



Universiteit
Leiden
The Netherlands

Nobody knows the risks i have seen: evaluating the gap between risk analysis and security operations

Schiele, N.D.; Egberts, J.; Gadyatskaya, O.

Citation

Schiele, N. D., Egberts, J., & Gadyatskaya, O. (2024). Nobody knows the risks i have seen: evaluating the gap between risk analysis and security operations. *2024 Ieee European Symposium On Security And Privacy Workshops (Euros&pw)*, 473-483.
doi:10.1109/EuroSPW61312.2024.00059

Version: Publisher's Version

License: [Licensed under Article 25fa Copyright Act/Law \(Amendment Taverne\)](#)

Downloaded from: <https://hdl.handle.net/1887/4214578>

Note: To cite this publication please use the final published version (if applicable).

Nobody knows the risks I have seen: Evaluating the gap between risk analysis and security operations

1st Nathan Daniel Schiele

LIACS, Leiden University
Leiden, The Netherlands
n.d.schiele@liacs.leidenuniv.nl

2nd Jari Egberts

LIACS, Leiden University
Leiden, The Netherlands
j.t.egberts@umail.leidenuniv.nl

3rd Olga Gadyatskaya

LIACS, Leiden University
Leiden, The Netherlands
o.gadyatskaya@liacs.leidenuniv.nl

Abstract—Maintaining cyber security is becoming increasingly difficult, forcing organizations to turn to security service providers. These security service providers use teams with different skills to provide security outcomes for their clients. These teams often consist of risk analysts (RA) and security operations (SO) experts. It seems natural to expect that these complementary team members will inform each other about relevant security threats. Yet, this is not always the case.

We have conducted a qualitative case study to examine the communication process between RA and SO within a large cyber security service organization. We conducted a series of semi-structured interviews with security experts ($n = 11$) who are either in RA or SO to find what tasks they regularly perform, how they communicate with other teams, and what areas of improvement they would find useful. The transcripts of these interviews were analyzed with template analysis.

We found that both RA and SO feel a communication gap, which they sometimes attribute to differences in teams' backgrounds. Use-case implementation has been identified as a potential major area of improvement as a means of communication between RA and SO.

Finally, we also observed that some issues are related to communication with clients and low levels of maturity in client's processes.

1. Introduction

With the ever increasing number cyber threats, it is increasingly difficult for organizations to meet their own cyber security needs. This situation is becoming more common, and the solution these organizations turn to are cyber security service providers, firms who specialize in ensuring security outcomes for their clients.

Considerable research has been devoted toward the people, technology and systems that enable these security service organizations to function optimally. Providing security services is a complex, multifaceted operation, requiring interconnected systems and people from many different backgrounds and skillsets to collaborate in order to ensure the best possible security for clients. For example, a typical service offer from such firms would include information security risk management and security monitoring. Information security risk management personnel are responsible for eliciting and managing cyber risks, providing advice on security strategy, including

implementation of security controls, and possibly compliance and audit-related tasks [9], [10]. Managed security operations centers (SOCs) are generally responsible for ensuring security via technical means such as detection and incident response [23]. Thus, SOCs are usually staffed with technical analysts who use an array of tools and work with threat intelligence data sources [7], [21]. Risk analysts are often client-facing, whereas SOC analysts tend to be internal technical personnel. Teams of risk and SOC analysts work together to provide effective security services for their clients. It would seem reasonable to expect that there is a constant communication going on between these two teams, to inform each other about the identified risks, implement the relevant detection mechanisms via use-cases, and observe that the registered security risks are relevant to the current threat landscape.

However, based on several recent studies, this is not necessarily the case. Kokulu *et al.* examined SOC workflows and found various issues [17]. For example, they found that limited visibility into a client's network often hinders the SOC in its duties, and that there is difficulty when communicating with teams outside of the SOC [17]. Webb *et al.* posed that security risk management should be informed by situational awareness and current threat landscape of the organization (what would be delivered by a SOC team) [25]; yet, Ahmad *et al.* found that, while communication between incident response and information security management teams is beneficial, few organizations enable integration between these security teams [1], [2].

There has been significant research into the challenges faced by SOCs or risk analysis as separate teams. Critically, these teams are not intended to work alone, so examining the needs of SOCs or risk analysts without examining the broader context does not enable understanding the challenges and issues facing security service organizations as a whole. Thus, in our study we specifically examine the communication process between risk analysis and security operations to find what key challenges are. To the best of our knowledge, ours is the first work to examine the gap between these two teams. We have done this in the context of a single security service firm; but, based on the existing literature [1], we believe that the lessons learned are relevant to any large organization, which has security teams that include both cyber risk management and security operations.

Our research is fundamentally guided by the question *What is the gap between risk assessment and security*

TABLE 1. STUDY PARTICIPANTS AND THEIR YEARS OF EXPERIENCE. OUR CLASSIFICATION OF THESE PARTICIPANTS INTO RISK ASSESSMENT (RA) AND SECURITY OPERATIONS (SO) IS BASED ON THEIR SELF-REPORTED ROLES IS INCLUDED IN THE ID COLUMN.

ID	Title/Role	Experience (# years)
RA-1	Security Operations Manager SOC	4
RA-2	Manager Security Advisory	6
SO-1	Security Operations Manager SOC	5
SO-2	Team Lead Security Operations	5
RA-3	Program Manager Cyber Security	5
RA-4	Business Security Consultant	6
SO-3	Security Operations Manager SOC	5
SO-4	Manager Blue Team	4
RA-5	Director Business Security	11
SO-5	SecOps Tech Lead SOC	13
SO-6	Security Operations Manager SOC	3

operations?. We have broken this overarching research question down into multiple more specific research questions:

- RQ1** How does the work done by risk assessment and security operations relate to each other?
- RQ2** How is the information exchange between risk assessment and security operations structured?
- RQ3** How can gaps between risk assessment and security operations be improved?

We apply template analysis (a type of thematic analysis) to answer these research questions in a qualitative study, where we interviewed 11 practitioners from risk assessment and security operations teams in a security service firm. The remainder of this paper is structured as follows. In Section 2 we introduce the methodology of our study. Section 3 summarizes results of our template analysis. Section 4 presents the study findings and answers the research questions. We then discuss limitations of the study (Section 5) and the related literature (Section 6). Finally, we state our conclusions in Section 7.

2. Methodology

Our research used a series of semi-structured interviews with both security operations and threat assessment staff within a large security service company. Semi-structured interviews are a common research practice within the security industry, as many prior works have conducted research with similar means [2], [12], [13], [17]. These interviews allow us to gain insight into the work and communications of risk assessment and security operations personnel. Ultimately, our research questions are centered around the actual work being performed and the communications necessary to both perform that work, and to communicate results. The only way to address these research questions is to speak to practitioners who have direct, relevant experiences.

The study design and execution plan was approved by the Ethics Review Board at our institution. Participants were asked if they consented to be interviewed for the study, and for their interview responses to be recorded, transcribed, and analyzed with selected, anonymized, quotes to be published.

2.1. Study participants

In all, we interviewed 11 participants, which has been shown by Hennik and Kaiser to be a sufficient number of interviewees to reach information saturation [14]. Five interviewees work as risk assessment (RA) personnel, with experience ranging from 4 years to 11 years in industry. Six interviewees work in security operations (SO), with experience ranging from 3 years to 13 years. Information concerning the participants can be found in Table 1. Note there that participant RA-1 performs risk assessment tasks for clients, while they are embedded in the SOC team. They remarked that this is an unusual nuance.

All participants work for the same cyber security service company in the Netherlands. This consulting firm works with many different corporate clients both domestically and internationally. It provides different types of security services to clients, including security and privacy risk management, security operations and incident response, pentesting, compliance certification and audit, and others. The firm does not specialize in any particular sector or application domain. Participants were recruited personally based on their willingness and availability.

2.2. Interview set-up and questionnaires

The interviews were conducted in April-May 2021. Due to the COVID-19 pandemic, ten of the eleven interviews were conducted via Microsoft Teams. One interview occurred in-person. All interviews were audio-recorded for later transcription. Ten of the experts were interviewed in Dutch and one was interviewed in English. All transcripts were later translated to English for coding and analysis. The translation was checked to confirm that the translated text was true to the original meaning offered during the interviews.

We used two sets of questions: one for those who work in threat assessment and another for those who work in security operations. These question sets overlap on questions that are relevant to both roles, and differ to concerning role-specific questions. Both questionnaires are provided in Appendix A.

The questionnaires start with demographic questions, to enable us to check our classification of employees as either primarily risk assessment or security operations personnel. We then ask all interviewees about the services their teams provide, both in terms of risk assessment and security operations. From there, the interviews diverge based on the role the interviewee fulfils. We ask risk assessment (RA) interviewees about the frameworks they use for risk assessment, how the long-term monitoring of risks is performed, and how this information is communicated to the SOC, as well as if any information is received from the SOC. We ask security operations (SO) interviewees about the threat intelligence and models they receive, the information about risks, the implementation of use cases, and communications they receive from risk assessment. We ask both about their communications with other departments, and what they believe could be improved. These questions allow us to gain insight into the work being done by both risk assessment and security operations. They are specific enough to guide interviewees to address the concepts that we wish to research, and

specifically, what tasks are being performed and how are these related to the tasks being done by others, which correspond to RQ1. We also gain insight into the communications between experts performing distinct tasks, which addresses RQ2 and RQ3.

2.3. Template analysis

To analyze the data, we performed a template analysis. Template analysis was first introduced by King in 1998 as a form of hierarchical thematic analysis in which codes are classified into themes and then applied to the data set [16]. This form of analysis was further expounded upon by Brooks and King to clarify that templates start from an initial form and are expanded and adjusted through further coding, with the final template being used to guide analysis [8]. The application of template analysis within security research is well-founded, and it has been applied in the context of SOC by, for example, [3], [4], [21].

Open coding. In our study, we used an initial round of open coding of the transcripts to gain insight into what the data was saying. This unstructured initial open coding was used to develop themes in order to establish the template which would be applied to the data. The themes were used to create the template entries for risk analysis and security operations (template items 1 and 2 respectively). Two members of the research team independently conducted the coding to establish inter-coder agreement.

The initial open coding process was completed by two researchers independently, who then discussed their coding, and aligned and merged code books. A total of 607 individual codes were assigned across all interviews during this initial process, with 41 unique codes present in at least four of the interviews; we considered these to be the starting point for the themes to be used in the template. The open coding process generated many different codes of work tasks, from which we selected only the most frequently used codes for the template described in the following section.

Template. After the open coding and themes identification, we developed and applied a template. As is standard practice with template analysis [8], the template was adapted and expanded as the template was being applied to the interview transcripts [16]. We created a total of 4 templates before arriving at our final (fifth) template. This template is provided in Appendix A, in Table 4.

One advantage of template analysis is the flexibility afforded toward the application of the template. As the template is flexible, it allows for expansion during the process of coding, which some coding processes do not allow [24]. Additionally, templates can be organized on many levels, with no maximum or minimum number of abstractions [16]. In our cases, some of our themes have multiple levels of abstractions with many resulting subcodes, which are directly applied to quotations, while some themes only have a minimal number of directly applied codes.

Specifically, we designed our template to take advantage of overlapping codes. Codes are designed to be applied in conjunction to provide new meanings and to modify other codes. The flexibility offered by the ap-

plication of template analysis allows us to encapsulate significant meaning within a compact codebook [16].

Ultimately, we found four major themes which are organized into our template. These themes are explained below.

Theme: Work Tasks. The codes and subcodes within this theme concern the actual work tasks being performed by our interviewees. For example, if an interviewee talks about implementing use-cases in their work, then this quotation would be coded with *Use-case implementation*. We have classified most of the work tasks into either risk assessment or security operations tasks. These codes were created from the open coding process discussed earlier. Additionally, we have *Other* codes to cover work tasks that don't have explicit codes included in the template. The *Other* code is also used when a participant refers to a task in a generic manner. Finally, the work task *Threat Intelligence* is not classified into either risk analysis or security operations, as this task can be done by various security personnel at many stages. This theme specifically addresses RQ1 as we are exploring the actual work being performed by the interviewees.

Theme: Process. The codes and subcodes within this theme refer specifically to the manner and style in which the work is performed. They can be thought of as modifiers to the **Work Task** theme codes. For example, interviewee SO-2 had the following quote: “*When a new customer starts, they start doing risk analysis with the customer*”. This quotation would be classified as the work task *Risk Management* as risk analysis is included as a part of risk management. From there, we would add the process codes *Process:Who:Other* as this is a work task being performed by someone other than the interviewee's own team, and *Process:When:Client Intake* as this work task is performed when the client is onboarded. In this way, we can classify the manner in which specific work tasks are performed without needing individual, unique codes to describe this, which reduces the overall number of codes.

We distinguish between the structure and scope of a process. Structure refers to the internal steps of a process: how the task is being performed. Scope refers to when the process is considered complete or the responsibility of another team. Put shortly, structure is a description of the internal aspects of a task, while scope is a description of the task boundary. Overall, this theme addresses RQ1 as we explore the manner in which work tasks are performed and how these work tasks relate to the work being done by other teams.

Theme: Communication. The codes and subcodes within this theme concern the communications between various participants both within and external to the organizations of the interviewees. Communication can be thought of as a process; however, as communication is a specific research question of ours, we have created a separate theme specifically for communication. Additionally, communication has some specific aspects that are not present in the process theme, such as directionality. As with the **Process** theme, these codes are designed to be used in conjunction with each other. For example, if a quotation was coded with *Communication:When:Scheduled*, *Communication:Participant:Client*, and *Communication:Direction:Sent*, we know that this

quotation refers to a regularly scheduled communication being sent to the client. Once again, we combine codes in this manner to allow us to encapsulate more meaning in a compact codebook.

We have created a distinction between content and structure for communication. This is similar to the distinction between process structure and scope. Communication structure refers to how information is organized in a communication: is the communication in a specified structure, such as a report, or is the communication entirely up to the transmitter, such as a text message? Communication content refers to what is being transmitted, whether all information is being transmitted or only some information. Otherwise put, content refers to the information itself, while structure refers to the organization of that information. This theme addresses RQ2 as we are exploring how different teams communicate with each other.

Theme: Perspective. This theme concerns the perspective of the interviewee on various concepts. *Improvement* refers to suggested improvements from the expert, with classifications for what kind of improvement is being suggested. *Issues* refers to issues identified by the interviewee, with subcodes providing information about the perceived source of the issue. *State* refers to whether the interviewee is describing the current state of what the organization is doing, what they did previously, or what they would ideally do. This theme addresses RQ3 as it aims to capture the perceived communication gap between the two teams and how it can be improved.

3. Template analysis results

After application of template analysis, we conducted a statistical analysis on the codes assigned to the interviews, in line with other works which use a similar methodology [3], [4]. In total, 432 individual quotations were marked across the 11 interviews. Each of these quotations had between 1 and 8 codes applied to them. As discussed in Section 2.3, codes were applied in conjunction in order to introduce a varied set of meanings.

In Table 2, we see the total number of quotations per interview. Additionally, we see the number of quotations in each interview marked with the code *Process:Who:Own*, which is an indicator that the quotation concerns a work task that the interviewee or the interviewees' team does themselves. The remaining columns show which work tasks specifically each interviewee has referred to as their own work task, and the number of quotations which were coded as such.

In Table 3, we see the work tasks codes per interview. These codes are additionally marked as either *Process:Who:Other*, *Process:Who:Own*, or *Process:Who:Unclear*, which is an indicator of who performed this process, with *Own* meaning that an interviewee themselves or their team conducted the work task, while *Other* indicating that it was another team, another organization or the client themselves that performed that work task. Only quotations in which an interviewee discussed one work task in both the context of their own team and another organization are included in this table. As such, the work task of *Alert Evaluation* has been omitted, as this never had both *Process:Who:Other* and

Process:Who:Own applied to quotations assigned with *Alert Evaluation* within the same interview.

We have represented the data as bubble charts to visualize how clearly the participants report their workflows and interactions with other teams, and how much improvement need did they voice. The bubble charts focus on the process (both scope and structure) and the communication structure. In these bubble charts each bubble is one interview, and the size of each bubble represents the number of quotations in this interview where participants voiced the need for improvement.

Figure 3 shows a bubble chart, where the size of the bubble indicates the number of quotations which mention a need for improvement with respect to *process*. There are only nine bubbles, as two of the eleven interviews did not mention process improvements needed. The placement of the bubble indicates how many quotes were coded with *Process:Structure:Framework* (meaning that there is a clear structure of the process) and *Process:Structure:Ad-hoc* (indicating that there is a lack of structure in the process and it is being done largely ad-hoc), provided by the x and y axis, respectively. The color of the bubbles indicates how many quotations were coded with *Process:Structure:Unclear*.

The intention is to illustrate any potential correlative relationship between interviewees making statements about their work processes being structured (*Process:Structure:Framework*) and process being unstructured (*Process:Structure:Ad-hoc*). Additionally, we also include the information of whether or not interviewees made statements to the effect of the structure (or lack thereof) of their work processes being not clearly understood by the interviewee, and the number of statements an interviewee made concerning improvements needed with respect to communication. We can see a slight positive correlation between number of quotes coded as unstructured processes and number of quotes coded as structured processes, which would indicate that there is no correlative relationship between these two, merely that if an interviewee spoke about their work process, they would roughly mention a structured and unstructured processes similar amounts of time. These quotes are likely more reflective of how often a participant spoke about their work process.

Figure 1 shows a bubble chart, where the size of the bubble indicates the number of quotations which mention a need for improvement with respect to *communication*. The placement of the bubble indicates how many quotes were coded with *Communication:Structure:Framework* and *Communication:Structure:Ad-hoc*, provided by the x and y axis, respectively. The color of the bubble indicates the number of quotations marked with *Communication:Structure:Unclear*.

Finally, Figure 2 shows a bubble chart, where the size of the bubble indicates the number of quotations which mention a need for improvement with respect to *process*, but in contrast to Figure 3, we are looking at the process scope here. There are only nine bubbles, as two of the eleven interviews did not mention process improvements needed. The placement of the bubble indicates how many quotes were coded with *Process:Scope:Limited* (meaning that the interviewee reports a clear delimitation between different tasks) and *Process:Scope:Expansive* (referring to


TABLE 2. “OWN” TASKS, TASKS DESCRIBED AS BEING COMPLETED BY THE INTERVIEWEE OR THEIR TEAM, DESCRIBED PER INTERVIEW AND PER CLASSIFICATION

ID	Total	Own	Alert Eval.	Monitoring	Use-case Impl.	SO: Other	Risk Manag.	RA: Other	Threat Intel.
RA-1	27	10	0	1	2	0	4	4	0
RA-2	54	25	0	3	9	1	16	6	2
RA-3	30	16	0	1	4	0	13	6	1
RA-4	47	25	0	5	2	4	22	2	0
RA-5	38	23	1	11	3	5	14	2	0
SO-1	41	14	0	3	3	4	2	1	2
SO-2	29	10	5	2	1	1	0	0	2
SO-3	64	16	4	3	5	5	1	1	1
SO-4	41	15	0	8	7	0	5	0	3
SO-5	34	17	1	5	4	6	4	2	0
SO-6	26	14	5	3	0	8	2	2	1

TABLE 3. CO-INCIDENCE OF WORK TASK CODES WITH *Process:Who* CODES (OTHER - OWN [- UNCLEAR]). “OTHER” INDICATES SOMEONE OTHER THAN THE INTERVIEWEE OR THEIR TEAM, “OWN” INDICATES THE INTERVIEWEE OR THEIR TEAM, AND “UNCLEAR” INDICATES THAT IT IS NOT CLEAR FROM CONTEXT WHO PERFORMED THE TASK.

ID	Risk Manag.	RA:Other	Threat Intel.	Monitoring	Use-case Impl.	SO:Other
RA-1	2 - 2	1 - 4	0	3 - 1 - 1	1 - 1 - 3	0
RA-2	4 - 14 - 1	1 - 6 - 1	4 - 1 - 2	8 - 3	0	0
RA-3	0	0	0	1 - 1	3 - 4 - 2	1 - 1
RA-4	4 - 22	0	2 - 4	3 - 5	1 - 2	0
RA-5	0	1 - 2	0	0	0	0
SO-1	1 - 2 - 1	0	0	0 - 2 - 1	3 - 3 - 1	0
SO-2	0	0	0	0	3 - 1	0
SO-3	1 - 1	0 - 1 - 1	1 - 5	0	1 - 5 - 1	0
SO-4	0 - 4 - 3	0	0	0	4 - 7	0
SO-5	0	1 - 2	0	0	0	0
SO-6	0 - 1 - 1	0	0	0	0	1 - 1


Figure 1. Distribution of quotations related to communication structure. Each bubble is an interview. The x -axis corresponds to number of quotations coded with *Communication:Structure:Framework*; the y -axis represents the number of quotations coded with *Communication:Structure:Ad-hoc*. The color and size indicates the number of quotations where the interviewee voiced that their communication structure was unclear or that the communication was in need of improvement, respectively.



a case where task boundaries are blurred), provided by the x and y axis, respectively. The color of the bubbles indicates how many quotations were coded with *Process:Scope:Unclear*.


In Fig. 3, we can see that there are fewer quotations mentioning an improvement to process with a higher number of quotations related to a well-structured process

Figure 2. Distribution of quotations related to process scope. Each bubble is an interview. The x -axis corresponds to number of quotations coded with *Process:Scope:Limited*; the y -axis represents the number of quotations coded with *Process:Scope:Expansive*. The color and size indicates the number of quotations where the interviewee voiced that their process scope was unclear or that the process scope was in need for improvement, respectively.



(with framework). We see the same relationship in Figures 1 and 2. This could suggest that if a process or communication is well-structured or has a clear scope, then fewer improvements are perceived to be needed. With quotations tagged with *Unclear*, we cannot see a relationship to the incidence of other quotations. A trend we see most strongly in Fig. 1 is that interviewees either had many *Communication:Structure:Framework* or *Communication:Structure:Ad-Hoc* quotes, but not both.

Figure 3. Distribution of quotations related to process structure. Each bubble is an interview. The x -axis corresponds to number of quotations coded with *Process:Structure:Framework*; the y -axis represents the number of quotations coded with *Process:Structure:Ad-hoc*. The color and size indicates the number of quotations where the interviewee voiced that their process structure was unclear or that the process structure was in need of improvement, respectively.



Indeed, these codes are mutually exclusive. These results suggest that participants either had a communication structure that was defined, followed and could be described, or that the structure of their communication was not defined or not able to be described. We see similar results in Fig. 2.

4. Study findings

4.1. Use-case handover

We can see in Table 3 that the *Use-Case Implementation* work task has a significant number of interviewees whom we classified as both risk assessment and security operations personnel, who indicated that *Use-Case Implementation* was both an *Other* and an *Own* task. As was mentioned in Section 3, each quotation that is coded with a work task is also coded with a *Process:Who* code to indicate who is performing that work task. As the distribution of codes between *Other* and *Own* is more equal, it is probable that interviewees are describing a task shared by multiple teams. We see this relationship most strongly with *Use-Case Implementation* with four interviews (RA-3, RA-4, SO-1, and SO-4) showing significant overlap between *Other* and *Own* processes. This indicates the *Use-Case Implementation* is being performed to a significant extent by multiple teams.

Interviewee SO-5 provided some insight into this by saying “when it comes to the way we normally do identify what use cases to implement in a SOC, we do it in a different way to the way that risk do it”. This indicates that the use-case implementation that is being performed by one team is inherently different from the use-case implementation being performed by another. Interviewee RA-5 added “we have your compliance controls and a description, but then you have your use cases and your technical description. I feel like in the middle there needs to be some kind of language that matches the two together”.

Overall, these results are not unexpected. Use-cases are often used in between risk analysis and security operations, being the deliverable following a risk analysis process or the starting point for conducting security operations [6]. As such, seeing overlap in the descriptions of the work tasks being performed by risk analysis and security operations personnel specifically in this work task follows from our understanding of use-case implementation. At the same time, as this is a shared task, it can become a natural opportunity for both teams to come together in order to build a shared understanding.

4.2. Cross-technical background communication

The most commonly identified source of issues concerned communication specifically between people of different background. 8 of the 11 quotes which concerned issues of communication were related specifically to this subset of communication issues. As interviewee SO-5 said about a risk analyst, “they don’t really know what we have done and if it is sufficient, because they don’t understand well enough the technical details to really know if that is sufficient or not”.

This is not solely an issue with communications between different teams within the organization. Half of the quotations concerning issues with clients touched on the difficulty of explaining complex technical concepts to clients. Threat models can be an excellent way to explain complicated risk concepts [20], however, no interviewee mentioned use of threat models. Interviewees did discuss methodologies for assisting with closing the communication gap between people with different technical backgrounds, many using use cases for this exact purpose. For example, interviewee SO-1 stated, “there are a lot of those little details behind such a use-case that make it easier for us to help the analysis and also make it easier for the customers to actually pick up on the useful things”. This indicates that use cases are being used as a means to bridge the communication gap between people with different technical backgrounds.

We use this distinction as a proxy to examine the communication between risk analysis which is traditionally less technical and security operations, which is traditionally more technical. By examining communication between differing technical backgrounds, we can gain insight into how the different teams of RA and SO communicate with each other, as the background of these teams is radically different.

4.3. Frameworks

We were particularly interested in frameworks supporting RA and SO teams, and we included questions about frameworks in both questionnaires. Over the course of the interviews, all interviewees talked about frameworks in some capacity.

In terms of specific frameworks being used, interviewees RA-1, SO-1 and RA-5 mentioned MAGMA [22], a framework for defining use-cases. All interviewees, except SO-2, SO-4 and SO-6, mentioned the ISO 27000 series [10]. However, no interviewee said that the used frameworks were implemented as described. The common

theme across the descriptions of framework implementations was that these frameworks are not implemented exactly as designed. Interviewee RA-1 stated *“that’s why I was just maybe being a little vague about whether or not we have a framework. It’s actually an mix of different things that have a whole thought process behind them, but these days we don’t implement it according to the letter of the framework. We’ve found a kind of method in which we take the best things from the frameworks, and that we roll it out that way.”*. Overall, the sentiment from the experts interviewed is that they use what they find useful.

What we do notice is that participants who are able to describe their internal framework more clearly have a much easier time answering our questions. Additionally, there is evidence from Figures 3 and 1 that the increasing incidence of *Structure:Framework* codes correlates with fewer *Improvement* codes, which may indicate that a more structured communication or process is related to a smaller perceived need for improvement. It is an open question why published frameworks are not being used more precisely. Are the frameworks too coarse-grained, or are practitioners overwhelmed by too many frameworks?

4.4. Client expectation management

The interviewees identified many challenges within their work, particularly concerning working with clients. Much of these issues surround client expectations, with interviewee RA-2 saying *“I can’t take over your responsibility, because I don’t own the risk”* when talking about client expectations of risk management. Interviewee RA-3 spoke about the expected monitoring capabilities of the organization by saying *“we didn’t actually have the log resources to be able to do that”*, as they were working with clients who expected them to be able to monitor more than what was feasible. Interviewee SO-3 had a similar issue saying *“a use-case is proposed that you really just don’t have the log sources for”*. On one hand, this can be a technology issue, where if the client organization had a greater capacity through automation, then the service company could potentially meet their clients’ expectations. However, it is also potentially a question of expectation management.

Interviewee SO-3 further said, *“If you’re at maturity level 0 then it’s up to a good security provider to first say to the customer ‘sorry’”*. This was said in reference to an example of a new client asking for monitoring when they had no other security solutions in place. The interviewee spoke about how clients often have difficulty understanding their own security needs. Interviewee SO-5 stated something similar about clients, saying *“I find that we don’t always do enough risk-based work. [...] I don’t know if that’s because the client isn’t always ready for it”*, in reference to clients who want to move straight into security operations without first following a risk analysis process. Overall, there seems to be a strong need for effective client expectation management, which would be an interesting area of future research.

4.5. Answers to research questions

To summarize, those are the answers to our research questions:

RQ1. In Table 2 we can see the tasks performed by both RA and SO. Further in Table 3 we can see the division of these tasks. Specifically we see that use-case implementation is a task that is heavily related to both RA and SO. We can see that apart from use-case implementation, there is very little overlap in terms of described work processes, which is evidence to suggest that the division is well defined and understood.

RQ2. As discussed in Section 4.1, we see that use-case implementation is a key means by which information is exchanged from RA to SO, as shown in Table 2 that all participants see use-case implementation as included in an aspect of their work. Further research is needed to fully understand the importance of use-case implementation and to develop use-case implementation into a more effective framework for communication between RA and SO.

RQ3. As discussed in Section 4.2, we find that there remains a communication gap between RA and SO, and both sides feel this communication disconnect. Specifically, we have identified issues related to different backgrounds of the teams. At the same time, we have observed that, while there is frequently a place for some *ad-hoc* processes and for modifications of established frameworks, participants who report their workflow as more structured are also able to share their knowledge and describe their processes more fluently. This opens a question whether there is a need for establishing a more structured framework or process for communication between RA and SO teams.

5. Limitations

Our study used qualitative research methods, and thus our findings have some limitations [5]. We now discuss these limitations and outline the steps we took to minimize them.

As was alluded to in Sections 3 and 4.1, coding is ultimately a subjective means of analysis [28], and this subjectively may introduce doubt as to our conclusions and threaten confirmability of the findings [11]. We took measures to minimize the subjectivity of our results, in line with commonly accepted steps when working with qualitative interview data [27]. For our open coding, two coders independently coded the data and then discussed afterwards to align their codes. Likewise for template analysis, two coders independently applied the template and afterward discussed their placement of quotes and application of the template codes in order to establish inter-coder agreement.

One issue with template analysis is that the design of the template itself can elicit meaning from the qualitative data that is not necessarily there [24]. To address this, we followed the best practices with regards to the application of template analysis, allowing the initial template to expand to encapsulate new information [8]. Additionally, the template was debated amongst the research team to ensure that the template only allowed the reflection of the information in the research data.

Another limitation is the modest sample size, with 11 participants. While researchers have identified that 9 interviews can be sufficient for code saturation [14], [15], what is important for template analysis, a larger sample size could have uncovered more information. Our participants

all work at the same security organization in the same country. Thus, our results can be considered as a case study, and generalization of the results (transferability) can potentially be limited. Our findings might be impacted by the organization's line of business (security services) and its policies and culture. Hence, we encourage independent reproducibility studies at other organizations and in other countries.

Our findings might also suffer from a self-selection bias, as respondents volunteered to participate. Finally, many of the interviews were conducted in a language other than English. There is no consensus on what is best practice when working with qualitative data in multiple languages. We attempted to follow guidelines by first translating this data into English, and checking this translation against multiple sources, and then coding the interviews entirely in English [27].

6. Related work

There have been numerous academic works focusing on studying security practitioners, including security operations and risk assessment. Much of this research has been in the form of interviews, e.g., [2], [12], [17], [26].

In the context of security operations, Kokulu *et al.* have studied technology- and knowledge-related issues at SOC [17]. They interviewed 18 SOC analysts and managers, finding issues with communication within SOC and suggesting means of improving communication outcomes. This study focused on issues internal to the SOC operations, while our focus is on the interplay between security operations and risk managers. Basyurt *et al.* conducted an interview study focused challenges SOC analysts face when processing and sharing threat-related data [7]. They have identified several important challenges related to using technology in this context: e.g., legal restrictions on sharing and receiving data, lack of APIs to acquire threat data, and automation of data processing. In our study, practitioners also mentioned challenges and tasks related to threat intelligence, but our focus is on improving organizational efficiency: to what extent is there a shared cyber risk picture between risk managers and security operations.

Sundaramurthy *et al.* have performed an anthropological study to identify burnout symptoms in SOC practitioners and propose mitigation strategies [21]. Sundaramurthy *et al.* also applied template analysis in an anthropological study of daily workflows in three different SOC [21]. They have reported that processes in different types of SOC diverge due to, among other reasons, their size and budget, and thus experience of SOC analysts with one set of tools is not easily transferable to another SOC context. Axon *et al.* studied how sonification (presenting data as sound) can help SOC operators, and they interviewed security practitioners within SOC and assessed that interview data with template analysis [4]. Further, Alahmadi *et al.* studied perceptions of false positive alerts in SOC using interviews and a survey of interviewed SOC analysts, with the interview and survey data being analyzed with template analysis [3]. We refer the readers to a systematic literature study by Vielberth *et al.* for more information on technical and non-technical challenges in SOC [23].

Practitioners performing security risk management and making security-related decisions have also been studied in the literature. For example, Parkin *et al.* have studied security decision makers, and they recommend to use structured cyber risk scenario representation to help executives to perceive cyber risks better [18]. Schmitz *et al.* researched how security managers make maturity level assessments for information security controls [19]. They have found that managers could use better structural support for making such decisions (in the form of detailed definitions and guidelines). Similarly, we have observed that having a structure enables practitioners to understand their own workflows better.

Relevant to our work, Ahmad *et al.* conducted interviews with incident response managers from four incident response teams and found that risk assessment information is generally not disseminated, indicating a communication gap [2]. More than 10 years have passed since this study, and our results show that there is still this communication gap that hinders usage of valuable threat information across the risk assessment, security operations and incident response teams. In a more recent study, Ahmad *et al.* also report that very few organizations enable integration and mutual learning between information security management (refers to RA in our study) and incident response (which includes SO) [1]. The benefits of such integration that they enumerate (including better defensive posture and more up-to-date threat intelligence) are closely related to what we learned from our participants.

Shreeve *et al.* have studied how different types of stakeholders (executives, legal and financial experts, managers, risk analysts and security specialists, among others) take security-related decisions in a table-top game [20]. They have found that background of participants does not have a strong effect on the game outcomes and that participating teams rarely apply systematic approaches to risk decision making, such as threat modeling, when faced with complex decisions, underscoring the need for organizations to think more systematically about this [20].

7. Conclusions

In our work we have studied qualitatively the interactions between risk analysts and a SOC team in a security service company. We have identified what these practitioners do and how they interact with each other, and we have found important gaps in their communication. The use-case implementation has been identified as an important means of communication, and we suggest that it could be further researched and developed as a means to exchange timely cyber threat information. Finally, our participants have named other issues, such as client expectation management, which are shared between both risk analysis and security operations. Those can be further explored to improve overall effectiveness.

References

- [1] A. Ahmad, K. Desouza, S. Maynard, H. Naseer, and R. Baskerville. How integration of cyber security management and incident response enables organizational learning. *J. of the Association for Inf. Sci. and Tech.*, 71(8):939–953, 2020.

- [2] A. Ahmad, J. Hadgkiss, and A. B. Ruighaver. Incident response teams - challenges in supporting the organisational security function. *Computers & Security*, 31(5):643–652, Jul 2012.
- [3] B. Alahmadi, L. Axon, and I. Martinovic. 99% false positives: A qualitative study of SOC analysts' perspectives on security alarms. In *USENIX Security*, 2022.
- [4] L. Axon, B. Alahmadi, J. Nurse, M. Goldsmith, and S. Creese. Data presentation in security operations centres: exploring the potential for sonification to enhance existing practice. *J. of Cybersecurity*, 6, Jan 2020.
- [5] R. Barbour. Quality of data analysis. *The SAGE handbook of qualitative data analysis*, pages 496–510, 2014.
- [6] S. Barnum. Standardizing cyber threat intelligence information with the structured threat information expression (STIX). *Mitre Corporation*, 11:1–22, 2012.
- [7] A. S. Basyurt, J. Fromm, P. Kuehn, M. Kaufhold, and M. Mirbabaie. Help wanted-challenges in data collection, analysis and communication of cyber threats in security operation centers. In *Wirtschaftsinformatik 2022 Proceedings*, 2022.
- [8] J. Brooks and N. King. *Doing Template Analysis: Evaluating an End-of-Life Care Service*. 2014.
- [9] M. Brunner, C. Sauerwein, M. Felderer, and R. Breu. Risk management practices in information security: Exploring the status quo in the dach region. *Computers & Security*, 92:101776, 2020.
- [10] A. Calder and S. G. Watkins. *Information security risk management for ISO27001/ISO27002*. It Governance Ltd, 2010.
- [11] D. Cruzes and L. ben Othmane. Threats to validity in empirical software security research. *Emp. Research for Soft. Security: Foundations and Experience*, 2017.
- [12] J. Goodall, W. Lutters, and A. Komlodi. I know my network: collaboration and expertise in intrusion detection. In *Proc. of CSCW*, 2004.
- [13] A. Hasenstab. An exploration of reflective practice in security risk management: how senior security managers experience reflective practice. *Reflective Practice*, 19(4):522–529, Jul 2018.
- [14] Monique Hennink and Bonnie N Kaiser. Sample sizes for saturation in qualitative research: A systematic review of empirical tests. *Social science & medicine*, 292:114523, 2022.
- [15] Monique M Hennink, Bonnie N Kaiser, and Vincent C Marconi. Code saturation versus meaning saturation: how many interviews are enough? *Qualitative health research*, 27(4):591–608, 2017.
- [16] N. King. *Template analysis*, pages 118–134. Sage Publications Ltd, Thousand Oaks, CA, 1998.
- [17] F. Kokulu, Y. Shoshitaishvili, A. Soneji, Z. Zhao, G. Ahn, T. Bao, and A. Doupé. Matched and mismatched SOCs: A qualitative study on security operations center issues. In *Proc. of the 2019 ACM CCS*, 2019.
- [18] S. Parkin, K. Kuhn, and S. Shaikh. Scenario-driven assessment of cyber risk perception at the security executive level. In *Usable Sec. and Privacy W.*, 2021.
- [19] C. Schmitz, M. Schmid, and S. Harborth, D. and Pape. Maturity level assessments of information security controls: An empirical analysis of practitioners assessment capabilities. *Computers & Security*, 108:102306, 2021.
- [20] B. Shreeve, J. Hallett, M. Edwards, K. Ramokapane, R. Atkins, and A. Rashid. The best laid plans or lack thereof: Security decision-making of different stakeholder groups. *IEEE Trans. on Software Engineering*, 48(5):1515–1528, 2020.
- [21] S. C. Sundaramurthy, J. Case, T. Truong, L. Zomlot, and M. Hoffmann. A tale of three security operation centers. In *Proc. of W. on Sec. Inf. Workers*, 2014.
- [22] R. van Os, F. Ladan, T. van Casteren, R. Toornstra, R. Metsemakers, L. Nieuwenhuize, and H. Grotenhuis. MaGMA use case framework, Apr 2018.
- [23] M. Vielberth, F. Böhm, I. Fichtinger, and G. Pernul. Security operations center: A systematic study and open challenges. *IEEE Access*, 8:227756–227779, 2020.
- [24] T. Waring and D. Wainwright. Issues and challenges in the use of template analysis: Two comparative case studies from the field. *Electr. J. of Bus. Research Methods*, 6(11), 2008.
- [25] J. Webb, A. Ahmad, S. Maynard, and G. Shanks. A situation awareness model for information security risk management. *Computers & security*, 44:1–15, 2014.
- [26] R. Werlinger, K. Hawkey, D. Botta, and K. Beznosov. Security practitioners in context: Their activities and interactions with other stakeholders within organizations. *Int. J. of Human-Computer Studies*, 67(7):584–606, Jul 2009.
- [27] D. Wicks. The coding manual for qualitative researchers. *Qualitative research in organizations and management: An Int. J.*, 12(2):169–170, 2017.
- [28] M. Williams and T. Moser. The art of coding and thematic exploration in qualitative research. *Int. Management Review*, 15(1):45–55, 2019.

General opening questions:

- Q1 How many years of experience in cyber security do you have?
- Q2 What is your current role? And what does it involve in terms of activities and responsibilities?
- Q3 Does your company provide services in terms of risk analysis and monitoring & detection to clients?
- Q3.1 If yes, how does this process looks like – from risk analysis to the implementation to monitoring?

TABLE 4. THE TEMPLATE AND EXPLANATION OF CODES (QUANTITY REFERS TO # TIMES THE (SUB)CODE HAS BEEN APPLIED)

Theme	Template Codes	Subcodes	Quantity	Explanation
Work Tasks			303	
	Risk Assessment		133	Risk assessment is a code for all tasks that would be traditionally performed by RA personnel. The <i>Risk management</i> sub-code is used to refer to the tasks traditionally involved with analyzing and managing risks (e.g., risk quantification or reduction steps). The <i>Risk Assessment: Other</i> sub-code is for generic risk assessment tasks that are hard to classify, for instance, an interviewee discussing "risk-based scoping". It is also used when someone describes the RA work in very generic terms.
		Risk Management	104	
		Other	35	
	Security Operations		188	
		Alert Evaluation	30	Security operations is a code for tasks that would be traditionally performed by security operations personnel. <i>Monitoring</i> refers to the general practice of checking for, being prepared to respond to, detecting and responding to an attack. <i>Alert evaluation</i> refers to the process of an analyst checking if an alert is valid, or just a false positive. <i>Use-case implementation</i> concerns the translation of risks into actionable, technical use-cases which are used to set alerts. <i>Security Operations: Other</i> is used similarly to the risk assessment version, in that it is for descriptions of security operations tasks that are non-specific or hard to independently classify.
		Monitoring	55	
		Use-case Impl.	67	
		Other	52	
	Threat Intelligence		16	Collecting information and ingesting data about cyber threats.
Process			321	
	Scope		140	These codes refer to whether or not there are clear limits on the work being done. <i>Limited</i> would refer to a clear delimitation between different tasks while <i>Expansive</i> refers to tasks which never clearly end or are under the purview of another operator.
		Expansive	40	
		Incidental	15	
		Limited	38	
		Unclear	56	These codes refer to the steps followed to complete the work. <i>Framework</i> refers to the task having clear steps that an operator can follow, <i>Ad-hoc</i> refers to a process that does not have clearly defined steps, <i>Unique</i> is a mixture of both, where a process has some clearly defined steps and portions that are completely undefined.
	Structure		165	
		Ad-hoc	34	
		Framework	56	
		Mixed	30	These codes refer to when a process is followed. <i>Client Intake</i> is an initial process, <i>Event-driven</i> is a process that occurs in response to some trigger, and <i>Scheduled</i> refers to a repeated or regular task.
		Unclear	81	
	When		61	
		Client Intake	24	
		Event-driven	20	These codes refer to whom the interviewee refers to as performing the task: <i>Own</i> is themselves or their team and <i>Other</i> is another team or external party.
		Scheduled	8	
		Unclear	13	
	Who		260	
		Own	188	An interviewee provides an example of their work/process.
		Other	54	
		Unclear	18	
	Example		20	
Communication			189	
	Content		91	The content of the communication being sent. <i>Complete</i> means that all available information is included, and <i>Abridged</i> means that some information is excluded from the communication.
		Complete	20	
		Abridged	13	
		Missing	18	
		Unclear	46	The direction of the information being transmitted from the perspective of the interviewee.
	Direction		54	
		Sent	34	
		Received	26	
	Participant		177	Who is participating in the communication. <i>Client</i> is the customer, <i>Internal</i> would be staff from the same organization as the interviewee, and <i>Third party</i> would be anyone other than internal staff and the client.
		Client	69	
		Internal	96	
		Third party	17	
		Unclear	10	How the communication is organized. <i>Framework</i> means that there is a defined structure to the communication; <i>Case-by-case</i> means that the communication occasionally has a defined structure, but not always.
	Structure		85	
		Ad-hoc	6	
		Framework	15	
		Missing	9	When the communication occurs: <i>Available</i> is for communications that are not active, but made accessible. <i>Event-driven</i> is for active communications that are sent in response to some trigger, while <i>Scheduled</i> refers to communication that occurs at regular intervals.
		Unclear	57	
	When		69	
		Available	14	
		Event-driven	25	Identification of current issues within the organization. Subcodes are provided to specifically identify what is currently problematic from the perspective of the interviewer.
		Scheduled	13	
		Unclear	21	
Perspective			131	
	Improvement		67	<i>Process</i> and <i>Communication</i> are for improvements with either work processes or with communication, while <i>Not needed</i> is for the cases where interviewees note that no improvements are needed.
		Process	32	
		Communic.	35	
		Not needed	2	
	Issues		48	A descriptor of the state of process, communications, etc. within the organization. <i>Past</i> refers to how the organization used to operate, while <i>Ideal</i> refers to what the interviewee thinks the state should be.
		Client	10	
		Communic.	11	
		Process	14	
		Technology	12	
		Other	9	
	State		42	
		Past	38	
		Ideal	4	

Risk assessment expert questions:

- Q4 Which methods / techniques / frameworks do you use to perform risk assessment/ threat modeling to identify threats at a client?
- Q5 What are the results / kind of information that you get after performing risk assessment / threat modeling, in terms of identified risks and the format of this information?
- Q5.1 What actions are taken after these exercises? To which persons do you communicate these findings? And how?
- Q6 How does the process of longer-term monitoring of the identified threats and risks look like? To see if the measures are taken and if the threats are in control.
- Q6.1 Do you use Key Risk Indicators to monitor this? If not, how do you monitor this?
- Q6.2 And who is responsible?
- Q7 Are there any activities that ensure communication of the risks to the SOC team in order to monitor these?
- Q7.1 If yes, how is the information exchange organized with the SOC?
- Q7.2 In what form and in what kind is the information being shared?
- Q7.3 Which person receives the information?
- Q7.4 If no, what is the reason and how could it be done?
- Q8 Can you describe the relationship of risk assessment with security operations / SOC?
- Q9 Do you use inputs from SOC for the risk assessment and threat modeling?
- Q9.1 In terms of data about likelihood or specific events for example?
- Q9.2 Anything else?
- Q10 Do you implement the countermeasures like use-cases / rules for the SIEM yourself or does another colleague do this?
- Q10.1 If yes, what kind of information do you need to do this?
- Q10.2 If no, what kind of information do you think is needed?

Security operations experts questions:

- Q4 We are now focusing on your role in the SOC/threat monitoring at a client. What kind of threat intelligence and threat models does the SOC receive?
- Q4.1 Which of these intelligence sources are internal? And which external?
- Q5 How do you receive information about risks & threats / controls / use-cases from risk assessment and threat modeling that have to be implemented for a client? How is information being exchanged?
- Q6 How do you integrate the information from the risk management and threat modeling team into your processes?

- Q6.1 E.g., are you tasked to implement recommended controls?
- Q6.2 Do you have a way to consume the identified threat scenarios and monitor for those?
- Q6.3 Do you receive any Key Risk Indicators for monitoring?
- Q7 What kind of information do you need from risk management and threat modeling to improve the security operations and to be able to better respond to actual attacks?
- Q7.1 In what kind of format?
- Q7.2 Do risk assessment and threat modeling provide this information to implement reactive and corrective controls when an attack occurs? Is there anything missing?
- Q8 Do you provide any data/information to the risk management and threat modeling team or participate in their activities?
- Q8.1 Do you supply likelihood estimates, consult them on observed threats or recent threats from the threat intelligence feeds?

General closing questions:

- Q11 Are you encountering problems / gaps in the process, communication between risk assessment and security operations?
- Q11.1 If yes, can you explain?
- Q11.2 What could be possible improvements?
- Q12 Are you available for follow-up questions?