interview guide were not asked literally nor sequentially; interviews took the form of a natural conversation, merely guided by the topics to discuss.

Ethical approval for this research, based on the design outlined above, was obtained from KU Leuven's ethical committee before potential participants were contacted.¹ All interviewed participants have signed an informed consent form. After an interview, each participant was offered a small gift voucher to thank them for their participation.

2.2 Recruitment process

Given the study's focus on organizations that are part of the target audience of NCSC, the NCSC provided a list of contacts at the relevant organizations to reach out to. As the goal of our research was to gain insight into the state of practice, we informed these contacts that potential interviewees should be directly involved in threat modeling. Some contacts (which had this experience) participated directly in the interviews, others redirected us to other members of their organization. All contacted participants with the relevant expertise agreed to participate in the study. Appendices A and B list the information sheet and informed consent form that were provided to the participants as part of the recruitment process.

In total, 13 participants from 7 organizations agreed to participate, resulting in 10 interviews (3 interviews were joint interviews, where two participants were interviewed at the same time). General characteristics of the interviewed organizations and participants are provided in Table 2.1.

¹KU Leuven Social and Societal Ethics Committee (SMEC), case ID G-2021-4578-R2

Sector	Focus	Participants
Energy	OT Systems	1
Finance	Software development	4
Marine	IT Infrastructure	1
Public sector	Software development, advice	3
Transport	Software development	4

Table 2.1: Overview of the organizations

2.3 Data collection process

Two interviews were conducted on-site at the participant's office location; the others were conducted online (through a Microsoft Teams video call). One researcher took the lead during the interview. For some early interviews, other researchers observed and took notes to ensure consistency with later interviews. Each interview lasted approximately 1 hour, except for joint interviews (approximately 90 minutes). All interviews were conducted in Dutch, except one (in English).

After a short introduction and repeating the agreements on confidentiality and data protection, the remainder of the interview was recorded (using a microphone for on-site interviews, and the built-in Teams functionality for online interviews). During or after the interview, some participants also briefly showed some reports of threat models that they have worked on to illustrate what was said before.

Afterwards, the (audio part of the) recordings was transcribed literally. The transcription process started with an automated transcription, which was subsequently checked and corrected by one of the researchers using the original recording. Automated transcription was initially performed using the built-in functionality of Microsoft Word; for later interviews, a fully-offline implementation of Whisper [Gerganov, 2023] was used.

The interview transcripts were subsequently anonymized manually, by replacing or scrubbing all information that would enable identification of the participant and/or the organization. All copies of recordings, non-anonymized transcripts, and notes are destroyed at completion of the research study.

All participants have received a copy of the anonymized transcript of their interview, and had the opportunity to add remarks, provide corrections, or highlight potentially identifying information. Two participants explicitly confirmed that the information in the transcript was still accurate; one other participant clarified some changes in the organization that occurred after the interview was conducted.

2.4 Analysis procedure

The analysis of the research data involved a systematic coding process of the anonymized transcripts. To facilitate efficient and organized coding, a software package for qualitative data analysis (ATLAS.ti) was employed.

The coding process used a mix of bottom-up and top-down codes, allowing themes and patterns to emerge from the raw data. Initial (top-down) codes were generated in alignment with the research questions (for example, related to demographics/sector organization, process/execution, etc.), augmented with the researcher's recollections from reading through and anonymizing the transcripts. Throughout the coding process, these codes were complemented with (bottom-up) codes that capture significant other concepts, ideas, or phrases relevant to the research questions. After coding, codes with similar meanings or concepts were grouped together into higher-level codes.

Based on the coded interview transcripts, several recurring themes were identified. For these themes, quotations were collected that were coded with relevant codes, and insights were extracted from these quotations. These quotations served as evidence to support the findings and conclusions drawn from the analysis. Finally, to answer the research questions, the observations from the relevant themes for each research question are combined.

This chapter bundles several observations from the threat modeling interviews along a range of different themes. Part of these observations relate to challenges encountered by the participants; these have been bundled in Section 3.10. The observations in this chapter constitute the basis for answering the four research questions in Chapter 4.

3.1 Definition and interpretation of threat modeling

This theme captures the interpretation of threat modeling by the different organizations, and which aspects or elements in their definition of threat modeling are considered important by the participants.

Threat modeling is recognized as an important analysis activity in the development process. This vision is shared by all participants. However, they do note that the activity is not always explicitly labeled or systematically executed. There are numerous ad-hoc or implicit analysis activities in which an assessment is performed and countermeasures are suggested or considered without that being an explicit part of a threat modeling activity (*“that happens a lot, but not structurally and not under the umbrella of threat modeling”*).

The second element that is frequently mentioned by the participants is using threat modeling as a means to analyze and map threats, vulnerabilities, or risks; combined with thinking about potential countermeasures, although some participants note the limited support in this regard.

Finally, some elements of the definitions of threat modeling show more variation between the participants: the importance of considering particular threat actors (*“know your enemy”*), explicitly thinking about key assets (*“what do we want to protect?”*), abuse cases (*“next to the use cases, to also define abuse cases [...] and think what could go wrong in the flow”*), and the supporting role of threat modeling in subsequent activities such as pentesting (*“it is also an excellent basis for [a pentest]”*).

Take-away. Threat modeling has no common definition and its focus (threats, threat actors, assets) varies among participants, yet is seen as an important activity by all.

3.2 Introducing threat modeling in the organization

This theme assesses the introduction of threat modeling as analysis activity in the organization, the timing, and the involved roles.

For most organizations, threat modeling has been introduced fairly recently (i.e. in the past 5–6 years) by the security team. Most profiles that take up an active role in threat modeling had prior experience with pentesting. While the execution of threat modeling is an entirely internal activity (none of the interviewees relied on external organizations to perform threat modeling activities), in one case, external expertise was consciously attracted to introduce threat modeling into the organization (*“I really followed [hired expert] around for 3 months, almost like a shadow, and that helped a lot too.”*), which enabled overcoming organizational challenges and habits (*“[the expert] does not have the bias of the organization and its processes [...] such that we can first determine what [threat modeling] is and what it adds, before it finds the right spot [in the organization]”*).

Take-away. Threat modeling was introduced about 5 years ago in most organizations, often based on prior experience of the person that introduced it.

3.3 Motivation for threat modeling

This theme dives into the motivation that organizations have to execute threat modeling activities as part of their development activities. It assesses the needs organizations have identified for themselves or the benefits they perceive in the execution of threat modeling.

There is a very strong focus on promoting threat modeling internally as a tool. The main goal is to raise awareness of threat modeling as a technique that is available for development teams to apply, rather than forcing the use of threat modeling through organizational policies or procedures. The incentivization or creation of awareness ranges from simply mentioning the existence of the technique (*“tell them once, let it simmer”*), providing support for applying it, to organizing internal workshops (*“we also started to develop workshops, so that people at least understand what threat modeling is and why we do it”*).

Most participants emphasize the need for development teams to internally recognize the relevance and usefulness of threat modeling. This ensures that the motivation comes from within the team (intrinsic motivation) rather than being imposed (extrinsic motivation), which would result in a compliance-like checkbox activity, for example *“The initiative to do [threat modeling] should come from the developers. [...] I would not like it if we would need to push it from management, because the moment you start forcing threat modeling, people naturally lose enthusiasm and do it because they have to and not because they see the usefulness and necessity of it.”* This internal relevance and usefulness is already

appreciated by development teams in several organizations (*“threat modeling is also well received, generally, by the teams”*).

The main benefits organizations perceive in the use of threat modeling are twofold. First, threat modeling is employed to gain understanding and insights into the security concerns of the applications being developed (*“to develop more secure products, and the focus is really on identifying design flaws”*). Second, threat modeling is also a useful tool to raise the overall security awareness of teams (*“they learn to think about threats”*), and to give the teams a way to talk about security (*“a way for them to discuss information security in a practical way within their team”*).

Another potential goal mentioned by one of the participants is to use threat modeling as a vehicle to communicate about potential security issues with non-technical people (*“so that [non-IT] people also get a good understanding of how certain things can occur and how certain things can happen”*). This was particularly apparent in the interviewed organization with an Operational Technology (OT) focus: *“if you look at risk assessments like in the ISO 27000 [series], then [OT engineers and operators] fall asleep, they don’t get it. It’s very IT-oriented. [...] Whereas, [if you ask an operator to] imagine that somebody else would get behind the controls, what can happen? He can explain it to you in great detail [...] and he also likes to explain it. [...] And if he doesn’t think it’s important, why would we act upon it?”*.

Take-away. Support by the development teams (intrinsic motivation) is deemed essential for the successful adoption of threat modeling. Furthermore, threat modeling serves a dual purpose: finding potential vulnerabilities, but also raising the security awareness among development teams.

3.4 Triggers for threat modeling

This theme assesses what the trigger is within organizations to perform threat modeling activities, the timing of threat modeling activities, and to which extent re-assessments are performed.

The majority of triggers are internal from development teams that reach out to ask confirmation or feedback on their threat models, or for training or support in the execution of their analyses. Other triggers are internal promotion of threat modeling by the security team, or the security team approaching development teams to suggest threat modeling. In some cases, threat modeling is explicitly required for certain types of applications (depending on the sensitivity of processed data or the business impact of the application, for example). In general, however, threat modeling is rarely imposed on development teams, as doing so would result in it becoming a checkbox activity (*“once you start having these compliance requirements [...] at that point, they will just not write stuff down*

anymore. So, the question is, what is the impact of that going to be?”)

Regarding the timing and execution of threat modeling activities, there is overwhelming consensus that threat modeling is a continuous effort and thus requires periodic re-assessments. The implementations vary from development teams reaching out for feedback on their updated models to the security team frequently checking in with developers to trigger them in performing a re-assessment if necessary. Furthermore, many participants recognize the opportunity for tooling and automation such as integration in CI/CD pipelines, as this would allow these activities to be better integrated with existing development practices and give development teams a reminder or trigger if changes introduce new threats.

Regarding the timing in the development lifecycle, the usefulness of performing threat modeling early are recognized, but this is in practice not always performed. One of the reasons for this is that there is still a backlog of high-risk applications which require a threat model, leaving less room for the security teams to support early-stage threat modeling sessions. Furthermore, especially in big projects, it is not always straightforward to determine to most opportune time, and an important element is determining the right scope of the analysis to keep the size and complexity of the analysis manageable. Even so, there are several instances where threat modeling was applied very early in the development lifecycle, in tender processes and procurement, leading to valuable feedback and concrete security requirements. For example, in one specific case mentioned by participants, threat modeling during procurement later prevented a specific ransomware attack.

Take-away. The development team typically takes the lead in initiating threat modeling. While it is ideal to (also) conduct threat modeling early in the development process and regularly repeat it, it may not always be possible due to challenges related to size, scope, and available resources.

3.5 Training

This theme assesses the role of training in organizational processes to support the execution of threat modeling activities.

Before teams actually start threat modeling, they need to learn the relevant concepts and methodologies handled during a threat modeling session. In general, the interviewees indicated that a threat modeling session usually starts with an introduction to threat modeling. The duration of such introduction varies from a couple of slides (*“a few slides, two or three, to shortly explain the methodology and the overall structure of a threat modeling session”*) to more lengthy ones (*“we first gave an introduction of about 40-45 minutes about threat modeling”*).

Providing separate learning materials or organizing workshops before the actual threat modeling session is also prevalent. While generally perceived as useful, interviewees indicate that separate learning materials do not suffice to learn

teams to threat model independently (*“I don’t see teams picking it up and doing this completely independently any time soon”*), and teams may not always go through them in advance (*“I don’t think they go through the materials we are sharing with them”*), so a short introduction in the session itself is still advised.

Besides learning the methodology and basics of threat modeling, the following aspects are usually covered during an introduction to threat modeling. First, teaching teams to think about what can go wrong was mentioned several times as an important part (*“worst-case thinking really needs to be taught to people”*). Second, while teams may be more comfortable with following a well-defined threat modeling method, several participants note that the exact methodology or process in general is of little importance, and that thinking about security at all is more important than doing it the correct way or following strict guidelines (*“the most important thing is to start [threat modeling]. You can’t really do something wrong”*). Third, teams may lack security expertise, so some examples or prevalent threats may also be covered during a training session.

A lack of security expertise was mentioned as the main reason as to why teams are not confident to independently start threat modeling (i.e., without the presence of a facilitator or security expert), as for example described by *“the intention is 95% that they could do it themselves. My impression is also that they are perfectly capable of doing it themselves if they have seen it once. That last 5% is indeed ‘what do we [as security experts] see?’. And they can’t do that themselves.”*

Take-away. While dedicated training sessions are both commonplace and essential for instilling the proper mindset, enabling a team to independently execute threat modeling can be challenging.

3.6 Use of models

This theme covers the creation and use of models during threat modeling.

The first step of threat modeling activities involves the creation of a model of the systems or applications being analyzed. Overall, the diagrams created or used in the context of threat modeling can take various forms, ranging from re-used architectural documentation to whiteboard diagrams. There is a balance between diagram quality conventions and the effort for the development teams to adhere to them, as the additional overhead can also limit the threat modeling initiatives. As a result, tool support for creating diagrams is mostly limited to drawing tools like Threat Dragon [OWASP, 2021], but in some cases more elaborate modeling support like Microsoft’s Threat Modeling Tool [Microsoft, 2023] is also used. While a lack of diagram conventions inhibits the use and interpretation of the threat modeling documents by other teams, a related and greater concern is the set of assumptions made by different teams. (*“it is up to [the threat modeling facilitators] to always connect with the teams at the right level, because ultimately*

you also want the teams to also take and feel some ownership using threat modeling.”)

A broadly recognized benefit of threat modeling as an activity, is that it forces the explicit consideration of architectural documentation which can be either non-existing or, more frequently, outdated. This analysis therefore provides an incentive to revise and update the existing architectural documentation. In some cases, existing documentation from other risk analyses or assessments can be re-used. (*“[if the team did not provide architectural documentation] we start with a high-level outline, and then it’s always the case that someone will still shout ‘oh, yes, but there’s also this...’ [...] And then there are always other people from the same team who didn’t know that.”*)

The creation of the documentation is the responsibility of the development teams and architects, although in some cases the security teams can, based on the inputs of the development teams, construct initial diagrams to bootstrap the threat modeling activities. An important concern for the construction of the diagrams is the scope of the analysis to ensure a focused discussion. This can be especially challenging for very large and complex architectures, for which it is difficult to obtain a complete picture of the system. (*“there is no single record, with the truth, not even on a conceptual level.”*)

In terms of model types, data flow diagrams (DFD) were most commonly used for software systems. In the organization focused on operational technology (OT), however, a map of the network layout was used as the primary model.

Take-away. (Up to date) architectural models are not always available for re-use, so (re)constructing them becomes an important part of threat modeling. Concerning the use of models, pragmatism prevails over conforming to standardized notations.

3.7 Eliciting threats

This theme assesses the threat elicitation process itself, including how threats are elicited, which types, and how they are prioritized.

For the threat elicitation step, STRIDE is most frequently mentioned as the main driver for the threat elicitation. While STRIDE is used frequently, the elicitation process itself is not necessarily performed systematically (e.g., using the STRIDE threat mapping table as described in [Shostack, 2014]). Indeed, organizations prefer flexibility in how the process is executed, giving development teams freedom in applying the analysis in a way that suits them. In this context, there are several other resources that are used, such as the background knowledge of involved experts, and interactive processes in which teams are challenged about security concerns. Furthermore, in some cases the methodology can deviate to other techniques such as PASTA [UcedaVélez and Morana, 2015], when the

need arises. (*"We chose STRIDE at the time mainly because it's very easy to explain and very accessible."*)

In addition to the regular model as input for the threat elicitation step, there are several other types of inputs that are frequently leveraged as part of the threat analysis: the ingress points in the applications, attack vectors, types of adversaries, and attack scenarios. A consideration in the use of these additional inputs is the desire to reuse organization-specific knowledge across multiple analysis activities. (*"[...] which threats, and which attackers do we think are interesting?"*)

The prioritization of the elicited threats is a challenging activity. Multiple organizations mention the use of background or domain knowledge as part of the risk assessment or focus mostly on the identification of threats rather than any subsequent risk assessment. (*"[...] even if we have some way to evaluate the risk, they will still be guessing it, it's not going to be accurate enough."*)

Finally, with regard to the process itself, participants perceive the value to be mainly in the process rather than in the quality of a threat model. That is, it is more important to do the analysis, rather than having a detailed and high-quality threat model. (*"[...] the most important thing is to start. You can't really do anything wrong."*) In this sense, there are also generally no strict criteria on when the analysis is finished. Usually, sessions end naturally when no new information arises or when all model elements have been covered.

Take-away. A pragmatic use of the STRIDE acronym is the most common approach for identifying threats during threat modeling. In this context, taking action and moving forward is considered more valuable than achieving a perfect threat model or prioritization of threats.

3.8 Output and follow-up

This theme assesses the outputs of the threat modeling activities in which ways organizations follow-up on the results of the analyses.

In general, threat modeling results in a report that contains the model of the system and the identified threats. In some cases, richer descriptions are made using attack scenarios. Additionally, the mitigations that are already present are listed together with advice or recommendations to resolve unmitigated threats. To reduce the amount of issues, threats can be ranked into a list of top risks, which can be updated over time as threats are mitigated. Diagrams are not always modeled in tools but can also be whiteboard pictures. (*"[...] a summary of the relevant risks, at the basis of which recommendations are made [...]"*)

Overall, organizations want to limit the reporting overhead as writing everything out in textual reports requires substantial effort with limited returns (*"[...] writing*

takes a lot of time, and I don't know if it's always worth the effort. Going through the process is perhaps the most fruitful.”). In some cases, presentations of the results are used to limit such reporting overhead. The execution of the threat modeling process itself is considered more important than the reporting. While tooling is something that is considered, linking the findings to business risks remains a challenge and requires manual effort.

For following up and mitigating the threat results, there are no strict processes in place. Follow-up is more of an ad-hoc activity for which the responsibility usually lies with the team itself (with the exception of some very severe issues where the security team actively follows up). How to monitor and follow up on the results more systematically is a recurring challenge. (*“That varies depending on the team, and also on the priorities of the product owner [...]”*) This will be explored in more detail in Section 3.10.5.

One activity that does frequently occur, and is a form of follow-up to verify the implementation of mitigations, is the use of pentesting. These analyses tend to resurface issues that were not resolved by the teams. Having access to a threat model was mentioned to simplify the pentest process. There is also an opportunity here for positive feedback. Analyses that do not uncover any findings often result in minimal reports, and stakeholders may think that they wasted time and resources without really gaining any value. The observation that the team did properly implement the right mitigations is, however, something that can also be actively communicated to them as positive feedback (*“[as a pentester,] it's not really accepted yet that you just go back to a customer, and say, 'gee, you guys just did a great job'.”*).

Take-away. In most organizations, no strict follow-up processes for the results of threat modeling are in place. As part of such follow-up, positive feedback appears underused.

3.9 People involved

This theme assesses which stakeholders are involved in threat modeling activities and to which stakeholders the outcome of threat modeling activities are reported.

The main stakeholders involved in the threat modeling activities are the development team and the product owners, which typically perform the threat modeling analysis together. The team is supported by a threat modeling facilitator from the security team. To a lesser extent, testers, architects, information security officers, and operations are involved. The lesser involvement of these other roles, such as for example the information security officer, is usually the consequence of their limited availability. Involving incident response people can be particularly useful, enabling the integration of additional insight into which types of security concerns are relevant and actively abused in incidents; however, they are even

more rarely involved (*"[...] they don't have the capacity [to attend threat modeling sessions]", "we share our threat models with [incident response] [...] but I think it would be better if they just join threat model sessions."*)

Because of the availabilities, the information security officers and management positions are often only involved in the communication of the results. However, it is often difficult to communicate these results and clarify the usefulness of threat modeling as an activity. Being able to demonstrate a clear business impact and having success stories can help to communicate the results (*"We share successful [threat modeling] stories from time to time, so that [management] sees the added value."*)

The situation becomes more complex when software is acquired from a third party. In such a case, the main stakeholders involved are not developers, but IT staff (operations), and involvement is required from a third party (the vendor). Not all third parties, however, provide equally detailed documentation regarding the security state of those applications (*"Then you depend on, on the one hand, [third parties] being able to provide information, and on the other hand also the level of maturity on security of those kinds of companies."*)

Take-away. Threat modeling primarily involves the development team and a facilitator from the security team. Involvement of others, especially security officers and management, is rare.

3.10 Challenges

This theme bundles the challenges, negative experiences, and potential pitfalls related to threat modeling mentioned by the participants throughout the interviews. It is subdivided in several sub-themes.

3.10.1 Planning

Several participants described difficulties to find the right time to start or revisit a threat model. Starting too early in the development lifecycle may not be desirable as the scope may not be fully clear yet. Starting too late, on the other hand, may lead to a large scope and therefore long threat modeling sessions. Furthermore, mitigating security issues, especially those rooted in the design, may be difficult or even impossible when applications are already fully implemented or deployed. (*"[...] many of the threat models we do, they tend to be more for stuff that's already out in the field, so then it also becomes, like, what can you still do, right?"*)

Second, simply finding a time slot that fits all the desired participants was mentioned to be challenging several times. (*"[...] the thing that takes the most time is, if you have a few crucial people, finding a hole in everybody's schedule."*) Security teams should therefore take care when approaching teams to perform a threat model. (*"[...] it depends on when you approach [the teams], just before*

[their software goes live], or when they have only just started.”) One participant also mentioned that the security teams themselves may experience difficulties to plan a threat modeling session if teams request it close to their deadline, for example when a threat model is mandatory. (*“Not all teams are aware of our schedule as [the security team], so sometimes teams ask for, yeah, within now and two weeks a threat model must be finished.”*) A more general challenge indicated by one of the participants is that security teams simply may not have the resources to provide threat modeling support to all teams (*“we simply don’t have the capacity for that yet, because we just have so many development teams.”*)

In general, participants described that the best way to tackle planning-related challenges would be to make threat modeling a part of the default workflow of the teams, as they themselves know best when a threat modeling exercise would be opportune. This requires the teams to be aware of the benefits of threat modeling, and potentially some changes to the organization of the overall workflow of development teams. (*“With us, the problem is mostly structure. Awareness among the teams themselves. So actually structure of the organization in the way development teams and applications are looked at at all.”*)

Take-away. As threat modeling is not a default part of the workflow, planning a threat modeling session often proves to be difficult and requires deliberate effort and consideration. A lack of capacity at the side of the security teams hampers scaling up threat modeling efforts.

3.10.2 Training

Participants mentioned multiple challenges related to training teams to become familiar with threat modeling. A first one is that development teams usually lack the attacker mindset which is required to analyze the security of a system. Thinking about what could potentially go wrong rather than in terms of the ‘happy flow’ is not something which developers are used to (*“People often think from the happy flow, like: that thing does this and it works and well, what could possibly go wrong? And if you then confront people with yes, but this could happen, or that could happen, then, well, for some people that seems far-fetched.”*)

Participants also mentioned that teams may be afraid to make mistakes during a threat modeling session, or that they do not correctly apply the chosen methodology. (*“[...] it’s seen as a big step to start doing it. Because most people are afraid of failing or not doing it right.”*) All participants, however, agree that the exact methodology or how threats are identified does not matter, and that every potential issue that is found is beneficial (*“[...] it’s really not so much about whether it’s done very well. The point is that we do it, and that we learn from it together and gain knowledge [...]”*) Making the participants of a threat modeling session aware of this may, however, be a challenge.

One participant mentions that creating worked examples for threat modeling is challenging, both because creating them is time-consuming (*“they tend to be*

very time intensive to actually create”), and because teams tend to focus on the specific material covered in the example, which may hinder them from finding other issues not covered in the examples (“[...] the only thing they’re going to be doing is regurgitating the exact same thing that you told them during the training, at which point, yeah, you can also just give them a checklist”).

Finally, besides teaching the teams to threat model, training people to facilitate threat modeling sessions was also mentioned to be challenging. Concretely, becoming proficient in threat modeling facilitation requires both security expertise and the ability to adapt sessions based on the team (*“every team is different and you have to be very focused on that”*).

One participant also mentions the lack of real-world experiences on how to introduce threat modeling to an organization (*“you rarely hear about, well, I did it this way, and you need this, and you need these contacts, and you need to arrange it this way.”*)

Take-away. It is challenging to get the team in the right mindset during training.

3.10.3 Timing

Managing the duration of a threat modeling exercise was seen as challenging by the participants. Teams may lose interest if a session takes too long, especially if a session is dominated by one or a few people or gets too technical, and teams may be reluctant to start threat modeling a large system or application due to the amount of time that must be invested (*“You have to keep the focus time short, right? So the time you work with the team, you shouldn’t make it too big. Otherwise the team gets bored or there’s no time left.”*) Rather than one long session, participants indicated that multiple shorter sessions are preferable (*“We do see that it is better to have multiple sessions rather than one session of half a day or longer.”*)

Take-away. Short sessions (of about 2 hours) are preferred over longer sessions.

3.10.4 Process

A process-related challenge mentioned by several participants is that architectural documentation is seldom available or up-to-date (*“the documentation we get is almost never up-to-date”*), which hinders the creation of models and diagrams (*“[...] the fact that we have to spend the beginning of a session on getting the model correct, or as correct as possible is, in my view, a bit of a waste of time.”*) An underlying problem is that multiple teams work on different parts of a system, and that a single overview of how it all comes together is usually not available (*“there is no single record, with the truth, not even on a conceptual level”*), not

even by the architects. Involving multiple architects during a threat modeling session to tackle this issue was also mentioned not to be favorable by one of the participants, as this may lead to lengthy discussions concerning the architecture. (*"[...] the risk if you invite three architects to a session is that the discussions go in all directions. So then we prefer to have only the architect who is most involved there."*)

Concerning the methodology, while certain situations may warrant a different approach, in general it is preferable to choose one methodology and stick to it. Otherwise, a considerable amount of time may be spent discussing and deciding on the specific methodology to use (*"If you aren't careful, you will have a lot of discussions about the form before you actually get started. And then you don't get anywhere."*) Specifically for STRIDE, one participant mentioned that it does not scale well, as even for smaller applications, the amount of threats may rise rapidly (*"[...] as the number of flows in and out of an application increases, the amount of time you have to spend on it increases exponentially."*) As a result, applying STRIDE during more agile work flows was indicated to be cumbersome (*"If you then, for example, wanted to apply threat modeling in an agile sprint or something like that, STRIDE is quite a cumbersome method, yes."*) One participant describes that the data flow diagram notation may not be ideal for more specific and technical types of analyses (*"For more the protocol related things, for example, this is where it kind of, kind of breaks down [...] because you really want to look at much more specific and technical issues."*)

Several participants indicated challenges concerning risk estimation. Determining the impact and likelihood of threats requires both security expertise and domain knowledge, and guidelines on how to do so are lacking in general (*"First, we don't provide a clear framework, how to do that themselves, and second, even if we had some way to evaluate the risk, they would still be guessing it, it's not going to be accurate enough."*).

Other process-related challenges include not thinking about the attacker (*"Knowing who you're up against... I notice that a lot of people don't talk about that"*) and approaching a threat model too much from a pentest point of view, which may lead teams to get stuck on the details (*"[sometimes] we treat the threat model a little too much as a starting document for our pentest, rather than a standalone thing. And that manifests itself, for example, by... Well, getting very technical in depth on certain things."*) Finally, multiple stakeholders from across an organization may be involved in a threat modeling session, which may make communication challenging. (*"Totally different sides of an organizations are suddenly going to be collaborating [...] Purely on language alone, you have to be very careful with that."*)

Finally, supply chain management with respect to security is becoming increasingly important. It no longer suffices for your own system to be secure, but all other systems and applications on which you depend must also be analyzed (*"I am starting to think about it more and more, yes we are fine, but what about our suppliers?"*)

Take-away. It is advisable to keep the sessions focused and free from overhead related to diagramming, architecture discussions, deciding on the methodology, or technical discussions.

3.10.5 Follow-up

The participants indicated that whether or not teams follow up on the outcome of a threat modeling session varies between teams. The consensus among the participants is, however, that follow-up is lacking in general. Security may not be a priority of the team or product owner, which may lead to the outcome of a threat modeling session being ignored (*“Some teams, [...] just ticked the box, like, ‘okay, there are recommendations in there, but, yeah, our product owner doesn’t think that’s exciting enough right now.’”*) This is especially the case when threat modeling is mandated by some policy (*“External supervisors, they just want a list, and ticked off, and then you’ve done well.”*) Participants do agree that this is not due to the lack of security interest, but rather because teams have limited time (*“It’s not that they don’t want to do security, but they have so many other things to think about besides security”*).

Following up on the outcome was also mentioned to be difficult for multiple reasons. First, acting on the results of a threat modeling exercise may require the help of people external to the team, for example when applications are hosted externally. In such cases, it may take time to get this on the agenda of the external entity (*“To solve an issue [with an external host] would involve creating a ticket, and most likely lengthy email conversations, phone calls, ...”*) Furthermore, as mentioned previously (Section 3.10.1), threat modeling sessions are planned late in the development cycle in some cases, which limits the amount of changes that can be made to an application (*“[...] and then we find out that there are actually quite insurmountable problems in the software”*).

Participants also described that following up is challenging if it involves taking into account other teams or stakeholders within the organization. For example, there is a risk of interfering with previously made (design) decisions, potentially taken by other teams (*“What a lot of organizations suffer from is that, they all take separate, siloed actions and don’t take into account what preceded it, or too late.”*) This is especially relevant when there is a business incentive to deploy an application as soon as possible. In such cases, deciding what to do or how to process the output of a threat modeling session (if at all) may become tedious and time-consuming (*“[...] that generates a lot of discussion, which in turn slows down the accreditation process.”*) Furthermore, even if teams want to take into account, for example, the outcome of a preceding threat modeling exercise, interpreting the results was indicated to be challenging by the majority of the participants (*“It might be a problem with other teams interpreting threat models, one team interpreting a threat model [differently from] another team.”*) Standardizing the threat modeling process and output within an organization may be one way to tackle this challenge, but too much standardization may deter

teams from threat modeling at all (“[...] then you do get some interchangeability of [threat modeling results], without immediately killing the whole enthusiasm by putting it in a straitjacket, because that’s not the goal either.”) Finally, one participant describes the risk of assuming that other stakeholders will take care of an issue (“Assuming that another team does something [...] is] more a problem than having the same circles, squares, arrows and whatnot.”)

Take-away. Follow-up and sharing of threat modeling outcomes is generally lacking, both by the development teams and other stakeholders within the organization.

3.10.6 Use of tools

Several challenges relating to tool support were described by the participants. First, tools (e.g., the Microsoft Threat Modeling Tool [Microsoft, 2023] and Threat Dragon [OWASP, 2021]) were indicated not to be user friendly (“I find that it lacks some things in terms of usability”). Microsoft’s Threat Modeling Tool specifically was mentioned to require a lot of detailed inputs in order to get to useful output (“You really have to fill out a lot to get useful information”; “You also don’t want to tire the team with all those details, like, what TLS version are you using, and stuff like that”). Second, interpreting the output of threat modeling tools was also indicated to be challenging, mainly because it requires security expertise (“At the very least you want to prevent [the teams] from, yes, not having the knowledge and, yes, then simply disregarding [the output]”). For these reasons, except to draw simple diagrams, using threat modeling tools during a session was generally avoided.

One participant mentioned that, in order to make threat modeling tools a part of the general work flow of teams, tools should be simplified (“[...] reduce to a simple implementation so that teams can start using it at all”). Another issue mentioned by one of the participants, specifically on why they stopped using [a commercial tool], is that it does not allow to model business logic well (“[...] it’s not really very easy yet to include business logic”). Finally, while participants indicated that integrating threat modeling tools in a CI/CD pipeline could be beneficial, none of them do so at the moment (“At the moment [...] I don’t see how you could integrate threat modeling specifically into your CI/CD pipeline.”). One participant described the idea to automatically create tickets for threats, but due to the amount of threats that are identified by threat modeling tools, this could also be challenging (“[...] just have ten thousand tickets automatically open... That’s not going to be nice.”)

Take-away. The use of tools is limited, because the effort required to use them outweighs the benefits.

3.10.7 Involving management

Participants mentioned several challenges related to the interest or involvement of (risk) management in threat modeling. First, management may not always be aware of the added value of threat modeling, which makes getting support, time and resources for threat modeling challenging (*“Getting resources to do it from the higher-ups, that always requires work.”*) Ideally, according to one participant, management should not push or mandate threat modeling, but rather be supportive when teams indicate that they would like to do so (*“I would hate to have to push that from a leadership role. [...] But management, according to me, does play a role in accepting it, seeing the added value of it and being able to translate that back to their stakeholders as well.”*)

Second, participants described that involving someone from management during a threat modeling session could provide useful insights, but doing so may not be straightforward for two reasons. First, management may not be aware of the benefits of them being present, and furthermore, management may think that threat modeling sessions require a strong technical and/or security background (*“They are very quickly afraid that it really becomes a very technical session.”*) Second, management simply may not have the time to join threat modeling sessions (*“For the whole [domain], we have a single ISO right now. [...] Yeah, that’s too few.”*)

Finally, management does not sufficiently follow up on the results of threats modeling sessions according to several participants (*“[...] that just doesn’t always happen or, at least, not consistently”*). Even if management would like to follow up, they may not always be able to correctly interpret threat modeling reports, because they are not always directly involved or familiar with the context (*“You need to be able to interpret a report [...] and [if] we are not in a position to give the voice-over there, then that can lead to differences in interpretation.”*) The result of the lack of follow-up from management is, in general, a lack of oversight across applications and an organization in general (*“[...] that leads to lack of oversight, where you can miss things”*).

Take-away. Involving management roles is often difficult, in part due to unawareness of their own added value and detachment from technical details.

3.10.8 Demonstrating effectiveness

Measuring the effectiveness of threat modeling, and security in general, is indicated to be challenging (*“Evaluating whether threat modeling helps to achieve security is very hard, because you can’t really measure security”, “It’s an article of faith and we are part of the threat modeling church.”*) However, in order to create awareness and motivate teams to do threat modeling, being able to communicate its added value may be crucial (*“What is the added value of threat modeling, right? And I think, making that clear and communicating unambiguously [and]*

empirically backed up [...] That will be decisive.”) One participant mentions that the results of a pentest could be a starting point to evaluate a threat model, for example to identify if some issues that show up during a pentest were missed during the threat modeling session. (*“Taking the insights from a threat model to a pentest and checking [...], like, does the pentest show up stuff that wasn’t in a threat model or assumption that were incorrect?”*) Other than evaluating the impact of threat modeling, evaluating the artifacts created and used during a threat modeling session itself is also indicated to be challenging (*“Looking at the artifacts themselves, and what can we say about those, and I think that’s also an area that’s still a bit open.”*)

Take-away. It is hard to unambiguously demonstrate the effectiveness of threat modeling, and as such the lack of evidence may prevent teams and management to invest in threat modeling.

3.10.9 Intra-organizational differences

While our interviews only include one participant with a focus on Operational Technology (OT, including for example industrial control systems), an important source of difficulties for that participant stems from the inherent (cultural) differences between the IT and OT domain. Mitigating certain threats or creating more secure systems may involve enforcing policies (for example related to patching), also on the OT side, even though IT policies don’t always translate well to an OT context (*“IT organization as I know them are often quite bold and understand little of the OT, yet they feel we must comply with their policies.”*) Understanding the differences between IT and OT, and good communication between both sides, is therefore seen as an important aspect of security in general, albeit challenging (*“let’s get closer to understand each other’s worlds better and, therefore, let’s not play any more yes-no games, but really embrace the fact that our worlds are different.”*)

Take-away. A difference in security culture between different parts of an organization may lead to friction.

This chapter revisits the research questions, outlined in Section 1.2, and answers them based on the thematic observations described in Chapter 3. The answers to the research questions are annotated with references to the relevant observations used to answer those questions.

4.1 Embedding of threat modeling activities

The first research question concerns the embedding of threat modeling activities in the organizations. This research question is split up into three sub-questions concerning (1) the motivation, (2) the benefits, and (3) using the results.

RQ1 How is threat modeling embedded in the organization?

RQ1.1 Why do organizations threat model?

RQ1.2 What are the biggest benefits perceived by organizations?

RQ1.3 How are the threat modeling results used?

RQ1.1. Why do organizations threat model?

While there is no consensus on the definition of threat modeling and what these activities specifically entail, all participants do recognize and agree on the importance of threat modeling (3.1). Additionally, the concrete motivation for performing threat modeling originates at the development teams themselves, albeit often with encouragement from the security team (3.3). In some organizations, threat modeling is required when an application is considered critical (3.4).

Furthermore, the intrinsic motivation of the development teams to perform such analyses was considered an important aspect by many participants, stressing the desire to have the teams want to perform such activities rather than a mandatory assessment that would be perceived as checkbox compliance exercise (3.3).

RQ1.2. What are the biggest benefits perceived by organizations?

The obvious benefits perceived by the participants is the identification of relevant security threats, as this is the primary reason to execute these types of analysis activities (3.3). An important secondary benefit that was recognized by many participants is the effect of raising security awareness among the development teams (3.3).

RQ1.3. How are the threat modeling results used?

Using the threat modeling results and especially the more systematic use and follow-up of the results is more of a challenge for organizations (3.8 and 3.10.5).

Organizations do not have any strict follow-up procedures in place to address the identified security threats (3.8): the results are owned by the development teams and following up on them is their own responsibility. Systematic follow-up of the results by higher organizational levels (management, executive, and/or risk divisions) does not happen. Participants also identify the missed opportunity of using it as a positive feedback mechanism (3.8).

4.2 Involved organizational roles

The second research question concerns the involved stakeholders, specifically (1) who is involved during threat modeling activities, (2) who introduced threat modeling at the organization, (3) the goal of management and operations, and (4) the involvement of third parties.

RQ2 Which organizational roles are involved in threat modeling activities?

RQ2.1 Who is involved in threat modeling activities?

RQ2.2 Who introduced threat modeling at the organization?

RQ2.3 What is the role of management and operations?

RQ2.4 Are any parts of threat modeling performed by external parties?

RQ2.1. Who is involved in threat modeling activities?

Promoting threat modeling and making development teams aware of its benefits is mostly done by dedicated security teams (3.3). In general, the development teams themselves are responsible to start threat modeling, but in some cases the security team may also suggest or mandate threat modeling, especially for high-risk applications (3.4).

The main stakeholders during a threat modeling session are the development team and the product owner. Usually, there is also a facilitator from the security team as the developers are not that confident about their expertise in performing these type of assessments (3.5 and 3.9). Testers, architects, information security officers, and operations are usually not involved, although their input may be valuable to, for example, create models (architect) and estimate risk (information security officer) (3.10.4).

Following up on the results of a threat modeling session is mostly the responsibility of the development teams themselves (3.8). Follow-up from management is lacking in general (3.10.7).

RQ2.2. Who introduced threat modeling at the organization?

In many cases, the introduction of threat modeling was triggered by prior (pen-testing) experience of a recently hired security team member. Hiring external expertise for this particular purpose was only done by one organization (3.2), and was well-received (3.2). The security team then further propagates threat mod-

eling within the organization (often as one of its services towards development teams) (3.3).

RQ2.3. What is the role of management and operations?

Management positions are rarely involved during threat modeling sessions (3.9). Furthermore, follow-up by management is lacking, and challenging in general (3.10.7). Operations, including members of the Security Operations Center (SOC), are also rarely involved, except in the case where applications are bought from third-parties and need to be integrated. In such a case, operations are the main stakeholder (3.9).

RQ2.4. Are any parts of threat modeling performed by external parties?

In all interviewed organizations, threat modeling is performed in-house, with support from the security team. In a single case, however, external expertise was consciously attracted to introduce threat modeling into the organization, which enabled overcoming organizational challenges and habits (3.2). Also, if applications are bought rather than developed in-house, it may be necessary to involve the provider of the application when making a threat model of the integration (3.9). Mitigating threats identified during threat modeling may require help of third parties, for example when applications are hosted externally, but this is also indicated to be challenging (3.10.5).

4.3 Threat modeling process

The third research question concerns the threat modeling process, including (1) the trigger, (2) methodology, (3) inputs, (4) model evolution, (5) output, and (6) quality controls.

RQ3 How is threat modeling performed within the organization?

RQ3.1 What is the trigger to start threat modeling?

RQ3.2 What steps are involved in threat modeling?

RQ3.3 What kind of inputs (models/abstractions) are used?

RQ3.4 How do threat models evolve over time?

RQ3.5 How are results documented and tracked over time?

RQ3.6 How is the quality of threat models assured?

RQ3.1. What is the trigger to start threat modeling?

Threat modeling sessions are mostly triggered by development teams wanting to examine the security of their system or application. In some cases, security teams may approach the development teams to start a threat modeling session, for example for high-risk applications. In general, however, strict requirements for threat modeling are not advised as to prevent it from becoming a checkbox activity. The consensus among participants is that threat models should be started early on in the development lifecycle and require periodical re-assessments, but this

not common practice as security teams are currently mostly catching up on a backlog of high-risk, operational systems (3.4).

RQ3.2. What steps are involved in threat modeling?

In general, a threat modeling session includes the following steps. First, the facilitator (usually a member of the security team) introduces the involved stakeholders to threat modeling, covering the goals, methodology and overall structure of a threat modeling session (3.5). The length of such introduction varies from a few slides to more elaborate presentations, and in some cases, stakeholders are asked to go through learning materials in advance. Second, a model of the system under analysis is created. Occasionally, the facilitator creates a draft model upfront based on existing architectural documentation, which is then discussed and finalized during the session itself (3.6). Third, the created model is analyzed and threats are elicited, usually using the STRIDE methodology (3.7). Threat elicitation may be systematic or in a more pragmatic form based on the context, and usually ends when all model elements have been covered. Finally, after the session, a report is created by the facilitator and distributed to the stakeholders (3.8).

RQ3.3. What kind of inputs (models/abstractions) are used?

Software models used during threat modeling take various forms ranging from free-form whiteboard drawings to structured notations like data flow diagrams (3.6). Models are based on existing architectural documentation, which may not always be available (3.10.4). Reconstruction of the architecture is therefore considered an integral part of a threat modeling session. In addition to the architecture, other inputs mentioned by participants include ingress points, attack vectors, types of adversaries, and other assumptions (3.7). The use of tool support is mostly limited to drawing diagrams for the system (e.g., using Threat Dragon [OWASP, 2021]), as using more features of these tools may complicate the drawing process and, as a result, also the threat modeling session (3.10.6).

RQ3.4. How do threat models evolve over time?

Participants agree that, in theory, threat modeling is a continuous effort and thus requires periodic reassessments (3.4). In practice, such reassessments, and following up on threat models in general, depends on the willingness of development teams and potentially the priorities of the product owner, and is overall challenging (3.8). Participants recognize the opportunity for tooling and automation such as integration in CI/CD pipelines to trigger reassessments if changes may introduce new threats (3.4), but none of the organizations do so at the moment, mostly due to the lack of tool support (3.10.6). In general, threat modeling therefore remains mostly a one-shot activity, and models are infrequently revisited or updated.

RQ3.5. How are results documented and tracked over time?

In general, threat modeling sessions result in a report describing the system model, identified threats, existing mitigations and mitigation advice (3.8). In

some cases, the identified threats are ranked or prioritized, but doing so is mentioned to be challenging as it requires both security and domain knowledge (3.10.4). Follow-up is mostly an ad-hoc activity for which the responsibility lies with the team itself, with the exception of some very severe issues where the security team actively follows up. Follow-up from management is lacking overall (3.10.7). While going through the process to create security awareness is, in some cases, the main goal, some participants expressed a wish for more frequent and standardized follow-up, but strict requirements may not be favorable and result in compliance-like checkbox activities (3.3).

RQ3.6. How is the quality of threat models assured?

Other than consensus between stakeholders, no controls were mentioned to manage the quality of the models and output (3.7). In general, participants perceive the value to be mainly in the process (and raising awareness) rather than in the quality of the models and output.

4.4 Threat modeling experiences

The fourth and final research question concerns experiences with threat modeling, including positive experiences (1), challenges (2) and causes of difficulties (3).

RQ4 What are the experiences with threat modeling within the organization?

RQ4.1 What are successful experiences with threat modeling?

RQ4.2 What are challenges with threat modeling?

RQ4.3 What are the causes of the experienced challenges?

RQ4.1. What are successful experiences with threat modeling?

Overall, a major success experience consists of teams becoming increasingly aware of security and the advantages of threat modeling and gaining valuable insights from it (3.3). In some cases, these insights directly prevented concrete attacks (i.e., ransomware attacks) (3.4). Furthermore, threat modeling is mentioned to decrease the effort required to develop pentests (3.8). Participants also indicate that teams are starting to threat model earlier in the development lifecycle, and do so more periodically, which has a positive impact on the complexity and duration of threat modeling sessions (3.4). Threat modeling during the design phase, although not very prevalent, was also indicated to be beneficial, leading to concrete security requirements which can be taken into account throughout the remainder of the development lifecycle (3.4). Finally, involving external organizations to introduce teams and organizations to threat modeling was also indicated to be beneficial (3.2).

RQ4.2. What are challenges with threat modeling?

A comprehensive discussion of the challenges is provided in Section 3.10. In summary, challenging aspects include the following:

- Finding the right time to start a threat model and finding a time slot that fits all stakeholders (3.10.1);
- Learning teams to think like an attacker, and dealing with the overall lack of security expertise in general when introducing teams to threat modeling (3.10.2). Because the presence of a security expert remains necessary, this challenge hampers scaling threat modeling to multiple teams;
- Managing the scope (3.6) and duration of a threat modeling session (3.10.3);
- Overhead during threat modeling sessions related to, among others, the lack of architectural documentation, discussing and deciding on the methodology, risk estimation, and long, technical discussions (3.10.4);
- The lack of follow-up (3.10.5), adequate tool support (3.10.6), and management involvement (3.10.7);
- Demonstrating the effectiveness of threat modeling (3.10.8); and
- Different (security) cultures between part of the organization, and IT and OT in particular (3.10.9).

RQ4.3. What are the causes of the experienced challenges?

Challenges concerning motivation, timing, and follow-up are mainly caused by product owners, information (security) officers, and other management roles not being aware of the benefits of threat modeling (3.10.7). A root cause for this is that demonstrating the effectiveness of threat modeling is challenging (3.10.8). Teaching teams how to do threat modeling is furthermore complicated by a lack of a security mindset and knowledge within the team (3.10.2). Finally, the limited use of software tools for threat modeling is due to the required effort that outweighs the gained benefits (3.10.6).

In this chapter, we extract advice for organizations based on the observations in this study, discuss the limitations of the study, and compare the findings of this study to other studies in the literature.

5.1 Advice

Based on this study's findings, the main advice for organizations is to consider and incentivize thinking about security in any shape or form, rather than mandating threat modeling and imposing strict requirements on the methodology. There is no one-size-fits-all threat modeling approach (nor software tool) that has worked for every organization that was interviewed. Especially for organizations that are yet to start or just introduced activities related to threat modeling, it seems that successful instantiations of threat modeling spring from giving some space and flexibility to the development and security teams to see if, where, and how threat modeling can provide value, and gradually building upon and expanding this expertise.

In an ideal scenario, threat modeling is done early in the development lifecycle, as mitigating discovered threats in large, existing systems that are already implemented and running may not be straightforward. Indeed, several organizations have highlighted difficulties in scaling up threat modeling, in particular for existing applications, and determining an appropriate scope for such endeavors. Furthermore, organizations agree that threat modeling should ideally be repeated upon important changes to the system (e.g., new features, or changes to the architecture). Participants recognize the potential benefit of using tool support to automatically trigger re-assessments of threat models when such important changes are made to a system, for example as part of a CI/CD pipeline, but currently available threat modeling tools do not offer such capabilities. An important insight from this study, however, is also that one of the major perceived benefits of threat modeling is to raise awareness about security with the development teams — which can be a valuable objective by itself, irrespective of the security improvements from doing threat modeling early and repeatedly.

Product owners and management roles in general need to be aware of the potential benefits of threat modeling, and allow for the necessary time in the planning of the development teams to learn and apply this skill. Therefore, besides incentivizing development teams, awareness campaigns aimed at management roles could be fruitful. Such raised awareness may also contribute to a better follow-up of threat modeling results; in many of the organizations, there appears to be a limited interest in the results from threat modeling by others than the development team, and follow-up actions remains limited. One possible avenue

to achieve this could be to involve these roles in the threat modeling activities; many participants have indicated that this would be valuable, yet care must be taken that such sessions then do not become too technically-focused. On the flipside, additional research towards the effectiveness and return of investment of threat modeling may be needed to fully convince management roles about the effectiveness and ROI of threat modeling, as there are currently very little studies on this topic, and therefore also little empirical evidence available.

5.2 Limitations of this study

This study has several limitations, which may prevent the generalization of its findings to larger populations.

First, this study is based on only a few organizations (7 in total), where often only one person of each organization was interviewed. Although that person was always well-placed and had a comprehensive view on the embedding of threat modeling in the organization (i.e., not just a lone developer, but a member of the organization-wide security team), they may not be fully aware of other initiatives, uses, or impact of threat modeling within the organization.

Second, this study is subject to several selection biases. It is performed on target organizations of the study's sponsor (the NCSC), which typically are large organizations with a dedicated in-house security team that provide critical services in the Netherlands; software development is not their main activity. Contacts within the organization were provided by the NCSC, based on previously established relationships. This means that the results are not necessarily representative for other (smaller and/or commercial) organizations and companies.

Furthermore, all interviewed organizations are already implementing some form of threat modeling and are willing to openly talk about it. In other words, the interviewed organizations already have gained some maturity in threat modeling, which is not necessarily representative for other (similar) organizations. Finally, most of the interviewees are threat modeling 'advocates' within the security team, which means that they appreciate the value of threat modeling, and actively push the practice in their organization. This study does not include any organizations that have tried and abandoned threat modeling, or where no threat modeling program is being developed — no such organizations were present in the initial selection of participants.

Third, because of the use of interviews as the only research method, there is a possibility for respondent bias, where respondents' answers provide an idealized or exaggerated version of the true state of affairs. Some interviewees showed threat modeling reports of projects in which they participated to illustrate what was said, which partially tackles this bias regarding the findings related to process and outcomes. Furthermore, with the duration of the interviews limited to one hour (or 90 minutes if two participants were interviewed at the same time) and the use of a responsive interviewing style, some of the topics listed in the interview

guide were not always explored in equal depth in each interview.

A final remark is that the focus of this study are activities under the name of ‘threat modeling’ (and ‘system-centric threat modeling’ in particular). Other organizations may perform similar activities under a different name (for example, performing a security design review, security risk assessment, or the creation of abuser stories). To obtain a more complete picture, a broader study that focuses on all design-level security activities would need to be conducted.

5.3 Relationships with other studies

The goal of this report is to provide insights on threat modeling within large Dutch organizations. A short overview of similar studies is provided here.

[Shostack, 2008] describes the STRIDE methodology as used by Microsoft back in 2008, and lists multiple challenges which are still relevant according to our observations, including integrating threat modeling in the development process (3.4), following up on the results (3.10.5) and teaching developers to think like an attacker (3.10.2).

[Dhillon, 2011] describes real-world experiences with threat modeling at EMC Corporation (now Dell EMC). They describe how the use of a threat library helps developers to gain security expertise and makes threat modeling more consistent (by making sure that well-known threats are identified) and predictable (by preventing far-fetched and hypothetical discussions). The major downside of relying on threat libraries, as described by [Dhillon, 2011], is that other threats are not considered, which is also mentioned to be challenging by participants in our study (3.10.2).

[Bernsmed et al., 2022] presents results from four different studies on threat modeling in agile projects. The paper describes how teams apply STRIDE, challenges with drawing diagrams and eliciting threats, the usefulness of Microsoft’s Threat Modeling Tool, and how to make the results from threat modeling more useful to agile teams. The challenges listed in the paper correspond to the observations in this report.

[Jamil et al., 2022] discusses findings from interviewing 11 security experts regarding threat modeling cyber-physical systems (CPSs). Similar to our findings, they describe that software threat modeling approaches and tools are not fit for an OT context (3.10.9).

Some of these studies focus on specific organizations (Microsoft/EMC), methodologies (STRIDE), and/or contexts (agile projects/cyber-physical systems). Furthermore, they primarily consider the process (i.e. how threat modeling is applied). In comparison, this report provides a more general overview across several organizations, considering not just how threat modeling is applied, but also the when and why, and who is involved.

This report describes insights into the threat modeling state of practice through a set of interviews with practitioners within large Dutch organizations.

In terms of organizing threat modeling activities, organizations tend to foster intrinsic interest in threat modeling rather than putting strict policies in place. The goals for threat modeling activities are finding and mitigating security threats, but also to raise the overall security awareness among developers. Following up on threat modeling results is indicated to be challenging.

The main stakeholders of a threat modeling exercise are the development team and a facilitator from the security team. Testers, architects, and operations are usually not involved, but participants indicate that their input may be valuable. When software is bought and integrated rather than developed in-house, however, operations are usually the main stakeholder, and the external provider may be asked to provide inputs to ensure a secure integration of the acquired system.

In general, a threat modeling session goes as follows. First, a facilitator from the security team provides an introduction of threat modeling, including an overview of the methodology (usually based on STRIDE). Then, a model of the system is constructed, the form of which ranges from whiteboard drawings to structured notations like data flow diagrams. Constructing a model may be time-consuming if architectural documentation is lacking. This model is subsequently analyzed, typically in a pragmatic manner. After the session, the facilitator creates a report which is distributed to the stakeholders. Follow-up is mostly ad-hoc, except when critical issues are identified. In general, this is a one-shot activity, but participants agree that threat modeling should be a continuous effort with periodic reassessments.

Positive experiences include the prevention of concrete attacks (albeit seldomly), and (much more commonly) the increased security awareness among developers. Challenges are related to, among others, timing, training, model creation, risk estimation, and follow-up. These are (at least partially) associated to product owners and/or management roles not being aware of the benefits of threat modeling, resulting in less allocated time for it by the development teams, as well as a lack of capacity of the security team to assist the development teams.

Organizations can use this report to assess their current threat modeling activities, and help inform decisions to start, extend, or reorient them. Furthermore, threat modeling facilitators and researchers may base future efforts on the challenges identified in this study.

Bibliography

- [Bernsmed et al., 2022] Bernsmed, K., Cruzes, D. S., Jaatun, M. G., and Iovan, M. (2022). Adopting threat modelling in agile software development projects. *Journal of Systems and Software*, 183:111090.
- [Braiterman et al., 2020] Braiterman, Z., Shostack, A., Marcil, J., de Vries, S., Michlin, I., Wuyts, K., Hurlbut, R., Schoenfield, B. S., Scott, F., Coles, M., Romeo, C., Miller, A., Tarandach, I., Douglén, A., and French, M. (2020). Threat modeling manifesto. <https://www.threatmodelingmanifesto.org/>.
- [Dhillon, 2011] Dhillon, D. (2011). Developer-Driven Threat Modeling - Lessons Learned in the Trenches. *IEEE Security & Privacy*, 9(4):41–47.
- [Gerganov, 2023] Gerganov, G. (2023). whisper.cpp. <https://github.com/ggerganov/whisper.cpp>.
- [Jamil et al., 2022] Jamil, A.-M., Ben Othmane, L., and Valani, A. (2022). Threat Modeling of Cyber-Physical Systems in Practice. In Luo, B., Mosbah, M., Cuppens, F., Ben Othmane, L., Cuppens, N., and Kallel, S., editors, *Risks and Security of Internet and Systems*, Lecture Notes in Computer Science, pages 3–19, Cham. Springer International Publishing.
- [McGraw, 2006] McGraw, G. (2006). *Software security: Building security in*. Addison-Wesley Professional.
- [Microsoft, 2023] Microsoft (2023). Threat Modeling Tool. <https://aka.ms/tmt/>.
- [Microsoft, 2023] Microsoft (2023). What are the microsoft sdl practices? <https://www.microsoft.com/en-us/securityengineering/sdl/practices>.
- [NIST, 2022] NIST (2022). Secure Software Development Framework (SSDF) Version 1.1: Recommendations for Mitigating the Risk of Software Vulnerabilities (SP 800-218). <https://csrc.nist.gov/Projects/ssdf>.
- [OWASP, 2021] OWASP (2021). OWASP Top 10 - 2021. <https://owasp.org/Top10/>.
- [OWASP, 2021] OWASP (2021). Threat Dragon. <https://owasp.org/www-project-threat-dragon/>.
- [OWASP, 2022] OWASP (2022). Software assurance maturity model. <https://owaspsamm.org/>. Version 2.0.3.
- [Rubin and Rubin, 2011] Rubin, H. J. and Rubin, I. S. (2011). *Qualitative Interviewing: The Art of Hearing Data*. Sage.
- [Sandhu, 2003] Sandhu, R. (2003). Good-enough security. *IEEE Internet Computing*, 7(1):66–68. Conference Name: IEEE Internet Computing.

- [Shostack, 2008] Shostack, A. (2008). Experiences threat modeling at microsoft. In Whittle, J., Jürjens, J., Nuseibeh, B., and Dobson, G., editors, *Proceedings of the Workshop on Modeling Security (MODSEC08), Toulouse, France, September 28*, volume 413 of *CEUR Workshop Proceedings*. CEUR-WS.org. <https://ceur-ws.org/Vol-413/paper12.pdf>.
- [Shostack, 2014] Shostack, A. (2014). *Threat modeling: Designing for security*. John Wiley & Sons.
- [UcedaVélez and Morana, 2015] UcedaVélez, T. and Morana, M. M. (2015). *Risk Centric Threat Modeling: process for attack simulation and threat analysis*. John Wiley & Sons.

This is a copy of the information sheet that the participants received prior to the study.

	GEDISTRIBUEERDE EN VEILIGE SOFTWARE (DISTRINET) DEPARTEMENT COMPUTERWETENSCHAPPEN CELESTIJNENLAAN 200A 3001 HEVERLEE, BELGIË	
Invitation to participate in study The use of system-centric threat modeling in Dutch organizations.		
Information sheet		
This information sheet provides you with additional information about this study which may help you in deciding on whether or not to participate.		
Goal. While threat modeling is perceived by the IT-industry as an important activity to secure systems during their design and development, little research is available on its use in practise, including potential obstacles and solutions. The goal of this study is to answer the main question: "<i>How is system-centric threat modeling used within the NCSC target audience?</i>". Through this study, we thus aim to investigate how (system-centric) threat modeling is used and experienced within Dutch organizations. The results of this study will contribute to a better, scientific view on threat modeling, which forms the basis for improved recommendations and techniques in the future. Conditions. Participation in this study requires experience with and/or knowledge about the use of threat modeling within your organization. Furthermore, you must be able to express yourself clearly in Dutch and/or English. Methods. If you agree to participate in this study, you will be sent an invitation to take part in a face-to-face interview in the first half of 2022, at your place of work. During this interview, you will be asked about your experiences with and knowledge about the use of threat modeling within your organization. The audio of this interview will be recorded and later transcribed for further analysis.		
If possible, this interview will be complemented with a short analysis of existing threat modeling documentation on a concrete project. If you wish to participate in this document analysis, you will be asked to bring one or more documents related to threat modeling, which will be examined by the researchers directly after the interview, under your supervision. The goal of this document analysis is to gain concrete insights into how threat models are composed and documented in practise.		
You may be contacted after the interview for any follow-up questions.		
Voluntary participation and withdrawal. Participation in this study is completely voluntary, and can in no way be obligated, in particular by your manager or organization. You can, at any point in time (before, during and after the interviews), withdraw from the study without any negative consequences for you or your organization, and you do not need to provide a reason for doing so. If you decide to withdraw from the study, no additional data will be recorded about you. Advantages, rewards and compensation. If you wish so, you will be informed about the general results of this study. You will also receive a small gift (20 EUR consumption coupon) at the start of the interview to thank you for your participation. Other than that, there will be no further rewards or compensation, and you will gain no other personal advantages from this study. Sharing your experience will, however, contribute to a clearer view on the use of threat modeling in practice, as well as its potential obstacles and solutions. The results of this study shall be used to share these insights and experiences with a broader public, which may lead to better recommendations and techniques to secure systems using threat modeling. Risks. The risks associated with participation in this study are limited. As a participant, you are always free to not answer certain questions, for example for confidentiality reasons. In addition, the risks of a potential data leak are minimized through technical and organizational measures, including disconnecting the identity information and replacing organization-specific elements. This is discussed in more detail below.		
1/2		

7. **Confidentiality.** All of your personal data, interview data, and other data gathered in the context of this study will be handled confidentially, and will not be shared with others (in particular, your manager or employer).

This data will be disassociated from your identity using a random pseudonym (e.g., S4X). A document linking your identity to the pseudonym will be stored separately and securely by the researchers, and will never be shared with others. This document will be destroyed once it is no longer required for the study. Your identity will never be revealed in publications or reports on the results of this study, and any element referring to your organization will be anonymized. The audio recordings of the interviews will be destroyed once they have been transcribed.

Once the study has concluded, the general results as well as the redacted transcriptions of the interviews will be shared with the study sponsor if you agree to this. Your data, as well as elements which specifically identify your employer, will be anonymized in the shared transcriptions.

8. **Data processing.** Public interest is used as the legal basis for processing your data in accordance with the GDPR. This means that this study will result in an increase of knowledge and insight which benefits society either directly or indirectly. As a data subject, you have a right of access (art. 15), right to rectification (art. 16), right to erasure (art. 17), right to restriction of processing (art. 18), and right to object (art. 21). The involved researchers can be contacted to exercise your rights (see contact information below).
9. **COVID-19.** The interviews will be organized in compliance with COVID-19 measures and constraints applicable at the time of the interview (e.g., self-testing, ventilation, ...) in order to limit the risk of infection. If the interview unexpectedly cannot take place physically, even after a delay of at most 2 months, a remote interview will be conducted instead.
10. **Conflicts of interest.** There are no conflicts of interest to be declared for this study.
11. **Contact information.** You can contact the researchers of this study at any time for further information. All involved researchers are affiliated with the Department of Computer Science, research group Distributed and Secure Software (DistriNet), KU Leuven (<https://distrinet.cs.kuleuven.be>).
- | | | |
|-------------------------------|--|-------------------|
| - ir. Stef Verreydt | stef.verreydt@kuleuven.be | (researcher) |
| - Dr. ir. Laurens Sion | laurens.sion@kuleuven.be | (researcher) |
| - Dr. ir. Koen Yskout | koen.yskout@kuleuven.be | (head researcher) |
| - Prof. Dr. ir. Wouter Joosen | wouter.joosen@kuleuven.be | (promotor) |

The ethical aspects of this study were examined and approved by the Sociaal-Maatschappelijke Ethische Commissie (SMEC) ("Social and Societal Ethics Committee") of KU Leuven. Any ethical complaints or concerns can be reported to them via smec@kuleuven.be.

This is a copy of the informed consent form that was used for this study.

	GEDISTRIBUEERDE EN VEILIGE SOFTWARE (DISTRINET) DEPARTEMENT COMPUTERWETENSCHAPPEN CELESTIJNENLAAN 200A 3001 HEVERLEE, BELGIË	
Invitation to participate in study The use of system-centric threat modeling in Dutch organizations.		
Informed consent		
You are invited to participate in the study "<i>The use of system-centric threat modeling in Dutch organizations</i>", carried out by researchers from KU Leuven (Belgium) and commissioned by the Dutch National Cyber Security Centre (NCSC), part of the Dutch Ministry of Justice and Security.		
Additional information regarding the study will be provided in the attached information sheet.		
By signing this form you certify the following:		
➤ I understand what is expected of me during this research. ➤ I know that I will participate in the following activities: ○ An interview spanning around one hour, concerning threat modeling in your organization. ○ (Optional) Guiding a document analysis directly after the interview. ○ Potential follow-up questions via e-mail or via audio or videocall. ➤ I know that my participation may be associated to risks or discomforts: ○ Questions could inadvertently involve potentially sensitive information about you or your organization. You are always free to not answer a question. ○ To limit the risk of a potential data breach, technical and organizational measures are taken, including disconnecting your identity information and removing organization-specific references. ➤ I or others can benefit from this research in the following ways: ○ The study results will be shared with you once the study is concluded. ○ The NCSC will gain insights into how the organizations within their target audience currently apply and experience threat modeling. Based on these insights, they can then better tailor their advice to the existing needs. ○ The study results will be published so that other researchers can build on them. ➤ I know that I will receive a small gift (20 EUR consumption coupon) for participating in the study, and that there are no other rewards or compensation associated with my participation. ➤ I know that my participation in this study is voluntary. I have the right to withdraw from the study at any point in time. I do not need to give a reason for doing so, and I know that this will not have negative consequences for me.		
Public interest is used as the legal basis for processing your data in accordance with the GDPR. Withdrawal from the study therefore implies that previously collected data can still be legally involved in the study and does not have to be deleted by KU Leuven.		
➤ The results of this study can be used for scientific goals and may be published. My name will not be published. The anonymity and confidentiality of the data will be protected in all stages of the research. ➤ An audio recording of my interview is made, which is transcribed afterwards. After this transcription, the recording will be deleted.		
1/2		



- A redacted transcript of my interview will be shared confidentially (without references to me or my organization) with the commissioner (the NCSC) if I agree to this.
☐ I agree that this information is shared with the commissioner.

- ☐ I would like to be informed on the results of this study. The researcher may contact me for this purpose via the following email address:

☐ I do not wish to be informed on the results of this study.

- For questions and for the execution of my rights (access to my data, rectification of the data, ...) after my participation I know that I can contact:

- ir. Stef Verreydt	stef.verreydt@kuleuven.be	(researcher)
- Dr. ir. Laurens Sion	laurens.sion@kuleuven.be	(researcher)
- Dr. ir. Koen Yskout	koen.yskout@kuleuven.be	(head researcher)

De promotor of this research is:
Prof. Dr. ir. Wouter Joosen
wouter.joosen@kuleuven.be
Celestijnenlaan 200A – bus 2402
3001 Leuven
België
+32 16 32 76 53

The promotor and all involved researchers are affiliated with the Department of Computer Science, research group Distributed and Secure Software (DistriNet), KU Leuven, Belgium (<https://distrinet.cs.kuleuven.be>).

More information with regard to privacy in research can be found at <https://kuleuven.be/privacy/en/>. With further questions about privacy issues I can contact the data protection officer: dpo@kuleuven.be.

- In case of complaints or other concerns with regard to the ethical aspects of this research I can contact the Social and Societal Ethics Committee of KU Leuven: smec@kuleuven.be.

I have read and understood the information in this document and I have received an answer to all my questions regarding this research. I give my consent to participate.

Date:

Name and signature of the participant

Name and signature of the researcher

C

Interview Guide

This is a copy of the interview guide used by the interviewer, detailing the topics and questions to cover. Note that, as the technique of responsive interviewing [Rubin and Rubin, 2011] was used, during the interview these topics and questions were not necessarily dealt with in this order, nor using the exact phrasing from below.

C.1 Introduction

- Introduction of the interviewer
- Verify whether informed consent form was clear and has been signed
- Repeat agreements with respect to recorded material

C.2 Demographics

These questions do not directly address the research questions but are used to contextualize and compare the responses of the different interviewees.

1. Personal
 - 1.1. Description of the current role
 - 1.2. How long in the current role
 - 1.3. Size of the team and goal of the team
2. Organization
 - 2.1. Number of employees
 - 2.2. Number of employees with security or privacy role
 - 2.3. Kind of applications and systems

C.3 Threat Modeling

Every question mentions the main corresponding research question. Answers may also contain aspects of other research questions, though.

3. General
 - 3.1. What does threat modeling mean to you? **(RQ1)**
 - 3.2. Where/when/how/from whom did you learn threat modeling? What are your first experience with threat modeling? **(RQ3)**
 - 3.3. Why do you threat model? What are the expected outcomes of threat modeling? **(RQ1)**

- 3.3.1. To which extent does your organization support you to perform threat modeling?
- 3.4. How would you describe the relation between threat modeling and risk management within the organization? **(RQ1)**
4. Process/execution
 - 4.1. How is threat modeling performed in your organization? **(RQ3)**
 - 4.2. Which kinds of models or abstractions are used? **(RQ3)**
 - 4.2.1. When and by whom are models created?
 - 4.2.2. Is the quality or correctness of the used models assessed and if so, how?
 - 4.3. Who is involved in during a threat modeling session? **(RQ2)**
 - 4.3.1. Who takes the initiative?
 - 4.3.2. What is role and involvement of management?
 - 4.3.3. To which extent are operational teams involved?
 - 4.3.4. Is threat modeling delegated to external organizations (e.g., consultants)? If so, is there an intention to eventually perform this in-house?
 - 4.4. How are models analyzed? **(RQ3)**
 - 4.5. How long does an analysis take? **(RQ3)**
 - 4.5.1. Are there any stop- or correctness-criteria?
 - 4.6. Which steps are explicitly documented and how? **(RQ3)**
 - 4.7. What happens with the results? **(RQ1)**
5. Experiences
 - 5.1. Do you experience any limitations with how threat modeling is currently performed? Steps that are laborious? **(RQ4)**
 - 5.1.1. How were/are these dealt with?
 - 5.2. What do you consider to be the biggest successes with threat modeling in your organization? **(RQ4)**
 - 5.3. How satisfied are you with the process, timing, efficiency? **(RQ4)**
 - 5.4. How satisfied are you about the follow-up on the results? **(RQ4)**

C.4 Closing/debriefing

- Any other elements not mentioned?
- Questions from interviewee?
- Repeat agreement regarding confidentiality.
- Ask for possibility of further follow-up.
- Provide contact information.

