Mindful balancing: Avoiding Alert Fatigue in Security Operation Centers

TERJE HEUM SELJÅSEN

ADRIAN MIKKELSEN

SUPERVISOR

Wael Soliman

University of Agder, 2024

Faculty of Social Science

Obligatorisk gruppeerklæring

Den enkelte student er selv ansvarlig for å sette seg inn i hva som er lovlige hjelpemidler, retningslinjer for bruk av disse og regler om kildebruk. Erklæringen skal bevisstgjøre studentene på deres ansvar og hvilke konsekvenser fusk kan medføre. Manglende erklæring fritar ikke studentene fra sitt ansvar.

1.	Vi erklærer herved at vår besvarelse er vårt eget arbeid, og at vi ikke har brukt andre kilder eller har mottatt annen hjelp enn det som er nevnt i besvarelsen.	**Ja**
2.	**Vi erklærer videre at denne besvarelsen:** • Ikke har vært brukt til annen eksamen ved annen avdeling/universitet/høgskole innenlands eller utenlands. • Ikke refererer til andres arbeid uten at det er oppgitt. • Ikke refererer til eget tidligere arbeid uten at det er oppgitt. • Har alle referansene oppgitt i litteraturlisten. • Ikke er en kopi, duplikat eller avskrift av andres arbeid eller besvarelse.	**Ja**
3.	Vi er kjent med at brudd på ovennevnte er å betrakte som fusk og kan medføre annullering av eksamen og utestengelse fra universiteter og høgskoler i Norge, jf. Universitets- og høgskoleloven §§4-7 og 4-8 og Forskrift om eksamen §§ 31.	**Ja**
4.	Vi er kjent med at alle innleverte oppgaver kan bli plagiatkontrollert.	**Ja**
5.	Vi er kjent med at Universitetet i Agder vil behandle alle saker hvor det forligger mistanke om fusk etter høgskolens retningslinjer for behandling av saker om fusk.	**Ja**
6.	Vi har satt oss inn i regler og retningslinjer i bruk av kilder og referanser på biblioteket sine nettsider.	**Ja**
7.	Vi har i flertall blitt enige om at innsatsen innad i gruppen er merkbart forskjellig og ønsker dermed å vurderes individuelt. Ordinært vurderes alle deltakere i prosjektet samlet.	**Nei**

Publiseringsavtale

Fullmakt til elektronisk publisering av oppgaven Forfatter(ne) har opphavsrett til oppgaven. Det betyr blant annet enerett til å gjøre verket tilgjengelig for allmennheten (Åndsverkloven. §2).
Oppgaver som er unntatt offentlighet eller taushetsbelagt/konfidensiell vil ikke bli publisert.

Vi gir herved Universitetet i Agder en vederlagsfri rett til å gjøre oppgaven tilgjengelig for elektronisk publisering:	**Ja**
Er oppgaven båndlagt (konfidensiell)?	**Nei**
Er oppgaven unntatt offentlighet?	**Nei**

Acknowledgments

This thesis was only possible with the guidance of our supervisor, Wael Soliman. Therefore, we would like to thank him for his support throughout this journey, for providing feedback, and for holding bi-weekly meetings.

We would also like to give a big thanks to the interviewees. They took time out of their schedule to provide us with empirical data that could be used in the thesis, and for that, we are very grateful.

Lastly, we give big thanks to our friends and family, especially Kristian and Daniel, who gave us unyielding support during the master's thesis period. This has motivated us to work continuously and improve the thesis throughout its progress.

We are forever grateful.

Kristiansand 2024

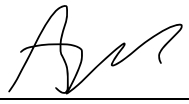

Adrian Mikkelsen

Terje Heum Seljåsen

Abstract

Analysts in Security Operation Centers (SOC) are experiencing large numbers of alerts that they must analyze. Out of all the alarms an analyst receives, between 50 and 90 % of them are false positives. Because of these numbers, they are prone to be affected by alert fatigue (AF), a condition where an individual is desensitized to alerts and suffers from cognitive overload. This issue could lead to a successful cyber-attack and cause damage to an organization. This thesis investigates how alert fatigue influences analysts and how SOCs should protect themselves from this issue. This is done by answering the following research questions: "*How does mindful organizing happen in the SOC environment?*", "*How does alert fatigue influence mindful conduction in SOC?*" and "*What strategies are being implemented to mitigate alert fatigue"*. We conducted a multiple case study interviewing 12 individuals from 5 different SOCs to answer these questions. We also utilized a High-reliability organization (HRO) framework as a theoretical lens since a SOC can be considered a digital HRO, which is an organization that builds both services and solutions to protect an organization and utilizes a cognitive mindset of mindfulness in its operations.

The main theoretical implication is our creation of a novel theory that we call Mindful Balancing, which entails how analysts use both mindful and mindless actions when they analyze alerts and that they are doing so to maintain their cognitive power. Failure to maintain that balance can cause them to drain their mindfulness, which is an exhaustible resource, and thus return to a default state of mindlessness. For the practical implications, we identified various mitigation strategies that could be used to prevent alert fatigue. They are split into categories based on how they can improve Mindful Balancing, which is assisting mindful balancing and facilitation of mindful balancing.

This thesis offers an interesting new way to see alert fatigue and provides a novel theory about how analysts use mindful and mindless actions to maintain their cognitive power. We hope this theory will assist organizations in understanding how analysts work and that it will assist further research on the topic.

Contents

List of Tables

List of figures

1 Introduction

The threat landscape in the realm of cyber security is ever-changing, especially in 2024 with war in Europe. Cybersecurity professionals are trying to keep track of emerging attack methods, Advanced persistent threats (APT), and are developing new countermeasures as new vulnerabilities appear (NSM, 2023). One can argue that keeping track of network activity and monitoring computers has become more important, and therefore more systems are being put in place to help security professionals. Security Operation Centers (SOC) are one of the units which do this. A SOC is a centralized unit or group in charge of enhancing an organization's cybersecurity and preventing, detecting, and responding to threats. Additionally, it performs proactive security by retrieving threat intelligence from various sources to maintain vigilance towards new emerging threats Microsoft (n.d.). They utilize intrusion detection systems (IDS) and endpoint protection, which are implemented to cover the attack surfaces, each providing data that are turned into alerts that could indicate an attack. As the attack surfaces increase with the digitalization we see today, more alerts appear, and with that, there is an increased number of false-positive alerts (Ban et al., 2023).

This increase in false positive alerts can introduce *alert fatigue*, a concept used in various first responder domains such as firefighting, health care, etc., that describes the overwhelming feeling caused by a vast amount of false positive alarms. Alert fatigue can compromise the first responder's ability to respond accurately to critical alarms because they have been desensitized by the sheer number of alarms (PSNet, 2019). The same concept is used in the domain of cybersecurity, where security professionals work with solutions for monitoring IT systems, such as SOC. If there are too many alarms to keep track of, employees might ignore the alarm and reduce the quality of the security. A study shows that one-third of IT professionals tend to ignore alarms due to the sheer number of false positives (Ban et al., 2023).

1.1 Research Aim

We have found a gap in the literature, where there is little research on how alert fatigue unfolds in SOC, even though there are multiple blog posts about the problem saying how dire it is. Therefore, this thesis aims to understand alert fatigue in SOC and how it affects SOC operations. To investigate it, we have defined SOC as a High Reliability Organization (HRO) and used that as a lens to find how mindful organizing happens in SOC. We have also investigated how alert fatigue affects mindfulness and what SOCs do to mitigate it. Mindfulness is seen as an exhaustible resource needed to avoid mindlessness (Ault & Brandley, 2023); therefore, it is essential to investigate how it can be preserved.

Our approach involves a systematic literature review to provide an overview of the current body of knowledge and a multiple case study consisting of 12 interviews from 5 SOC environments.

1.2 Research Approach

We found that a qualitative approach would be the best fit for our research, as it will allow us to have in-depth interviews to investigate this socio-technical phenomenon. Before the interviews, we used a systematic literature review to explore the current body of knowledge. We investigated the topic of alert fatigue in both the health sector and IT since the topic is a broad term not specific to IT. The literature review resulted in 79 articles**,** where 24 articles were included in the final thesis. We also acquired a lens that would help shed light on our topic and angle our research. We proceeded to acquire 5 cases, which would result in 12 interviews, with 2 and 3 interviews from each case ranging from 30 minutes to 1 hour in length. The interview questions included questions about their role, seniority, and their experiences on their SOC. We performed the interviews in March and April as semi-structured interviews, which means we had some predefined questions and some made during the interviews to gain better insight into the particular interviewee. We tried to avoid informing them about the topic of alert fatigue to avoid them projecting their thoughts about it on their answer. Our efforts were only sometimes successful as most of them had been informed by their leader about the topic. The topic was also always disclosed at the end of the interviews to get their final thoughts about alert fatigue. By performing the systematic literature review and the semi-structured interviews, we were able to get results that enabled us to answer our research questions.

1.3 Structure of the thesis

1 – Introduction: Gives an overview of the current threat landscape and the situation of technology and alert fatigue, including the aim of the research and approach.

2 – Background and related work: This section gives an overview of the systematic literature review that has been performed and insight into the theoretical articles that support the thesis.

3 – Theoretical framework: Explains the theoretical framework that is used and why it is being used.

4 – Research approach: Discuss the selected research approach and why it was selected. It also discusses the data collection and analysis.

5 – Findings: Presents the empirical data that have been collected and analyzed in the previous chapter.

6 – Discussion: We address our findings in relation to the literature and put forward the theoretical implications, practical implications, limitations, and future research.

7 – Conclusion: Provides this thesis conclusion and the closing remarks of our thesis.

2 Background and Related Work

The literature review, which will be covered in this chapter, offers an overview of the body of knowledge that will aid in addressing the research questions. We will put forward and explain the literature review methodology, the steps taken to conduct the literature search, a summary of keywords, the screening procedure, and a quality assessment. Following that, a summary of the included articles provides insight into the topic, and at the end, the chapter concludes with a more in-depth discussion of the findings.

2.1 Literature review

When conducting research, it is important to collect earlier work and investigate the current body of knowledge to advance the knowledge and identify gaps to explore. One can test hypotheses or create new theories by summarizing, analyzing, and synthesizing a group of collected literature that is related (Xiao & Watson, 2019). We will therefore in this chapter explain the method used in finding, selecting, and collecting the literature.

2.1.1 Literature Method

In our research, we chose systematic literature review (SLR) as our methodology. This methodology is a rigorous approach to retrieving literature that enables others to recreate what we produce (Okoli, 2015).

"A rigorous standalone literature review must be systematic in following a methodological approach, explicit in explaining the procedures by which it was conducted, comprehensive in its scope of including all relevant material, and, hence, reproducible by others who would follow the same approach in reviewing the topic." (Okoli, 2015, p. 880)

When conducting the SLR, we will be following the Xiao & Watson model. This model uses eight steps, which are as follows: formulate the problem, develop, and validate the review protocol, search the literature, screen for inclusion, assess quality, extract data, analyze and synthesize data, and report findings (Xiao & Watson, 2019). By following this model and its step-by-step guidelines, we can accurately perform an SLR which will collect data that is valid and extensive to answer our research questions.

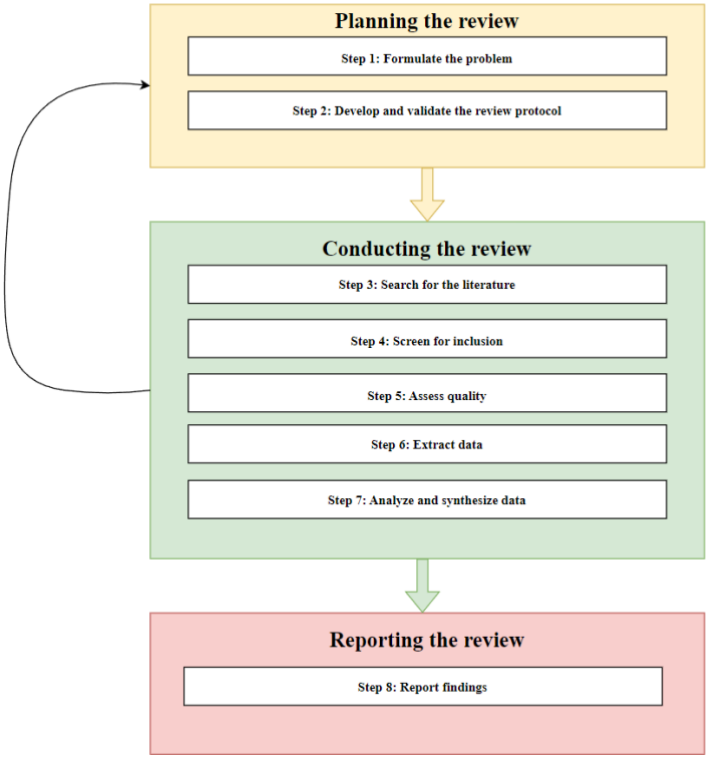

Figure 1: Literature method (Xiao & Watson, 2019)

2.1.2 Literature search

Broadly speaking, the three major sources to find literature are electronic databases, backward searching, and forward searching (Xiao & Watson, 2019). It is also important to remember that "*No database includes the complete set of published literature collections*" (Xiao & Watson, 2019, p. 103), which is why we have collected data from various sources. Our demand for electronic databases was that they needed to be able to apply sophisticated filtering, reducing the number of publications and retrieving those more consistent with the research. We ended up with the following databases:

Google Scholar provides a simple way to broadly search for literature. It retrieves articles from a wide range of sources and disciplines (*About Google Scholar*, n.d.).

Web of Science is a platform formed by several different literature search databases (*Introduction - Web of Science Platform - LibGuides at Clarivate Analytics*, n.d.).

IEEE Xplore is a digital platform for finding and accessing scientific and technical literature published by the IEEE or its publishing partners (*About IEEE Xplore*, n.d.).

To find relevant information on the topic we had to create keywords that would aid in the search. Some of the first keywords were "Alert fatigue in Cybersecurity", "Alert fatigue", "Alert fatigue AND Blue team" and "Alarm fatigue". However, to establish more information on what it means to be alert fatigued we had to investigate fatigue and different types of fatigue. The table below provides an overview of the keywords used

and the number of articles the databases have, including settings that have been used to refine the search to get the article number down.

		Database		
Keywords	Refining the search	Web of science	Google Scholar	IEEE
“Alert fatigue”		409	8640	121
Alert fatigue cybersecurity		5	7230	13
"Alert fatigue" AND "Blue Team"		0	65	0
"Alarm fatigue" OR "Alert fatigue"		795	14500	126
“False positive” AND “Security analyst”		7	1720	2
False positive AND Security operation center		85	107000	169
SOC AND "alert fatigue" OR "alarm fatigue"	2023, 2024	58	1220	18
Fatigue and "Shift work"	Web of science category = Nursing	161	55 700	11
Mental Fatigue		67	304000	879
Fatigue AND Workload	2023, 2024	316	15500	56

Table 1: Searches with keywords

During the search for articles, more keywords that are relevant to the research question were found and used to discover more articles. We also used features in the search engine to narrow the scope of potential articles, such as choosing only articles from a certain category like “Computer science information systems” or changing the year. Additionally, a quick look at the article's title gave us enough information if the article was somewhat relevant. Regarding Google Scholar, the number of articles in the results was too broad. We, therefore, sorted by relevancy and looked at the first three articles to mainly look at the articles that were most relevant for our search. Backward and forward searches were also performed when articles highly relevant to the research were found.

Due to our limited knowledge in the area, we used the search to build up knowledge on the fields that were relevant to alert fatigue and fatigue. As our search commenced, we read articles that we found interesting and added them to our database. For each

search the database grew, and as we reached a large enough quantity, we saw ourselves content with the result and could continue to the next step in the process.

2.1.3 Literature screening and Quality assessment

After completing the list of references, we conducted a screening to establish if each article should be included for analysis and data extraction. In the screening process, we used the two-stage procedure that Xiao & Watson mentioned. The articles that were first found were picked based on the relevance of their titles, and afterward, the different articles were assessed by doing an in-depth look into their abstracts and conclusions to evaluate if they could be used in the study. If the articles were found interesting, we performed a full-text quality assessment. If the article was irrelevant to our research question or other criteria, then it was removed.

During the screening process, we set different inclusion and exclusion criteria to ensure that the literature we are searching for is relevant to our research questions. The table below shows the selected inclusion and exclusion criteria.

Inclusion criteria	Exclusion criteria
The article includes information about fatigue or alert fatigue.	Not available in English or Norwegian
The articles come from a reputational source	Not scientific articles
The article is relevant to our research	Not published in journals that are ranked in the Kanalregister

Table 2: Inclusion and Exclusion criteria

79 articles were found based on their title. We methodically studied the abstract and conclusion of every publication we found that was relevant to our research to have a better grasp of the work and determine whether it could be useful. 26 of the 79 articles were eliminated after being reviewed. After that, we read each article thoroughly to evaluate the article's overall quality. 29 of the articles were eliminated after their quality was evaluated; the result was the 24 remaining.

The information flow and process used in the literature screening are depicted in the PRISMA flow diagram below.

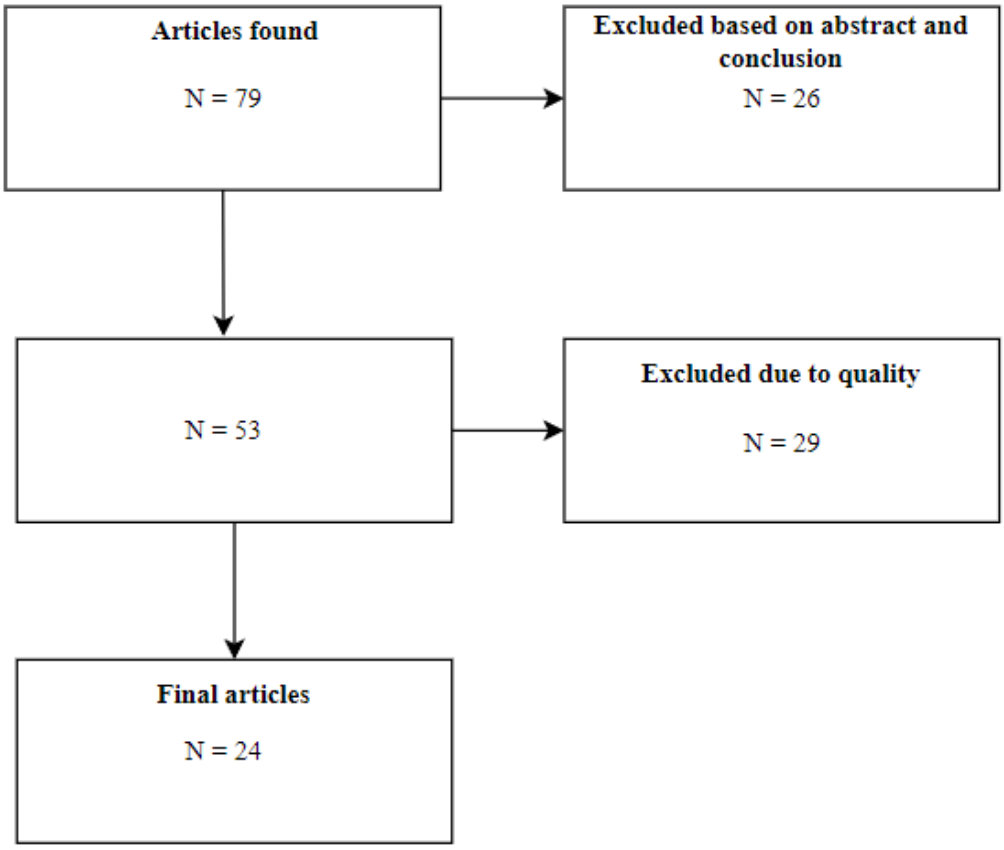

Figure 2: PRISMA model

2.1.4 Data extraction and Analysis

The information that was found relevant to our research was extracted from the articles. This includes themes, frameworks, topics, or statements that would provide value to the literature review.

2.2 Literature findings

This section aims to provide a comprehensive overview of the found literature on the topic of alert fatigue, fatigue, and mitigations for alert fatigue. The first part of the findings will provide an extensive view of alert fatigue and give a definition that we have created based on the literature. Further, we will delve into general fatigue and its causes. Lastly, we investigate strategies that mitigate or prevent alert fatigue. The mitigations are split into two categories, which are technical and human factor mitigations. This is to provide an overview of the current focus areas in the literature. A theoretical framework, high-reliability organizations (HRO), has also been discovered. This framework is explained in Chapter 3 and will be used as this thesis's theoretical lens. An overview of the 24 articles used in the findings can be found in Appendix A.

2.2.1 Alert fatigue

Before providing and discussing the literature about alert fatigue, we will propose the definition for alert fatigue which we have created based on the upcoming literature, to provide an understanding of alert fatigue that allows the literature to be easier to digest. As a result of the literature, we have defined alert fatigue as:

> *"Alert fatigue is a state where individuals are exposed to excessive numbers of false positive alerts, and become desensitized to the alerts, meaning that they are less likely to analyze them adequately, and suffer from cognitive overload due to its numbers, meaning the individual is mentally exhausted"*

The phenomenon of Alert Fatigue is a term used to describe the state of an individual who can no longer respond appropriately to alerts. The leading cause for this is having to respond to large numbers of non-actionable alerts (Alahmadi et al., 2022), in addition to the vast amounts of information mixed with uninformative details and a lack of time or cognitive resources to separate relevant information from the irrelevant (Ancker et al., 2017). When this occurs, an individual will end up being both desensitized and cognitively overloaded. When an individual is cognitively overloaded and desensitized, they are alert fatigued, which leads to a lack of, or reduced action from the respondent (Alahmadi et al., 2022).

According to a survey in a German university, 56% of all alarms in a hospital are non-actionable or false positives (Wunderlich et al., 2023). The fact that there are too many false alerts is agreed upon by many other studies, and the consequences are affecting their ability to care for their patients. Seok et al. (2023) state that "*Excessive exposure to false-positive alarms that occur without accurate physiological data violations causes sensory overload in nurses, lowering their sensitivity to alarms*" (Seok et al., 2023, p. 946). This is similarly stated by Hravnak et al. (2018), as the number of non-actionable alerts increases, *"… there is an incremental increase of clinical response delay*" (Hravnak et al., 2018, p. 545). Casey et al. (2018) found the same results, where there was a high level of agreement that there are too many alarms that are disrupting them from caring for their patients and leading them to disable the alarms.

Aside from these factors, some studies point to heavy workloads and long shifts as causes of alert fatigue. Ding et al. (2023) state that night shifts have a negative effect on performance in the hospital. Movahedi et al. (2023), agree and point out that "*Heavy workload, long hours of work shifts, and high noise levels in the workplace are associated with desensitization due to alarm fatigue*" (Movahedi et al., 2023,p. 981). Interestingly, this statement indicates that more variables can increase the probability of becoming alert fatigued, showing that factors other than the number of alerts could have an impact. Additionally, he further emphasizes a connection between alert fatigue and desensitization.

Comparing the health sector to the IT security sector, more specifically the SOCs', we find many similarities. Studies like Ndichu et al. (2021) and Ban et al. (2023) state that the number of security alerts produced by their tools is causing analysts to become desensitized to the alerts. A study by McRee, (2022) found that the primary challenges are noise and minor problems, chasing false positives, overwhelmed team members,

and triaging alerts, underlining the problems with the vast amounts of false positives. One of the participants in the interview from the study by Kokulu et al. (2019) points out that the information their systems gather increased, from 17 to 20 terabytes in just two months. They say that they are at a point of information overload and are having to filter through 50% of the first-level threats they collect. Additionally, the same study found that at least 5 participants found the information in some of their feeds to be filled with "*uncorrelated, industry-unrelated, and low-quality data*"(Kokulu et al., 2019, p. 1960)

While there are many studies and articles that very much point toward false positives as the root cause of alert fatigue, the study by Kokulu et al. (2019) that investigated issues in SOC, found that false positives were not a major concern for the analysts. Fifteen interviewees responded with the same answer, which contradicts most other studies about alerts. This study was, however, not centered around alert fatigue, and alert fatigue was not mentioned by the interviewer or the interviewees.

When looking at alert fatigue it's also important to look at what the literature says about fatigue in general, thus we explore two definitions that have been identified. Aaronson et al. (1999) defined it as "*The awareness of a decreased capacity for physical and/or mental activity due to an imbalance in the availability, utilization, and/or restoration of resources needed to perform the activity"(Aaronson et al.*1999, p. 46). Their explanation of the resource aspect of their definition is that it can either be biochemical properties and other physiological capacities to social and cultural factors that can affect a person's reaction or fatigue experience. Thus, under this definition of fatigue, where replenishment and utilization take place, the resources are constantly used to support an individual's activities. If any disturbance occurs to either the needed amount of resources or the replenishment, fatigue sets in. A different definition of fatigue is also put forward by the author Lerdal et al. (2005)

> *"Fatigue is a symptom of both physical and mental disease. It also occurs as a side effect of medical therapy and as an imbalance between individuals' energy potential and performance of daily activities and exists to a significant degree in the general population." (Ledal et al., 2005, p123)*

This explanation in many ways resembles the definition created by Aaronson et al. (1999), in which energy is seen as a resource. Additionally, an individual needs the correct amount, or fatigue can occur. While looking into some of the causes of fatigue Yuan et al. (2011) state in their research that the lifestyle factor that had the most significant impact on the contribution of fatigue is poor sleep quality. However, some of the most vital lifestyle factors including the sleep quality for chronic fatigue were higher workload perception, lack of exercise, and non-availability of support. According to the article from Åkerstedt et al. (2004), it is said that fatigue usually occurs as a result of either long periods of stress or disease. Additionally, they reference studies that aim to explain the causal chain of the specific features of stress that are involved. Some studies identify role conflicts and work demands as causes of fatigue, as well as being overweight and lack of exercise, as predictions for male individuals. The authors, Åkersted et al. (2024), found that fatigue could be predicted by "*high work demands, low social support, being a supervisor and being a female, inability to stop thinking about*

work during leisure time, snoring, and disturbed sleep"(Åkerstedt et al., 2004, p. 431). They also found that fatigue was strongly predicted by disturbed sleep and that fatigue was less likely among people who were older and more physically active.

2.2.2 Consequences of alert fatigue

Claudio et al. (2021) and Alahmadi et al. (2022) have found that alert fatigue can cause personnel in the health sector to have a lower response time and might be unwilling to respond to alerts due to desensitization. An additional factor they found was that false positives can lead to distrust, leading to the same outcome. Ndichu et al. (2021) show the same findings for IT personnel who often will ignore alerts or have decreased response time due to desensitization. The effects of alert fatigue can have critical consequences in the health sector, as stated by Seok et al. (2023), it might be *"…rendering them unable to differentiate between false-positive and actual alarms in real life-threatening situations. This may further lead to their inability to recognize alarms as serious and requiring immediate action"* (Seok et al., 2023,p. 947). This is not specific to the health sector as we find similar consequences in the IT sector where Ban et al. (2023) inform that desensitization can lead to insufficient response and missed critical incidents, which causes increased security risks.

Other unwanted consequence of alert fatigue is that cognitive overload can lead to missed alerts due to silenced or altered alarms. Multiple studies, like Claudio et al. (2021)**,** Casey et al. (2018) and Movahedi et al. (2023) found that health personnel are often prone to silence alarms or alter them to avoid false positives and ease their shift. This is also a common practice in IT, where SOC personnel will tweak alerts to avoid false positives. Kokulu et al. (2019) found in their interviews that this practice is common and that the interview stated they must find a balance when altering the alarms. A quote from one of the participants says "*False positives are always a balancing act. You can tune out false positives, but you have to be very careful how you do it. If you want, I can eliminate 100% false positives, but I am going to miss some of the true positives*" (Kokulu et al., 2019, p. 1961).

Missing true alerts can have severe consequences, like critical data breaches or even death. In the study by Hravnak et al. (2018), they found that death has been the outcome in multiple cases, which also has led to self-blame and burnout in health personnel. They say that patients are the second victims of alert fatigue, after the health personnel. Ding et al. (2023) point out statistics from the US where between 2005 and 2010, 566 deaths have been blamed on alert fatigue, which shows how dire the consequences can be. Wang et al. (2024) mention the Target incident, where a critical warning was missed due to cognitive overload and led to what they describe as a "devastating data breach incident" (Wang et al., 2024,p. 2), which goes to show that the consequences are severe in the IT sector as well.

2.2.3 Mitigations

The literature discusses a variation of mitigating actions that could have a positive effect on reducing or preventing alert fatigue. We have split them into two themes: technical and human factor mitigations.

Human factor mitigation

Mitigations that fall under the human factor focus more on the human to help prevent or reduce alert fatigue. This can include education, routines, personality, and other factors where the human is in focus.

Ding et al. (2023) found that the work schedule can affect a healthcare worker's perceived alert fatigue. They explain that night shift workers are describing their shifts as stressful and are often more prone to fatigue and sleepiness. They pay more attention to patient alarms to avoid safety issues during their shifts. A solution proposed by Claudio et al. (2021) is to assess the length of the night shifts and rotate the responsibility for monitoring the alarms. This way one can reduce the probability of alert fatigue by limiting the person's exposure to the alarms.

There is also a variation of other mitigation that focuses on human factors and how they can have an effect on alert fatigue. An example from Robinson (2023) is to hire a consultant, who can evaluate present factors that are leading to alert fatigue and can come up with solutions to improve those factors. Another variation from Claudio et al. (2021) is to use personality assessments to find personality traits that have a relation to the perception of alert fatigue. They used the Big Five personality test to find which personalities work well together and which will be more prone to alert fatigue than others. This evaluation can be important to identify traits in candidates in an interview process. Wilken et al. (2017) are supporting this theory, and point out that factors like attitude, discipline, competency, and communication are human factors that can affect alert fatigue. They add that there is a need for more research into actionable advice to influence these factors.

Technical mitigations

The technical mitigations focus on the technical side to prevent or reduce the occurrence of alert fatigue. The ones that use either machine learning (ML) or artificial intelligence (AI) can be described as cognitive automation where the technology free humans of the burden of physical work (Rinta-Kahila et al., 2023).

One of the most prominent technical mitigations is the tuning of alarms. Tuning of alarms means that the respondents are altering the settings for the alarm so it will react differently to its input, thus removing false positive and nonactionable alarms(Movahedi et al., 2023). If this is done right it can help reduce the number of false positives, but as shown in the previous chapter it must be done with care.

In the health sector, this seems to be used differently depending on the patient's condition and the respondent's opinion. An example from Hravnak et al. (2018) is the

widening of alarm parameter thresholds. Depending on the condition of the patient they can alter how the alarms should evaluate their inputs to avoid alarms for fleeting instabilities in the patient's condition, which would be a form of false positive alarm. Also, Movahedi et al. (2023) identified alarm customization as an effective method to reduce alert fatigue, but they added alert hazard and customization education as an important part of it.

Contextual knowledge is also important when it comes to alerts. In the study by (Alahmadi et al., 2022), they found that the analysts also utilized knowledge about the business, like working hours and information from other third parties, implementing this into alerts will help improve the validation process.

Some technical mitigations attempt to utilize ML to combat alert fatigue. In the article by Ndichu et al. (2021), the authors explore a mitigation method that utilizes the class imbalance problem in multi-appliance security alert data and how to automate the alert analysis process in security operation centers. They utilize different machine learning like the Neighborhood cleaning rule (NRC), which removes noisy and redundant false positive alerts. Support vector machine synthetic minority oversampling technique (SVMSMOTE) is a technique that is used to oversample the true alert data. Decision tree (DT) and random forest (RF) are used after the removal of the false positives and oversampling of true positives, and then DT and RF perform the alert classification. Alert data that has been collected from eight security appliances is used to demonstrate the results of the method they have used. Some of their findings conclude that a majority of alerts are compromised by false alerts, which makes verifying true positives difficult. Using the model, they managed to generate significant improvements that reduced the need for manual auditing and achieved a high recall performance of around 99%. This means that the model has a high probability of catching true positive alerts.

On the other hand, Wang et al. (2024) look at the challenges that are associated with alert triage. They deem that the most important factor is alert prioritization and not alert reduction to mitigate alert fatigue in security analysts. The authors propose a framework called AlertPro that utilizes reinforcement learning to optimize the different results by re-ranking alerts based on historical features from expert feedback. This allows it to go further than just retrieving basic features of an alert but retrieving context features from both the alert sequences and the historic feature from a security expert's input. This framework contains three main modules. The first one is feature extraction, where the basic features are extracted from raw alerts, context features from the alert sequence, and history features from alerts that have already been investigated by a security analyst. The second module is Alert ranking, where the module utilizes the contextual and basic features that have been previously extracted and uses an IF algorithm to calculate an anomaly score; this score is then used to rank the alerts. The third module is Alert re-ranking, which utilizes the historic features and the anomaly score as input and then utilizes reinforcement learning algorithms and expects feedback to re-rank alerts. Five real-world datasets were used to test the framework. AlertPro was able to increase the precision of multi-step attack identification in anomaly detection by utilizing the context features. Additionally, it successfully filtered out alerts that were seen as low risk and highlighted those that were seen as high-risk. This was done by

taking advantage of the active learning model and incorporating the expert knowledge from their history features feedback.

McElwee et al. (2017) Investigate a solution for the problem of too many alerts and the challenge for security analysts to manually review IDS alerts where systems like IDS and SIEMS are inadequate for the analyst to quickly determine if it is an incident. The authors, therefore, propose a system called Federated Analysis Security Triage Tool (FASTT), this system utilizes TensorFlow Deep Neural Network (DNN) classifier to automatically classify IDS alerts, decide which ones a security analyst should examine, which ones can be distributed to interested parties, and which ones should be disregarded. The system has been built on previous research that exists within IDS machine learning and deep neural networks. To determine the efficiency of the system it was evaluated in two ways. The first one was classification accuracy, the second was usability, and the value it provided in automation. FASTT had a classification accuracy of around 98% percent, which was computed using a 5-fold cross-validation using an out-of-data sample. The neural network that matched the fold that obtained the best accuracy was the final neural network produced by this study. The usability and automation were also deemed excellent; this allowed the security analyst to save time when it came to investigating alerts, narrowing down the scope of their review, and making it easier to prioritize alerts.

In the article by Raff et al. (2020) they investigate how a type of machine learning called passive aggressive (PA) can help fix false positives within a model already in production. They do not disagree with the claims made by numerous security firms that their machine learning is 99.99 percent successful. However, they mention that these systems are trained on the global representation of data, and this could be very different when it comes to the local environment of a company. In previous studies, it has been shown that different anti-virus systems (AV) have different false positives (FP) and false negatives (FN) for malware from different areas of the world. This could generate a lot of false alerts and lessen the trust in the malware detection of the AV. Their primary approach utilizes the PA algorithm where a user will submit a false positive to the model. The model must then, based on a set of variables, fully correct the error so it's no longer seen as a false positive. It also utilizes Area Under the Curve (AUC) to prevent users from mislabeling alerts. Additionally, as their foundation the research uses a MalConv based model, also known as malware detection model. As a result, they found that the model could be used in addition to other techniques like Gradient Boosted Decision Trees (GBDT) to utilize its full potential. When it comes to global values the model that utilized MalConv + GBDT had better performance than the Malconv + PA. They therefore recommend using the MalConv + GBDT and when hard FP were identified they recommended switching to Malconv + PA.

McRee (2022), focuses on usability and ease of use when it comes to the security alert output from both data science and machine learning. A dimensional study showed that minor problems or noise, wasted time chasing false positives, feeling overwhelmed, excessive time used to prioritize alerts and increased overall security risk were some of the primary challenges. Prioritization and summarization are therefore profound to reduce the excessive amount of information presented. Thus, the writers aimed to

determine if there could be increased efficiency by looking into their perception and usability of text-based alert output versus visualized alert output. They utilized a theoretical framework known as the Technology acceptance model, in an attempt to answer their research questions, and utilized a quantitative, quasi-experiential, explanatory methodology for the study. They received their data through an online survey, and their results proved that there was a significant difference in the perception and that it favored the visualized alert output significantly.

3 Theoretical framework

In this chapter, the theoretical framework of High-reliability organizations (HRO) will be introduced as our theoretical lens. By utilizing a theoretical lens one can examine a topic or domain from a different perspective. This could illuminate relationships that could in other ways be hidden because of too much data/information (Niederman & March, 2019). We are looking at a phenomenon whose fundamentals haven't been thoroughly investigated in the information systems or security field, therefore, HRO will serve as our theoretical lens.

The lens was chosen based on our presumption that SOC is considered a digital HRO. A digital HRO is an organization that develops and builds both services and solutions used to protect an organization's IT systems from threats like computer malware and malicious software (Salovaara et al., 2019). One of the main traits of HROs is that they utilize the cognitive mindset of mindfulness in their operations. By using the lens of HRO, we can look at mindful operations in SOC and find the relationship between mindfulness and alert fatigue.

3.1 HRO Model

The foundation of the HRO framework is based on companies that handle complicated systems almost entirely error-free. A primary priority for HROs is to focus and learn from an operation's failure instead of its successes. Additionally, prioritizes reliability over efficiency (Salovaara et al., 2019). HROs can manage this kind of reliability of performance because they exhibit the characteristics of these five capabilities together, referred to as "mindful organizing". This means that the HROs spend:

"(a) more time examining failure as a window on the health of the system, (b) more time resisting the urge to simplify assumptions about the world, (c) more time observing operations and their effects, (d) more time developing resilience to manage unexpected events, and (e) more time locating local expertise and creating a climate of deference to those experts." (Weick & Sutcliffe 2006 p. 516)

Mindful organizing is the defining aspect that makes an organization an HRO. This cognitive mindset helps to handle normal, accidental, and inevitable threats on a collective level. HROs focus on having this mindset throughout the organization, which can be seen manifested in its five system-level characteristics:

Preoccupation with failure (more time examining failure as a window on the health of the system): HRO maintains importance on the risk of normal failures. Viewing almost failures as learning opportunities. Additionally focusing on understanding events and situations that rarely occur or have never occurred (Salovaara et al., 2019).

Reluctance to simplify interpretations (more time resisting the urge to simplify assumptions about the world): HROs attempt to minimize blind spots by utilizing healthy skepticism about clear interpretations in their sense-making when studying failures and their pathways. Conversely, they encourage actors to be vigilant toward subtle anomalies and early warning signals (Salovaara et al., 2019).

Sensivity to operations (more time observing operations and their effects): HROs create a holistic view of their operations and environment, asking participants to continuously evaluate any potential relationships between occurrences for more well-informed decisions on consequent actions(Salovaara et al., 2019).

Commitment to resilience (more time developing resilience to manage unexpected events): HROs manage surprises "In the moment" by rapidly creating untried approaches, existing in a make-do mentality, and learning from previous experiences that usually only involved a few samples (Salovaara et al., 2019).

Under-specification of structures (more time locating local expertise and creating a climate of deference to those experts.): HROs create multi-functional processes, which encompass a wide range of experts, that can smoothly take responsibility for different aspects of problem-solving as a necessity to generate flexible, dynamic responses to the escalation of a failure (Salovaara et al., 2019).

The concept of mindfulness is seen as a state of active awareness distinguished by the ongoing creation and clarification of categories, receptivity to new knowledge, and readiness to consider situations from several angles. The other concept is called mindlessness. When fewer cognitive processes are running it usually results in one of the characteristics of mindlessness where one is dependent on past categories, acting on "automatic pilot" and only focusing on a single perspective without the realization that things could be differently (Levintha & Rerup, 2006).

To be able to identify mindful activities, we have adopted the concept of Mindful Conduction by Rinta-Kahila et al. (2023). The concept is a combination of three facets that describe "*modes of working where the worker is consciously engaged in the* task" (Rinta-Kahila et al., 2023, p. 1387) .These three facets are "*activity awareness*", "*competence maintenance*", and "*output assessment*". In contrast to Mindful Organizing, which focuses on mindfulness on a collective level, Mindful Conduction focuses on mindfulness on the individual level. By looking for these facets in the way an individual operates, we can identify if they are mindful or mindless in their activity. For "*activity awareness*" we can identify if they are aware of what they do and why they do it. "*Competence maintenance*" will involve learning more about the topic of their tasks and wanting to evolve their knowledge of it. "*Output assessment*" are actions they can perform by evaluating the result of their work to see if it meets a set of criteria, or if the desired outcome is achieved (Rinta-Kahila et al., 2023). In the absence of these facets, we can be sure that an individual is acting mindlessly.

3.2 Mindfulness and Mindlessness

Salovaara et al. (2019) have created a framework to analyze HROs. The framework is split into three dimensions: Nature of operation, nature of cognition, and purpose. The first two dimensions are based on the characteristics of HRO. They explain that the nature of the operation is either human-based or digital, and the nature of the cognition is mindful or mindless, meaning "*the sensitivity to and anticipation of surprises*" (Salovaara et al., 2019, p. 7). The last dimension is adopted from the classifications found in Kirsh et al. (1994), where they document two types of operations in the

cognitive system: epistemic and pragmatic. They define these as "*the purpose of the operations making up the organization's activity flow*" (Salovaara et al., 2019, p. 7).

Three dimensions of analyzing high-reliability digital operation		
Feature	Feature type	
Nature of operation	Human-based: approximate, error-prone, limited by memory capacity and processing speed, of varying precision, context-sensitive	Digital: exact, transferable, editable and programmable via expression of binary data.
Nature of cognition	Mindful: heedful, with anticipation of surprises and prioritization of safety in operations, unconstrained by the frame problem	Mindless: constrained by the frame problem via algorithm-use or reliance on highly structured routines
Purpose	Epistemic: interpreting and analyzing information	Pragmatic: performing decision-making and acting

Table 3: Dimensions of a digital HRO (Salovaara et al. (2019))

The table above describes the dimensions and features, which as a whole describe two opposing operations: human and digital. To elaborate on the framework, Salovaara et al. (2019) explain that the area on the left side of the table depicts the "*mindful human-based operations with capacity for context-sensitive processing, imagination, and bricolage* [...]" (Salovaara et al., 2019, p. 7), and that it can be victim to the entrenchment problem. The entrenchment problem is a human´s inability to see the problem in a context outside their immediate proximity. They describe the right side as the "*operations characterized by algorithmic processing and mindlessness stemming from the frame problem*" (Salovaara et al., 2019, p. 7). The frame problem is that digital-based operations cannot see context outside their predefined rules.

The frame problem and the entrenchment problem are similar problems that occur in human- and digital-based operations. People can struggle to switch perspectives, react to unexpected events, and dismiss new or contradicting information. The difference between the two problems is that people can be reminded to think mindfully if there is a trigger that points out their mindlessness, while we are unable to do the same for digital-based operations (Salovaara et al., 2019).

We have made an adaptation of this model with alert fatigue in mind to fit our definition of alert fatigue and its effect on operations. We believe that the model created by Salovaara et al. (2019) could be used since we find the nature of the operation to be well described and could be fitting with the human and digital operation in the SOC: digital alarms go off and are acted upon by the predefined rules or are sent to a human who will use the information and its context to act. This description also identifies the purpose of the human and digital operations, which is to act or to interpret. We think the SOC might have a different nature of cognition. Based on the framework by Salovaara et al. (2019)

humans have the ability to act mindfully, and they are anticipating surprises and alert. However, we believe when humans are affected by alert fatigue, they are no longer able to act mindfully and are starting to act mindlessly. They will instead of being heedful start to act on autopilot, taking decisions without the prioritization of safety. Therefore, we have created a potential framework that considers this possibility as well, based on the framework by Salovaara et al. (2019).

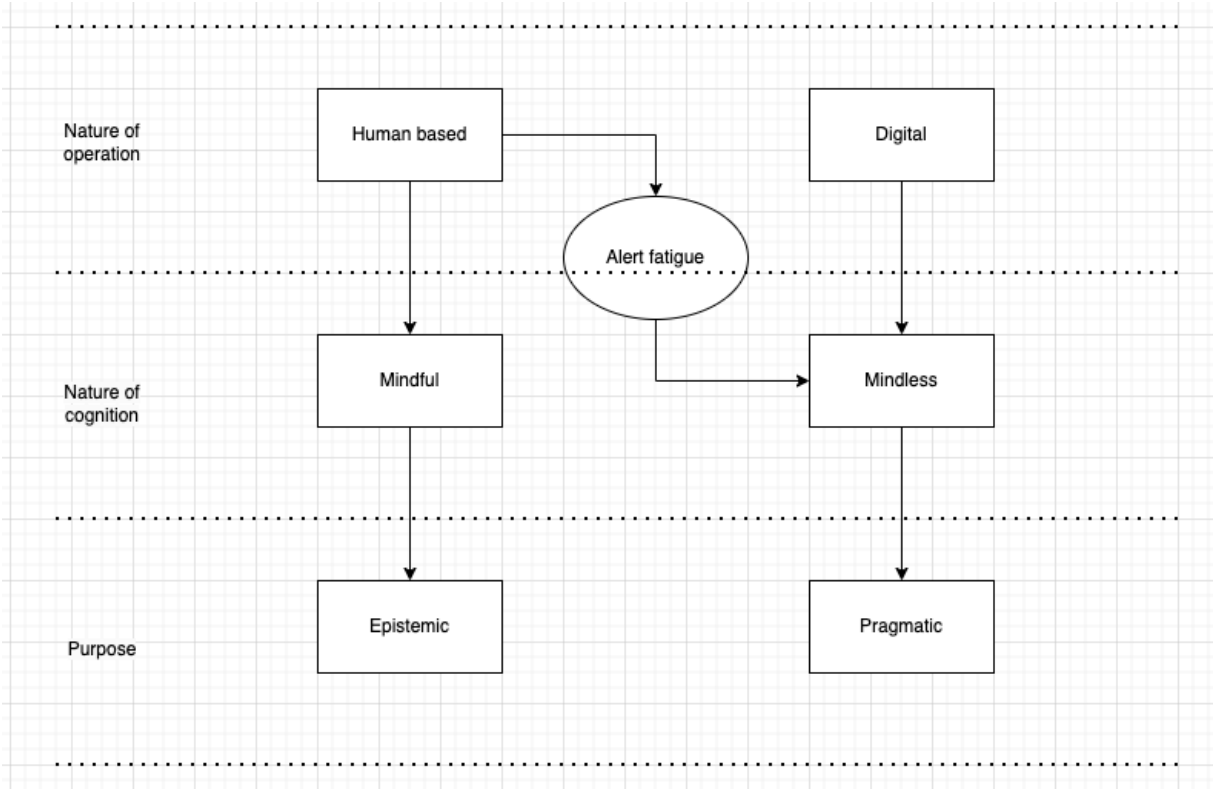

Figure 3: Adaptation of dimensions flow diagram

Our model is a flow diagram that shows the connection between the dimensions. These dimensions have been redefined to specifically represent the operations found in a SOC rather than in generic terms for a Digital HRO.

Features	Feature type	
Nature of the operation	**Human-based**: operations performed by an analyst based on information provided by a digital system.	**Digital-based**: operations performed by a digital system, performing actions, or providing information to the analysts.
Nature of cognition	**Mindful**: operations that are performed in a heedful	**Mindless**: constrained by the frame problem,

	manner, anticipating surprises and prioritization of safety in operations, unconstrained by the frame problem.	unaware of context outside of its predefined or assumed rules. This can be a result of alert fatigue or when analysts preserve cognitive stamina to stay mindful in higher-prioritized situations.
Purpose	**Epistemic**: interpreting and analyzing information	**Pragmatic**: performing decision-making and acting

Table 4: Adaptation of dimensions table

Table 4 shows how mindlessness can occur in human-based operations as well as in digital-based operations. We think that mindful conduction in combination with this framework will aid us in visualizing how alert fatigue can affect mindful operations.

3.3 Negative implications of mindfulness

Ault & Brandley (2023) cite Weick & Sutcliffe(2007) in their article, where they suggest one should always be mindful in an HRO, even when it's calm. Ault & Brandley (2023), who have researched the drawbacks of HRO, have found that there is a cost to staying mindful all the time. They found that mindfulness is a resource that will be exhausted, and when it’s exhausted, one will return to a default state. This default state is, according to Ashforth & Fried, (1988) cited by Ault & Brandley (2023), the mindless state. They also found that it takes cognitive power to be able to maintain their mindfulness and that it will cause negative effects when that cognitive power is depleted.

4 Research approach

This master thesis goal is to understand how mindful organizing occurs in the SOC, how alert fatigue potentially influences this, and what kind of mitigations are currently utilized to combat it. The following research questions (RQ) need to be answered to address this.

- RQ1: How does mindful organizing happen in the SOC environment?
- RQ2: How does alert fatigue influence mindful conduction in SOC?
- RQ3: What strategies are being implemented to mitigate alert fatigue?

The following chapter will discuss the approach used for this master's thesis research, which includes planning, designing, preparing, data collection, and data analysis. The selected research design is a multiple case study. When conducting a case study, one usually follows an iterative process. Our process was done following Yin's 2009 iterative process diagram (Yin, 2009).

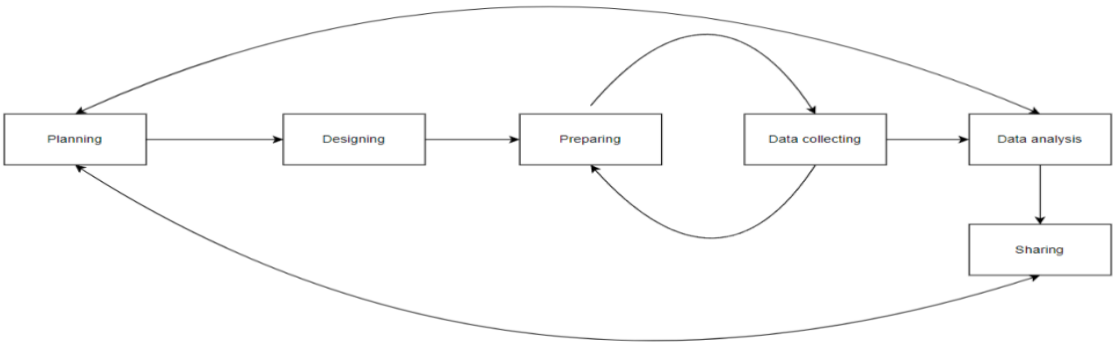

Figure 4: Case study procedure (Recker, 2021)

4.1 Planning

In the planning phase, the identification of research questions and rationales for choosing a case study is usually done (Recker, 2021). We have previously stated the research questions we want to answer through this thesis. A research approach is typically split into two categories: the approach of data collection and the approach of data analysis or reasoning. The data collection section is divided into qualitative and quantitative (Chetty, 2016). After the selected approach, a rationale is given for why the case study approach was utilized.

4.1.1 Qualitative approach

A qualitative approach has been selected to help answer the research questions. Recker (2021) explains qualitative research as: "*Qualitative research empathizes understanding phenomena through direct observation, communication with participants, or analysis of texts and may stress contextual subjective accuracy over generality*" (Recker, 2021, p. 115).

This study seeks to obtain a deeper understanding of how security analysts face alert fatigue through the theoretical lens of HRO. Therefore, it is important to directly communicate with participants who are in situations where this issue might occur. Individuals and organizations can potentially have different ways of mitigating or preventing the issue including different views of what alert fatigue is. Therefore, there is a need for more contextual subjective accuracy than attempting to solve this by generality.

To investigate and understand the phenomenon and the types of mitigations implemented, a group of individuals who work at a SOC will be interviewed. The interviews will be semi-structured. A semi-structured interview is usually performed with an incomplete script where some questions are prepared beforehand, but there is a need for improvisation (Myers & Newman, 2007). This will help us to get answers to some important questions we will prepare and let the interviewee speak more freely on the topic and how they experience it.

4.1.2 Rationale

The table below will help to rationale why the multiple case study was the selected approach for this thesis. The description of the different categories is from Recker (2021).

	Spectrum	
Aim	Exploratory	Explanatory
Method	Qualitative	Quantitative
Boundary	Case	Statistical
Setting	Field	Laboratory
Timing	Cross-sectional	Longitudinal
Outcome	Descriptive	Casual
Ambition	Analyzing	Designing

Table 5: Case study spectrum (Recker, 2021)

- **Aim:** The aim section of the table explains the overall aim of the research. The purpose of Exploratory research is to discover new insights, find out what is happening, or one can attempt to assess a phenomenon in a new light. Meanwhile, explanatory research aims to determine the connections between various variables (Makri & Neely, 2021). This study will mainly be exploratory research since we want to explore how mindful organizing happened in the SOC environment and look at the phenomenon of alert fatigue and its influence on

mindful conduction. Additionally, we want to seek out what kind of mitigation strategies are used.

- **Method:** The method specifies the approach that can either be qualitative or quantitative. In this study, we will use a qualitative rather than quantitative method to get in-depth data to gain a better understanding of how mindful organizing occurs, how they experience alert fatigue and its influence on mindful conduction, and their mitigation strategies.

- **Boundary:** Specifies the kind of study performed whether it will be case related or statistical. As we are interviewing multiple subjects over multiple organizations our type of study will involve multiple cases.

- **Setting:** This is about the place where the research will be conducted. As our interview will be at the location of the respective interviewees, our setting will be field.

- **Timing:** This is about the length or period of time the research will be conducted. As time is limited, the research will fall under the Cross-sectional category for timing as opposed to Longitudinal.

- **Outcome:** This part is about the outcome of the research whether it's descriptive or casual or in between. Our research is in between since it is on the casual side and the descriptive side. Firstly, investigating mindful organizing and mitigation strategies is more on the descriptive side, attempting to describe them, while alert fatigue influence on mindful conduction is more on the casual side to see a cause and effect.

- **Ambition:** Ambitions are seen as the end goal of research and are split into Analyzing and Designing. Our ambition falls into the category of Analyzing, as we want to understand mindful organizing and alert fatigue influence on mindful conduction and mitigation strategies.

The table below shows what kind of spectrum this study falls on.

This study	
Aim	Exploratory
Method	Qualitative
Boundary	Multiple Cases
Setting	Field

Timing	Cross-sectional
Outcome	Descriptive / Casual
Ambition	Analyzing

Table 6: Research design spectrum

4.2 Designing

The design phase involves defining a unit of analysis, the number, and types of cases to be studied, and theories or propositions to guide the study (Recker, 2021). When designing a case study there are three different variants: Descriptive, Exploratory, and Explanatory. The descriptive case study attempts to describe a phenomenon, the exploratory looks for causal factors to describe a specific phenomenon, and the exploratory attempts to study a phenomenon with the intent of exploring or identifying new research questions (A. Priya, 2020). Our study falls into the exploratory variation as previously mentioned.

As the thesis will examine several businesses operating in the same cybersecurity sector, this study will make use of multiple case studies. Investigating mindful organizing and how alert fatigue influences mindfulness, including the different ways an organization attempts to battle the phenomenon. By utilizing a holistic view, we can explore different factors in separate organizations and see how they contribute to and or attempt to mitigate alert fatigue (Recker, 2021).

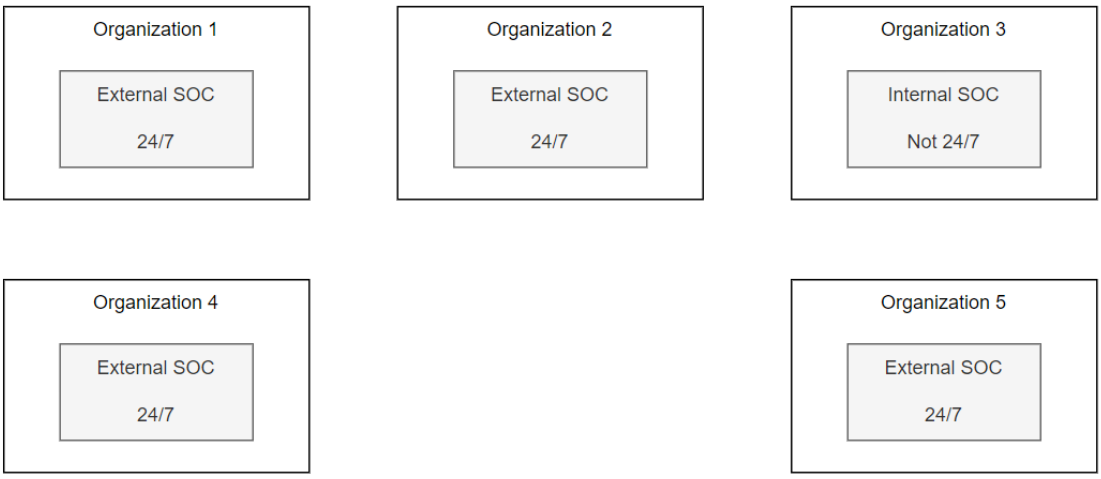

Figure 5: Holistic case study (Recker, 2021)

Most of the cases are similar when it comes to having active surveillance 24/7 and working as external SOCs for other companies that work in different sectors. Case number 3 is the only one that differs as they are an internal SOC. They protect their organizations and they do not operate 24/7, however, they have a different external SOC

that protects them in the evenings and at night. As we have promised each SOC full anonymity, we have removed any information that could be used for Open-Source Intelligence (OSINT). There are few SOCs in Norway, and including any more information about the cases could compromise their identity.

4.3 Preparing

In this phase, the researchers hone their data-collection skills both in interviewing and observation. The case study protocols are revised, developed, and finalized the method of data collection is practiced, and pilot-tested (Recker, 2021). We conducted a pre-study before this thesis where we performed interviews with similar topics. This process provided some insights into our topic and interview style that we could take with us in the next iteration of interviews. Additionally, all questions that were prepared beforehand were tested and changed both during testing and after interviews when interesting answers occurred.

4.4 Data collection procedure

When using a qualitative approach, there are various methods for gathering information. This could either be interviews, observations, other documented sources, or combining several data sources(Recker, 2021). The data collection method that will be used in this study is interviews.

4.4.1 Selection of interview subjects

A total of twenty-two companies were contacted to gather the required number of candidates. Purposive sampling was performed during this process. Purpose sampling is "*Identification and selection of individuals or groups of individuals or group of individuals that are proficient and well-informed with a phenomenon of interest. In addition to knowledge and experience*" (Etikan et al., 2015, p. 2). Following a brief screening procedure with some firms failing to respond, the list of companies was reduced to five. From these five companies, twelve interviewees were found ready to conduct the interviews.

Case	Interviewee	Role
1	1	Security analyst
1	2	Security analyst
1	3	Security analyst
2	4	Security analyst

2	5	Security analyst
3	6	Security analyst
3	7	Security analyst
4	8	Security analyst
4	9	Security analyst
5	10	Operations
5	11	Security analyst
5	12	Detection Engineer

Table 77: Overview of interviewees

Since there is a need for a high level of confidentiality, much of the information on the companies will remain anonymous to protect both the interview objects and the companies that they work for. The information found could be both negative and positive for the companies, so a high level of confidentiality was one of the conditions that were agreed upon for them to take part in the study.

4.4.2 Semi-structured interviews

The interview method that was utilized was semi-structured interviews. This was done since we wanted to use more open-ended questions to better understand individual ideas about alert fatigue and their actions. In a semi-structured interview, the script is not fully complete, however, some questions are prepared beforehand that can be used, and some improvisations are needed (Myers & Newman, 2007). When performing a semi-structured interview, the questions are usually open-ended and followed up by some why or how questions. The discussion is usually around the topic, however by doing so it also may fall into unforeseen issues during the discussions (Adams, 2015).

4.4.3 Limitation of semi-structured interviews.

Interviews as a type of data collection have various degrees of limitations depending on what kind of interview the researcher aims to use. Semi-structured interviews are very time-consuming and labor-intensive which makes it difficult to conduct data collection on a bigger scale and makes it hard to yield a large enough sample to have precision on the "*Plus or minus n percent*" (Adams, 2015). It also requires a skilled interviewer who has a lot of knowledge about the issue being discussed (Adams, 2015). Lacking in this area could lead to a more biased interview where the interviewer can guide the subject to say what they want to hear creating a bias for the data, which makes it less reliable and valid. Additionally, depending on different levels of confidentiality could lead to more or less information from the subjects.

It is crucial to understand the limitations and take measures to reduce the likelihood of them occurring, the data cannot be corrected after it is collected. Therefore, in this study, we are two researchers so we can lessen the labor intensity of the interviews by sharing the workload and making it less time-consuming. We have conducted previous interviews to sharpen our skills in conducting interviews and expand our knowledge in the area by performing a literature review of relevant literature.

4.5 Data analysis

One of the important parts of qualitative research is data analysis. This part has a major influence on the result of the research as it represents one out of many steps in the research process (Mayer, 2016). A key attribute of performing data analysis in qualitative research is the amount of data that needs to be analyzed. There is a variety of analytical strategies since qualitative data can be seen from any number of different perspectives (Recker, 2021).

When identifying what strategy will be best suited for the research different aspects have been considered. The strategy should fit into the type of data collection that has been used, how difficult it is to use, and lastly if it will provide sufficient results for the research problem. Thematic Content Analysis (TCA) is the chosen strategy to be conducted in this study's data analysis section, as it works well with semi-structured interviews (Anderson, 2007). Additionally, this allows us to create different themes from the literature regarding our theoretical lens, investigating things in the interview that could be seen as mindful, mindless, and any type of mindful organizing, including themes that are indicators for alert fatigue.

Thematic content analysis is a descriptive presentation of qualitative data. Qualitative data can take different forms of information; however, TCA is mainly limited to textural data (Anderson, 2007). In this strategy, the researchers will pinpoint themes in the text and create a thematic overview of the interviews. There are two main ways of doing this, which are inductive and deductive. In an inductive TCA, one should create themes as they emerge in the data you are analyzing. The themes will change as you get to know the data better and will merge and split into new themes as the process proceeds. In a deductive TCA, the themes are made prior to the analysis and are fixed during the analysis. In this study, we have performed a combination of the two. We have considered the related literature when coding the interviews and letting the themes emerge from its contents. Later, we merged them into broader themes that are closer to the literature, such as the codes in a deductive approach.

The TCA can be split into 6 steps (Caulfield, 2019)

1. Transcribe
2. Code
3. Generate themes
4. Review themes
5. Defining and naming themes

6. Writing up

In the first step of the analysis, the interviews must be transcribed into text. Teams were selected to use during the interviews since it has a tool that transcribes everything that is said. However, we had to go through the text and fix typos or grammatical errors the transcriber made.

After transcribing the interviews, the next step is to code the information which was done using the software NVivo. Here the main effort was to find information that was relevant to the research. Information that could be beneficial for the research was inserted into different codes as they emerged in the interviews. This step was performed by each researcher individually to be able to find more information and to reduce the chance of something being overlooked. We had a meeting after the first two interviews were coded to compare our codes and to ensure that we were looking for the same things. For example, both of us had coded the following sentence as *suppress*. "*It is in that way we suppress, if there is something we have reported again and again, then we have the opportunity to add a note to it before, if we need to, and close it and the offense then stops. the offense from collecting events*" (Interviewee 11, SOC5). By comparing the code and the content for multiple codes, we knew that we were looking for the same information.

The third step is to generate themes. Both researchers read the codes that were found in the prior step, and the code was placed under themes. An example of a theme that was created and its code is:

Themes	**Codes**	**Raw Translated**
AF Mitigations	Suppress	"It would be natural to turn on a filter that removes unnecessary information. This allows you to focus on alerts that could potentially be real threats." (Interviewee 2, SOC1)
AF Mitigations	Rules	"Yes, I would dare to say so, since you can be very general in a rule. You create it so it can cover a lot, however, it would then generate a lot of false positives". (interviewee 12, SOC5)

AF Mitigations	Attitude	“We have worked hard to turn that perspective in our analysts, they should not think that they are under time pressure, but rather perform a good analysis.” (Interviewee 5, SOC2)
AF Mitigations	Personal mitigations	“Take a little break, take a breather. Where I currently work, we are lucky enough people to keep the analysis quality up, so there is room to take a breather, maybe also just going for a walk to get some fresh air” (Interviewee 7, SOC3)

Table 88: Overview of themes, codes and raw translate

Some themes are extracted from an existing code, and others are created as a generalized version of more specific codes. This step also allowed the researcher to discuss the meaning of the codes, what their importance was, and how they could be fitted to a theme.

Following this discussion, we enter step 4 where the themes are reviewed. The review will make sure that the newly defined themes are represented in the data and backed up by the codes.

The next step is to name and explain the themes, where the researchers will rename themes and add descriptions to the themes to make them more understandable. The themes and their description are in the table below.

Theme	**Description**	**Raw Translate**

Reluctance to simplify interpretations	Indicators for the company's efforts to be skeptical of anomalies and early warning signs.	"yes, usually I solve it the way I think it works, however, I often want an extra set of eyes to see if they agree with my assessment. Just to be completely sure, additionally, I won't be standing in it alone if something were to happen" (Interviewee 8, SOC4)
Commitment to resilience	The company's efforts to be resilient and their ability to turn around and react to the unforeseen.	"Yes, and people have responsibility for it, they can be contacted 24/7. We perform surveillance every second of the day and year. So if a system goes down, then it has to have been up "yesterday" (Interviewee 11, SOC5)
Preoccupation with failure	How the company focuses on not failing and taking almost-failures seriously. Having pre-made plans if the unexpected should happen.	"We have one that is writing detection rules, and the analysts that are using these rules. We have weekly meetings, or when we see the need for them. But a weekly meeting has been set up so that everyone can go through new rules, and see if changes need to be made to improve anything." (Interviewee 7, SOC3)
Sensitivity to operations	The company's ability to take new information and adapt the way they work to mitigate potential new risks.	"It can be bad if we make a mistake. So, we mustn't make mistakes" (Interviewee 9, SOC4)

Under specification of structure	The company's flexibility when it comes to escalation and decision making.	"We got different escalation processes, it's based on what kind of incident it is. We have an Incident Response Team (IRT) that could be sent out when real attacks occur and help the customer. I can also escalate to blue team personnel who can give more support or have more access to the system. So, we usually use 2 main processes, that's at least what I usually use." (Interviewee 11, SOC5)
Mindful	Actions that the analysts do that can be described as mindful.	"Of course, sometimes green ones are just as dangerous, however, we just can't see it that well. So then it's more about coming back to it performing an analysis and seeing more of the total picture of what has occurred." (Interviewee 3, SOC1)
Mindless	Actions that the analysts do can be described as mindless.	"It's a danger for that, that you can lose a bit of the context in the alert, so you might analyze it a bit quicker than what you should have done." (Interviewee 7, SOC3)
AI and ML	The occurrence and use of AI and ML in SOC systems.	"To my knowledge, AI is not used today. So, if this is specifically aimed at AI then no, there is no AI." (Interviewee 2, SOC1)
Overwhelmed	Indicators and reflections from the participants that point towards them being	"It does happen from time to time, it depends on the week and the time of the day. The

	overwhelmed in situations in the SOC	main reason one can feel overwhelmed in this type of work is having to work on large number of alerts." (Interviewee 2, SOC1)
Trust	Reflections from the analysts on trust in existing systems, alerts, criticality, etc. in the SOC environment.	"Yes, it's something about the subconscious, so if you trained yourself up not on purpose, however, this occurs every day. This is something that is not dangerous, looking at things with that mindset, including low levels of trust. Then when investigating a case, you will probably not look that hard after dangerous things." (Interviewee 12, SOC5)
Alert fatigue	Indicators and reflections related to alert fatigue.	"I seen it become a big problem, people have gone on leave." (Interviewee 11, SOC5)
False Positive	How often the analysts experience false positives and what percentage the total alerts are false positive.	"There are a lot of false positives, to a very big degree. I don't know the direct percentage, however it's very high." (Interviewee 7 SOC3)
Context Switching	Experiences related to switching contexts during shifts at the SOC.	"I feel less fatigued, if I don't have to context switch too much" (Interviewee 2, SOC1)

Table 99: Overview of themes, description and raw translate

The last step is to write up the findings. The findings chapter will contain the themes as headings and will consist of what was found in the interviews. The chapter will try to answer the research question and provide further knowledge on the subject. Since all the data collection has been performed in Norwegian, the quotes that can be found in the findings have been translated from Norwegian to English. During the translation of

the quotes, some readjustments and grammar changes have been performed to not lose the essence of what they said.

4.6 Validation of findings

Validating the findings is important to increase the validity and credibility of the research that has been conducted (Bans-Akutey & Tiimub, 2021). Data source triangulation was used to validate the findings of this study. *"Data source Triangulation involves collecting data from different types of people, including individuals, groups, families, and communities, to gain multiple perspectives and validation of data"* (Carter et al., 2014, p. 545).

This has been conducted by reviewing the answers from the interviewees and comparing them to look for similarities. Additionally, purposive sampling has been used. The selection of interview subjects has been mainly personnel in a security operation center; this was done to retrieve data that would be of most value. When conducting the data analysis, it is important to note that we, as the researchers, are biased and could potentially mainly look for data that would support our assumptions.

4.7 Ethical consideration

When conducting scientific research, it is important to follow some ethical guidelines to ensure that one follows the principles of good research practice (Mirza et al., 2023). Some of the ethical guidelines that were followed were ethics of respect, Informed consent, Confidentiality, and Anonymity.

Ethics of respect: Everyone who has been part of the study must be treated with respect and trust. This includes that the research that is being conducted must be done with the respect of the individual in mind regardless of their sex, age, race, religion, political belief, lifestyle, or other differences that are significant when it comes to differences between the induvial and us the researchers (Mirza et al., 2023).

Informed consent: As researchers, we must send out a "*Voluntary informed consent*" to all the individuals taking part in the study before collecting their data. It should include the main aim, and objective of the research and some of the ethical information like confidentiality and anonymity. This must be done to seek their consent for the collection of data and be signed before the data collection occurs (Mirza et al., 2023).

Confidentiality and Anonymity: The researchers should do their utmost to protect the anonymity of the subjects taking part in the study, and the privacy of their data. This must be communicated to the individuals taking part in the study, including inserting it into a written agreement between the researchers and the individuals (Mirza et al., 2023).

In this study, the personal information of the different participants was deleted very rapidly after the interview was transcribed. This was to enforce a high level of confidentiality when it came to the information of the subjects and their respective

companies. Any information that mentions the names or size of the company or any identifiable critical system information was also anonymized. This information was given verbally to the subjects, and they were given a document from SIKT that explained what kind of information would be collected, where it would be stored, and for how long. This was done to ensure that the participants understood what was collected and the objective of the study. This was also done so the participants were more willing to share information. They were also given a verbal explanation of what the information would be used for including how it would be anonymized and how quickly it would be deleted during the study. Lastly, most of the quotes that were used in this thesis were sent back to the interviewees to receive confirmation of their usage and to ensure that what stood was correctly interpreted.

5 Findings

In this section, we will discuss the findings that have been found in relation to our research questions:

- RQ1: How does mindful organizing happen in the SOC environment?
- RQ2: How does alert fatigue influence mindful conduction in SOC?
- RQ3: What strategies are being implemented to mitigate alert fatigue?

Our findings show that the difference in how the SOC environments work is insignificant. The major differences are company size and age, but because of confidentiality we will not share any more specifics that can disclose their identity. The findings chapter will therefor present the findings from the interviewees point of view and will lay focus on their experience.

We have discovered actions that exhibit how mindful organizing occurs in their environment in all the five different capabilities. This includes, among other things, testing rules before they are implemented, having people on standby 24/7, and to be ready for unexpected events. We have also made discoveries when it comes to mindful conduction and alert fatigue. We have identified that when analysts face a large number of alerts, they can begin acting more mindlessly, and start having trouble with mindful conduction. This is indicated by interviewees discussing something they call “zombie - mode” where they perform actions automatically. Interestingly some analysts might knowingly perform mindless actions, but they do this so they can be more mindful when something more serious occurs. There have also been identified various mitigation strategies, which are split into three categories: mitigating of causes, mitigation of desensitization and cognitive overload, and mitigating consequences. Surprisingly, there is almost no use of AI or ML in SOCs currently. Some of the reasoning behind this is the lack of trust, and they might even cause more noise. Additionally, one surprising finding was that alert fatigue is not observed in internal SOC, but it does appear to be a problem in their counterpart: external SOC.

5.1 What is it like to work in a SOC?

The 5 SOCs we have studied are similar in most areas. All except SOC number 3 are external SOCs with 24/7, 365 days surveillance, while SOC 3 is an internal SOC managed from 8 to 16 each day. The latter also use an external SOC that informs them about security incidents that are not monitored by themselves. What they do have in common with the external SOC is that they have shifts that they call dayshift, evening shift, and night shift, which range from 8 to 12 hours.

A lot of the analyst's tasks in the different SOCs are similar as they hold similar positions, but there are some who are given more responsibilities or have more time allocated to perform additional tasks. Interviewees 2 and 11 explain their normal tasks as analysts. They will analyze alarms, evaluate the incident, contact customers, and perform some action if it’s relevant to the incident.

“[...] My job is to receive alarms from various sources, analyze the alarms, and decide if it’s a security incident and if so, notify the customer. Also, potentially give tips and information about recommended measures or things like that.” [Interviewee 2, SOC1]

“[...] My job is to investigate the network alerts and, should we say, all the artifacts that possibly would be part of that security incident and then evaluate whether or not it's a real security incident and then evaluate on how serious that threat is, if it's a true positive.” [Interviewee 11, SOC5]

Some of them have additional tasks that do not involve analyzing alerts. Interviewee 3 uses some of their time as a detection engineer, which means that they are creating or editing rules used to detect anomalies that could be potential attacks and thus should result in alarms.

“I'm a security analyst [...] I'm also doing a bit of detection engineering” [Interviewee 3, SOC1]

Interviewee 4’s additional tasks are similar to Interviewee 3's, as they monitor the SOC environment and fix areas that can be improved. They are also involved in meeting with the SOC's product owners, whose job is to add new products, alarms, and signatures used for anomaly detection.

“When I'm not on duty, I do everything from meetings with product owners, who are the ones pushing products and signatures and alarms to the SOC. And I am working with statistics [...] and monitoring the incident flow in our SOC and looking for things to fix or improve”. [Interviewee 4, SOC2]

One of the SOC's main tasks is to determine whether something is out of the ordinary and a security incident or just normal activity that poses no threats to the organization. Interviewee 9 talks about how this was an issue when they started their position at the SOC. They say the experience has made the job easier over time, but new alerts can still make it a bit harder.

“[...] it was incredibly difficult at first, I think. And I can figure out what's what, but over time, it's really gotten easier, but it's still... There will be new alarms that I haven't seen before, and then you sit there and wonder, where exactly should I put this? How should I classify this? But it has become easier to do so”. [Interviewee 9, SOC4]

To determine if an alarm is a true positive, meaning it is a real threat, the analyst needs to look at the context in which it emerges. The context will tell the analyst what happened in that system during the time the alarm was triggered and will include traces of activity. Based on the context, they can determine whether the alarm was triggered by malicious activity and should be escalated or if it was a false positive. How difficult this is can vary from alarm to alarm. Too much context means that there is too much information related to the alert, and if it’s too little, then there is not enough information. Interviewee 2 discusses that too much and too little context can make it hard to determine the nature of the alarm, in addition to the technical complexity of the alarm.

“[…] it can be very demanding if there is little context. However, there could also be a lot of context, but it's hard to know if it’s something serious depending on how technically complicated it is” [Interviewee 2, SOC1]

In the SOC, there are other roles besides analysts. Interviewees 12 and 10, for example, are working with SOC operations and detection engineering. Their job is to create new rules for detection and make sure that the SOC is up to date with new emerging threats and newly implemented customer-facing services and products. They are also responsible for keeping their systems up to date when it comes to patches and new errors to make sure their platform is safe.

"[…] The detection engineer for the SOC, that's my main role and my job, it's pretty straightforward to explain. I'm going to make new rules to make sure our detection is good in terms of new vulnerabilities, new products, new services, and stuff." [Interviewee 12, SOC5]

"[…] they need to be patched, there are errors, keep the platform up”. [Interviewee 10, SOC5]

5.2 Mindful Organizing in SOC

Our first research question is, “How does mindful organizing happen in the SOC environment?”. In this section, we will provide evidence from the SOC environments for each of the five capabilities of mindful organizing. These characteristics are*: (a) more time examining failure as a window on the health of the system, (b) more time resisting the urge to simplify assumptions about the world, (c) more time observing operations and their effects, (d) more time developing resilience to manage unexpected events, and (e) more time locating local expertise and creating a climate of deference to those experts”* (Weick & Sutcliffe 2006 p. 516)

5.2.1 Preoccupation with failure (more time examining failure as a window on the health of the system)

For this capability, we have identified how they will test rigorously when creating new detection rules, update their knowledge of the threat landscape, and investigate new or rare alerts.

Interviewee 12 discusses detection rules, and when new detection rules are implemented, they will be tested rigorously to make sure that they work as expected. When a new rule is introduced, it is written, created, tested, and then put into production. Later, they will go back to take samples of potential false positives that have been caught and also perform purple teaming to test and see if new attacks are caught by the rules. This process is continuous and will be repeated during the detection rules' lifetime.

“they're written, they're made, they're tested and then they're put into production, and then we always go back again after a while and verify and test. If there are tests for false positive [we] sample on all the alarms that have gone off on that rule. Or if you spin it up

in a new environment, and then you run a pen-test, purple team, to see if new attacks are taken by the rules you've written and so on. It's a life cycle that keeps going. It's not just such a once-off, then you're done with what you're doing." [Interviewee 12, SOC5]

To detect new and emerging attack methods, analysts and detection engineers need to maintain an understanding of the threat landscape. They will, therefore, try to keep up with news within the cybersecurity domain and use this in their purple team tests, as stated by interviewee 5:

"[...] It's one thing to stay on top of the threat picture, what's creeping and going outside." [Interviewee 5, SOC2]

"Then the employees work in different disciplines like threat intelligence [...]" [Interviewee 5, SOC2]

Interviewee 5 further explains that they also have procedures where the second level reviews more unique or special alerts that have occurred during evenings, night, and day shifts to verify that they have been correctly handled.

"[...] Then we have procedures as well like where the second level analyst looks over some more unique and special alarms that have popped up both in the evening and at night and during the day, simply to just double check that it's correct." [Interviewee 5, SOC2]

5.2.2 Reluctance to simplify interpretations (more time resisting the urge to simplify assumptions about the world)

In the interviews, we looked for evidence that the SOC is avoiding making simple assumptions about what they experience. We also examined how the analysts analyze alerts and what they do to ensure that they are making the right conclusions.

We found that the analysts take precautionary measures to ensure that they make the right conclusion. They are critical and question their own work, have other analysts look over their work, try to interpret the context of the alarm, and escalate the alarm to senior analysts.

The analysts handle large numbers of alarms each day and make decisions based on the information they receive from the alarm. Previous analysis and experiences can influence how the analyst analyzes similar alerts, making it crucial for them to question their analysis. By questioning their work, they can re-evaluate their decisions and find potential areas where they have made assumptions based on prior experience rather than the evidence they have been presented, as explained by interviewee 2:

"I think that it is quite natural that you are colored by different things, whether it is your own experiences or whether it is previous alarms you have examined [...] and therefore I think it is important to be critical of the decisions you make". [Interviewee 2, SOC1]

All the different SOC's also allow the analysts to contact a co-worker or escalate an alert to more senior analysts to retrieve feedback on their analysis. They can make a

new analysis and potentially find more information that was missed in the first iteration. Interviewee 2 explains how it works at their location.

"One sends the alarm to another source to get it analyzed again, and get an enrichment from that person, if they manage to find other information that may be relevant." [Interviewee 2, SOC1]

Interviewee 7 says there can be situations where the analysts are analyzing alarms from systems they are unfamiliar with. This can be an issue when attempting to gather the relevant data from logs to make correct conclusions. To avoid missing critical information, the analyst can contact other analysts or employees with expertise in the relevant system, who can provide additional information or perform a second analysis for that alarm.

"If there is something we are unsure about or something like that, and then we can [...] check with a system managers or department leaders if they are familiar with the system we have investigated. Then they can either give us the name of a person we can cross-check with you, or, yeah, to get a better context for an alarm." [Interviewee 7, SOC3]

To better understand the alarm and why it triggered, the analysts will often read the detection rules. The rules are a set of keywords or artifacts that need to be present for an alarm to be triggered, and this can make it easier for the analyst to understand the context of the alarm, as stated by interviewee 7:

"So, maybe I can see how the alarm is written to see what that rule is looking for to give me a bit more context on, yeah, on what I should look for further in order to then come to a conclusion". [Interviewee 7, SOC3]

Malicious activity can be challenging to detect since attackers are often quite clever. That is why the analyst needs to look outside what's obvious and look for other indicators of an attack. If the normal indicators of a false positive are present, the analysts need to look outside the normal checklist to see if anything is out of the ordinary. The analysts indicate that this is something they strive to achieve, and by doing so, they can avoid missing true positives, as explained by interviewee 8:

"So they are indeed skillful. They try not to be seen, so if you only look for the obvious things, that's what you miss out on. You need to look into detail to see the actual attacks, and that might easily be overlooked when the other indicators are good." [Interviewee 8, SOC4]

Interviewee 6 empathizes with something similar:

"You do that as you don't want to ignore something that may be real, or you think it's a false positive, but then it's actually real. So you always have to keep that in the back of your mind a little bit, and I checked what can be checked [...]." [Interviewee 6, SOC3]

A contradictory point for SOC to be reluctant to simply interpretations is that sometimes the analyst can lean a bit too much on previous analyses when making decisions. In an attempt to be more productive or save their cognitive power, they can use prior analysis of similar alerts as inspiration and potentially forget to look at the alert's context.

“[...] Yes, it can happen, and it sometimes happens that you lean a little too much on previous analyses.“ [Interviewee 4, SOC2]

5.2.3 Sensitivity to Operations (more time observing operations and their effects)

The relevant findings for this characteristic are tuning and suppression of alerts, which can help increase the relevancy of alerts and lower the total workload, and severity classification, which can help analysts focus on the most critical alarms first. Having meetings or discussions with co-workers can help generate a more holistic view of the system.

In SOC, they use the term “tuning” when they alter detection rules for alarms that have been triggered unnecessarily. Often, this involves the analyst marking an alarm for tuning and describing why the alarm needs tuning, and the alarm is escalated to other employees in the SOC who are trained in detection engineering. The detection engineer will then alter the rule corresponding to the analysts' description. Interviewee 5 notes that it happens continuously:

"Tuning is something that really happens, it's a continuous process that happens all the time." [Interviewee 5, SOC2]

In addition to tuning, some of the SOCs have the option to suppress alerts. When they suppress the alert, this particular alert will no longer appear for the analyst. This is used when an alarm can be noisy and disturb the analyst and is often suppressed in combination with tuning. When the analyst is waiting for the alert to be tuned, they can suppress it in the meantime. Interviewee 2 discusses that they have this capability.

"[...] we have the opportunity to change in real-time whether a certain type of alarm comes in or not." [Interviewee 2, SOC1]

We found that all the SOCs are working with classifications of alerts to indicate the severity of an alarm. The classification allows the analyst to prioritize alarms more efficiently and will help mitigate the chance of severe incidents going unnoticed. They will often be classified from low to critical and have corresponding colors like green and red. Interviewee 9 talks about their different severity categories:

"[...] in our system, we have Low to Critical. Low, Medium, High, and Critical already as Severity when you look at the case." [Interviewee 9, SOC4]

Some of the SOCS are offered meetings with other parts of the organization so the analysts can better understand the infrastructure or systems to help their professional growth. This can make it easier for them to understand alarms that come from systems they are unfamiliar with. This is especially noted by interviewee 7:

"[...] they, for example, offer to run a meeting or just a little bit of information on, like, how they work and what they have or are sitting on of infrastructure or systems, so it's very open like that with sharing of knowledge for you to get better." [Interviewee 7, SOC3]

5.2.4 Commitment to resilience (more time developing resilience to manage unexpected events)

It is commonly understood that mistakes will happen, and certain SOCs have strategies in place to keep operations running in the event of unforeseen circumstances. Some even have employees who are always available to help in case of an emergency. SOCs often operate with a service level agreement (SLA) that requires them to have next to zero downtime. This is due to the critical consequences their downtime might result in. Therefore, they strive towards minimal downtime and the ability to get to a normal state as quickly as possible. They will often have employees available 24/7 to assist in these situations. This is explained by both Interviewee 5 and Interviewee 11.

"[...] attempt to restore the normal state as quickly as possible. We work according to the kind of Service level agreement that so many other operations centers do, so it is quite time-critical that things work as they should and as intended." [Interviewee 5, SOC2]

"There are people who are responsible for it, they are contacted 24/7 because we have running surveillance every single second of the day and year. If a system goes down it has to be up "yesterday" again.". [Interviewee 11, SOC5]

As some of the SOCs have distributed offices, interviewee 12 states they have redundancy since they can provide SOC for their customers even though one of their locations is made unavailable. The analysts in the other offices will then receive the alarms that were supposed to go to the office with downtime, and the customers will still be given the service they are promised.

"If one of the offices goes down or the internet in the region is down, then we still have people who can deliver the SOC. If the infrastructure was a part of the SIEM tools that we host for the customer goes down, then we have redundancy on that." [Interviewee 12, SOC5]

As noted by interviewee 2, their tools often have backup solutions. In case of emergency, they will use these tools to remain operative, even though they might not be optimal for the situation.

"It has happened in the past that tools have gone down, and we had to do without them for some time, we had to use alternative solutions, that may not work nearly as well, but become a crisis tool in that setting you are in." [Interviewee 2, SOC1]

5.2.5 Under-specification of structures (more time locating local expertise and creating a climate of deference to those experts)

SOCs have varying degrees of expertise that can be contacted when faced with scenarios where this particular knowledge is needed to avoid failures, errors, or making uninformed decisions. From most of the comments from the interviewees, there seems to be an industry practice where they divide the SOC into levels of expertise or experience. That includes experience from other parts of the incident response team and seniority. Often, they have three different levels of employees. If particular expertise

is needed to make decisions, then they can easily get in touch with them without going through a lot of different measures, as explained by interviewee 4.

“We have our first line which is the ones who sit and do the initial analysis, and then we have the second line which is the ones we use for analysis assistance and they sit right behind us, So you can turn around and ask for help, or you can escalate [to] them with a bit of a workflow like that, as well as third level.” [Interviewee 4, SOC2]

Interviewee 1 stated that they follow the industry practices.

"My company follows industry practices, and with that, it will typically consist of a first, second, and third level. And, where it escalates and triages after that practice." [Interviewee 1, SOC1]

We also found in one of the SOCs that when any of them need backup they can contact a group chat with others who can jump in and assist. Even though it might entail overtime for other employees, it can be done without having to go through management for confirmation. This will also result in them being able to contact the necessary expertise when they need it, this is explained by interviewee 4.

"yes, and everybody has the opportunity to do that's what the very first level can say we need a backup and even those who work in the evening, can then take the opportunity to send a message to sort of like a group with the shift pool without having to ask the boss's permission first, even though it entails extra compensation." [Interviewee 4, SOC2]

5.3 Mindful conduction and alert fatigue

In the previous section, we saw that the analysts are attempting to act mindfully, as this is a crucial part of their job. There is clear evidence that shows how mindful organizing happens in SOC environments. This section will answer the second research question: “How does alert fatigue influence mindful conduction in SOC?”. To be able to do so, we will put forward the findings that are coherent with the symptoms of alert fatigue and how these cause the analyst to act mindlessly instead of mindfully. We see clear signs of how alert fatigue has a negative effect on mindful conduction by how alert fatigue affects the individual. Since alert fatigue consists of both desensitization to alerts and cognitive overload, we will present findings related to these conditions separately.

The first section will explain how desensitization will lead the analysts to make decisions based on previous experiences, make hasty decisions, and act more automatically. These effects are counterparts to two of the mindful conduction facets: *activity awareness* and *output assessment*.

The second section explains how cognitive overload will affect mindful conduction by overwhelming the analyst, so they no longer are able to perform proper analysis and are unable to grasp the context of the alerts. This will cause them to be less aware of their activities and follow predefined rules. These factors are also counterparts of the two facets mentioned above: *activity awareness,* and *output assessments*.

5.3.1 Desensitization

The analysts are exposed to repeating alarms daily. They are often the same or very similar, making it hard to distinguish one from the other. As they see these alerts multiple times during their shift, they are prone to becoming desensitized to them, which will cause them to draw quicker and less fact-based conclusions. Interviewee 5 states that they will often lean on previous experiences, like their own or other analysts' analysis, to make conclusions instead of analyzing the alert themselves.

"[...] if you see the same thing over and over again. [...] It has happened that you make decisions based on past experiences without necessarily checking everything". [Interviewee 5, SOC2]

As interviewee 7 says, this can lead them to miss important clues about the alarm, which can result in true positives being missed.

“It could make you make a hasty decision and think that if you have seen many false positives, then in the same category or similar, it could make you miss something that is not” [Interviewee 7, SOC3]

These repeating alerts can lead the analyst to begin reacting more automatically, leaving their mindful behavior behind and no longer being aware of the alert's context. One of the analysts referred to this state as “zombie-mode”, where they will analyze the alerts that they are familiar with based on what they did previously. New alerts that they have not seen before, however, will not be affected by this automation and will be adequately analyzed. Interviewee 9 explained it like this:

"It's kind of like zombie mode and it's just sitting and especially on the stuff you recognize that you've seen before, it's really automatic. The new ones, then it kind of falls out of that automatic, and it's pretty weird, so it's 100% automatic". [Interviewee 9, SOC4]

Interviewee 2 explains that as they receive more false positives at a high rate, they will become tired. Not just because of the high rate, but also because they are unable to conduct a proper analysis of the alert. So, when they receive false positives, they will conduct poor analyses and be drained of energy as they do it. They describe the situation as a “whack-a-mole” game, where they will get one alert out of the way, and another will appear.

“[...] I became very tired of continuously just staying on the surface and not being able to use my capacity to analyze properly; it just ends up as a kind of ”whack-a-mole” atmosphere [...]” [Interviewee 2, SOC1]

5.3.2 Cognitive overload

Cognitive overload has been found to be a problem in four out of the five SOC environments. Most analysts can refer to situations where they have been overwhelmed by the number of alerts, and this has reduced the quality of their analysis. This, too, can lead the analysts to make quicker decisions without grasping the whole context of the alerts before submitting their analysis.

Interviewee 9 had an experience where they had so many alerts that they were unable to take a break and ended up not analyzing an alert thoroughly enough to make a proper decision. Instead of reporting it, they decided to leave it be.

"I know I don't have time to take a break. I don't have time to eat lunch, then I sit there and know that, with low blood sugar and quite tired and foggy, that now, now I can make mistakes because now I am quite tired, so my judgment, and it goes down, I can, I can honestly say that I have looked at an alarm that is in the gray zone, should I report, should I not, I have been so tired that I can't be bothered". [Interviewee 9, SOC4]

Another analyst, interviewee 11, reported that they had to send an alarm to another analyst at the end of their day because they were no longer able to perform a proper analysis of that alarm. This was due to a shift with a high number of alerts, which caused the analyst to be deprived of their mental capacity.

"*[...] you don't have the mental strength to go through it, but it's kind of like that then I often hand it over to another analyst or just 'can you take this one for me I don't have the willpower to go through it'". [Interviewee 11, SOC 5]*

As the number of alerts increases, the attention of the analyst will be extended over more alerts. The analyst's ability to investigate the alert properly will decrease and they will struggle to understand its context, as explained by interviewee 2.

"You get spread out thinly when the alarm pressure is so high [...] So, it is as you say, I think you lose the ability to establish that context on a deeper level if it is not very clear". [Interviewee 2, SOC1]

Some of the analysts mention that they use a mental checklist to analyze the alerts. This can save their cognitive power and make their analysis easier, but they must keep the context in mind when they analyze them. Interviewee 11 elaborates on this problem, saying that there have been incidents where an analyst has relied too much on the checklist and missed out on relevant information.

"*I have seen it before, but it has happened that people have missed stuff because they are using a checklist on these use cases and then missed stuff*". [Interviewee 11, SOC5]

5.4 Mitigations

This section will answer the third research question: "What strategies are being implemented to mitigate alert fatigue". There are a fair number of mitigations that emerged from the interviews. The literature chapter divided the mitigations into technical and human factor mitigations, but since we only found one technical mitigation, we have instead divided them into three other categories. These are: mitigating of causes, mitigations of desensitization and cognitive overload, and mitigating of consequences. At the end of this section, we will also elaborate on the use of AI and ML as a mitigation and why they are rarely used in SOC.

5.4.1 Mitigating of causes

The leading cause of alert fatigue is the high number of false positives. To mitigate this, the detection engineers attempt to only provide the analysts with the information they need and create rules that do not result in false positives. This can help to reduce the chance of cognitive overload and desensitization.

Interviewee 12 explains that alert fatigue is due to poorly defined detection rules that do not provide the analysts with the necessary information and have inadequate automation. Automation, in this sense, could mean grouping similar alerts and removing irrelevant information.

"[...] one of the most important points about alert fatigue is that alert fatigue actually comes from bad rules or bad automation." [Interviewee 12, SOC 5]

To avoid overwhelming the analysts with information, the detection engineers try to limit the information they include with the alerts. Irrelevant information should be excluded as it has nothing to do with the current incident, as explained by interviewee 3:

"We're trying as best as possible not to present things that we don't think are important for an analyst to see." [Interviewee 3, SOC1]

Interviewee 12, who is a detection engineer, explains his goal is to create the best possible rules and make it easier for the analyst to decide if the incident is a true positive or not.

"My goal is to create rules that are more concentrated than the ones we already have so that when an alarm comes, it should be well described with exactly what happens. Since I have very concentrated rules which means that there have to be a lot more rules that put more of the pressure on knowing scary things in those environments on me and my department then, so that responsibility is shifted from the analyst 100% and more onto that the detection should see if there is anything scary and let you know if there is anything, and [...] then you should be able to make a good choice or easily see where can I find information about if this is dangerous or a false positive?" [Interviewee 12, SOC5]

5.4.2 Mitigate desensitization and cognitive overload

Some of the SOC environments can suppress repeating alerts to avoid unnecessary noise. They will place a filter that excludes specific alerts for a period, which allows them to focus solely on new and more severe alerts. This can be helpful to minimize repeating alerts and thus mitigate desensitization. This is explained by Interviewee 5:

"The first line has the opportunity to, [...] put a filter in place to, yeah, simply suppress and escalate down." [Interviewee 5, SOC2]

In a situation with high pressure that demands more attention than one analyst can give, interviewee 3 explains that they have the option to call in extra staff. This way they can divide the demand and make sure that all the alerts are analyzed properly and avoid cognitive overload.

"[...] So when it gets too much, you can turn to other parts of the organization and then hopefully bring in enough personnel to handle the amount that comes in." [Interviewee 3, SOC1]

Similarly, there is the option to escalate the alerts, which will take more time to analyze, to a second or third-level analyst. If these analysts are available, they can relieve the first-level analyst from some of the workload. One of the analysts explains:

"[...] If there's big, big pressure, a lot going on then and there and you get an incident that you see that OK this one is going to take some time to analyze. [...] Then you also have the opportunity to kind of lift it up to someone else.". [Interviewee 5, SOC2]

Interviewee 10 talks about how sometimes a message or post in a group chat is supposed to raise awareness. It mentions they should remember to focus on one alarm at a time and avoid multitasking to make sure that they are analyzing the alarm correctly. They emphasize that they should focus on thought analysis rather than being fast and getting the alerts out of their way.

"[...] A post on Teams now and then, which we say makes you aware of doing one analysis at a time, one customer at a time, not multitasking and making sure it's doing it right" [Interviewee 10, SOC5]

Both interviewees 11 and 12, are working on other projects to get a break from analyzing alerts. They agree that this is a helpful measure to stay focused on the job. Interviewee 11 explains:

"[...] you vary on projects, somehow just looking at a different user interface and not looking at the same logs all the time and staring at payloads from a log source." [Interviewee 11, SOC5]

Interviewee 5 has a similar opinion. They say that it is nice to have a balance between challenging tasks and tasks that are less demanding.

"[...] I can sit here and talk to you about an interview instead of having to sit and deal with security incidents which are essentially more chill, but not as academically challenging and educational I would say, but you need a nice balance between challenges and a little bit more laid back tasks." [Interviewee 5, SOC1]

Interviewee 11 will mute notifications on Teams to avoid distraction. This way they will have their full attention on the alerts, and nothing else.

"[...], so I'm busy on Teams so I don't get notifications on Teams, yes, it's full concentration." [Interviewee 11, SOC5]

Another interesting finding was that the analyst will work by established routines and checklists to save their cognitive power, and not become tired as fast. We would categorize this as a mindless way of working, but as interviewee 2 says, they will also be aware of the alert's context. This way, they avoid alert fatigue by performing seemingly mindless actions.

"Then I try to catch myself in it and do that extra check then, but I would say that the established routines that you end up working out. Hopefully, they will be so good that

they might roll in the extra check, so that you then have the opportunity to OK, so even if you're tired and you're fed up and you get the same alarm for the seventh time, it's OK." [Interviewee 2, SOC1]

Interviewee 11 explained that they will take some time off to work on other projects or to relax at home by playing video games, watching TV, or reading.

"Taking some time off and using some of my days off, and then sitting and playing and relaxing for one or 2 days or watching shows." [Interviewee 11, SOC5]

5.4.3 Mitigate of consequences

After alert fatigue has manifested itself, some mitigations can still be effective. Most of the analysts say that they will take breaks and step away from their desks. Having a break will refresh their mind and help mitigate the feeling of overwhelmedness when there are a lot of alerts. Interviewee 12 says:

"I'm trying to take, get some fresh air, get some breaks, good food." [Interviewee 12, SOC5]

Interviewee 8 explained that they would work at their own pace and ignore the fact that there are a lot of alerts. They will focus on doing a good analysis and work by criticality to have a structured workflow.

"There's nothing I can do about the amount coming in so just take the break and then if that I, if there's 100 alarms there and I'm hungry then I'll take lunch regardless of there being stuff there. Also, I just work by the criticality." [Interviewee 8, SOC4]

Interviewee 2 has a similar statement.

"[...] Usually, I try to neglect the fact that there are a lot of alarms coming in, that it's going to affect the capacity I have to analyze correctly." [Interviewee 2, SOC1]

5.4.4 AI and ML

There is little use of Artificial intelligence and machine learning in the SOC environments we have investigated. There is a healthy amount of skepticism to both, even from the ones who use it in their SOC.

Interviewee 12 talks about how their organization has anomalous behavior detection, which is a type of ML that reacts when something is out of the ordinary and is seen as anomalous.

"We've got some machine learning on now the kind of user behavior is part of that anomaly stuff, right? So, we have. We have the technology running the normalized state and so any anomalies that go outside of it can just classify as an alarm." [Interviewee 12, SOC5]

However, they are skeptical about leaving everything to machine learning and AI. They do not believe that these things are mature enough to be fully trustable. They explain further:

"[...] Machine learning and AI, those things. I don't know if we're mature enough to just put everything over to them and hope for the best. It will probably take a while before I at least trust it, but I'm kind of a skeptical person [...]." [Interviewee 12, SOC5]

Both interviewees 2 and 3 state that no AI is currently being used in their SOC. Interviewee 4 adds that they are not convinced that AI is the solution, even though they have some AI and ML in their SOC. They add that the AI term is broadly used, making it hard to say what technology is AI. They even mention that it could generate more alerts and be counterproductive for their SOC.

"I am very unsure about the AI solution. I'm not entirely convinced. I think often AI can create even more alarms." [Interviewee 4, SOC2]

"So AI is a very broad term. There are things that can be defined as AI, but pure machine learning doesn't have much of it yet. [...] It's not reliable enough for us yet, in a way." [Interviewee 4, SOC2]

Interviewee 5 agrees with Interviewee 4, saying that it is a buzzword used by their suppliers and that they do not trust it fully yet.

"[...] Anomalous behavior detection is not exactly AI or machine learning but can be sort of a form of it then because it is you go beyond normal behavior, so some suppliers will throw that mark on this type of detection mechanism as well. We have something to some extent. And for my opinion and from my point of view, it's not working optimally yet at least" [Interviewee 5, SOC2]

Interviewee 5 also makes an interesting statement: AI will probably be helpful to attackers before it will be helpful to defenders.

"[...] I'll probably see AI be useful to attackers before it will be useful to us. I think they can take advantage of it before the defense can do that." [Interviewee 5, SOC2]

5.5 Unexpected findings

Prior to writing this thesis, we assumed that alert fatigue would be a problem for all the SOCs. Interestingly, the only SOC not experiencing alert fatigue was the internal SOC. The rest of the external SOCs were experiencing it to some extent. Interviewee 6, who has experience only from their current employer, says they have not felt overwhelmed.

"I haven't really felt overwhelmed, I think." [Interviewee 6, SOC3]

The other analysts from the same company had felt overwhelmed before. However, this was at their old employer. At their current employer, they say that they have not felt overwhelmed, as interviewee 7 states:

"How often have you felt overwhelmed, and what was the main reason?" [Researchers]

"No, I haven't really felt it" [interviewee 7, SOC3]

These analysts operate in an internal SOC, meaning their primary responsibility is defending the company they work for. This is the primary distinction between their SOC and the others that have to protect external companies. It is interesting to see that there

are no major differences other than this SOC being an internal one and that there still is such a large difference in their experience of alert fatigue.

6 Discussion

In the discussion section of this thesis, we will discuss the theoretical implications of the research performed, including the practical implications, limitations, and recommendations for future research.

6.1 Theoretical implications

This part of the thesis will discuss the empirical data and the previous literature relating the data to previously known concepts and theories. Some of the main theoretical implications are mindful balancing, alert fatigue and the entrenchment problem, and the state of AI and ML.

6.1.1 Mindful balancing

We consider mindfulness and mindful conduction to be intertwined; they both represent the same thing. The mindfulness theory by Levintha & Rerup (2006) defines it as being in active awareness, trying to understand the context and view the situation from multiple angles, whereas Rina-kahli et al. (2023) divide the concept into three different facets: “active awareness”, “competence maintenance” and “output assessment”. These definitions refer to the same goal, which is to be able to grasp the whole picture of the situation an individual is in. We will refer to this as being mindful and mindfulness.

As an opposite of mindfulness, we have mindlessness, which in Levintha & Rerup (2006)s theory, is when individuals act more on autopilot and lose the context they are working in, including only seeing a situation from one angle. According to Ashforth & Fried, 1988 (as cited by Ault & Brandley, 2023), what we describe as mindlessness is an individual’s *default state*, where they will perform mental routines at a superficial level without contextual awareness. To exit the default state, they must use cognitive effort to act mindfully, which Ault & Brandley (2023) describe as an exhaustible resource that causes negative effects when exhausted. As we see it, this negative effect is to go back to the default state of mindlessness. Interestingly, this notion of exhaustible resources is also found in the description of general fatigue by Aaronson et al. (1999), which states fatigue is the imbalance in the restoration or utilization of resources needed to perform an activity. This could mean that the mentioned resource is necessary to perform the activity of being mindful, and if this imbalance occurs, the individual becomes fatigued and therefore falls into the actions of mindlessness.

Our findings show that the analysts attempt, to the best of their ability, to remain mindful by using their cognitive power when investigating alarms and attempting to discover anything malicious. We have also found that some analysts begin to work more automatically when faced with large numbers of alerts. They will try to get rid of the alarms faster by conducting quicker and shallower analyses, thus acting more mindlessly. Regardless of the number of alerts, the analyst will often follow a predefined

mental checklist to determine whether the alert is a true positive. Using this list, they can check a range of common indicators that often would make them able to draw a conclusion. This action alone would also be categorized as a mindless action as it does not take the context of the alert into account; however, the analyst says that they will have a last check for other malicious indicators at the end of their analysis, which we see as a mindful action as they are investigating the context of the alerts and assess their initial analysis.

Mindfulness is a fundamental part of the cognitive mindset of an HRO (Salovaara et al., 2019). However, as we have seen, they are not operating mindfully 100% of the time. What seems to be occurring is that the analysts try to remain mindful, but to maintain their cognitive power they are sometimes acting mindlessly. By, for example, following this routine with the checklist, they are saving their cognitive power to preserve their mindfulness. In other words, to avoid their cognitive power being exhausted, they are preserving it by balancing their mindlessness and mindfulness and thus avoiding alert fatigue. It is also interesting to see that the analysts not only will take breaks during their shifts, but also in between shifts. They will have some days off work or perform other tasks at their SOC as a way of getting away from the alerts. They will have these small and long breaks to replenish their cognitive power so they will be able to be mindful.

Based on our readings, we could not find any prior research on how the delicate balance between mindfulness and mindlessness is handled in SOC environments. Therefore, we will propose a novel theory of "*Mindful Balancing*", where we propose that the security analysts working at SOCs balance their mindfulness and mindlessness to avoid cognitive exhaustion. There is a sweet spot between the two poles where the analysts are in balance with mindlessness and mindfulness, where they are able to conduct proper analysis while preserving their cognitive power. On the opposing end, they will either use too little of their cognitive power and act on autopilot, while on the other side, they will use too much of their cognitive power and become cognitively exhausted.

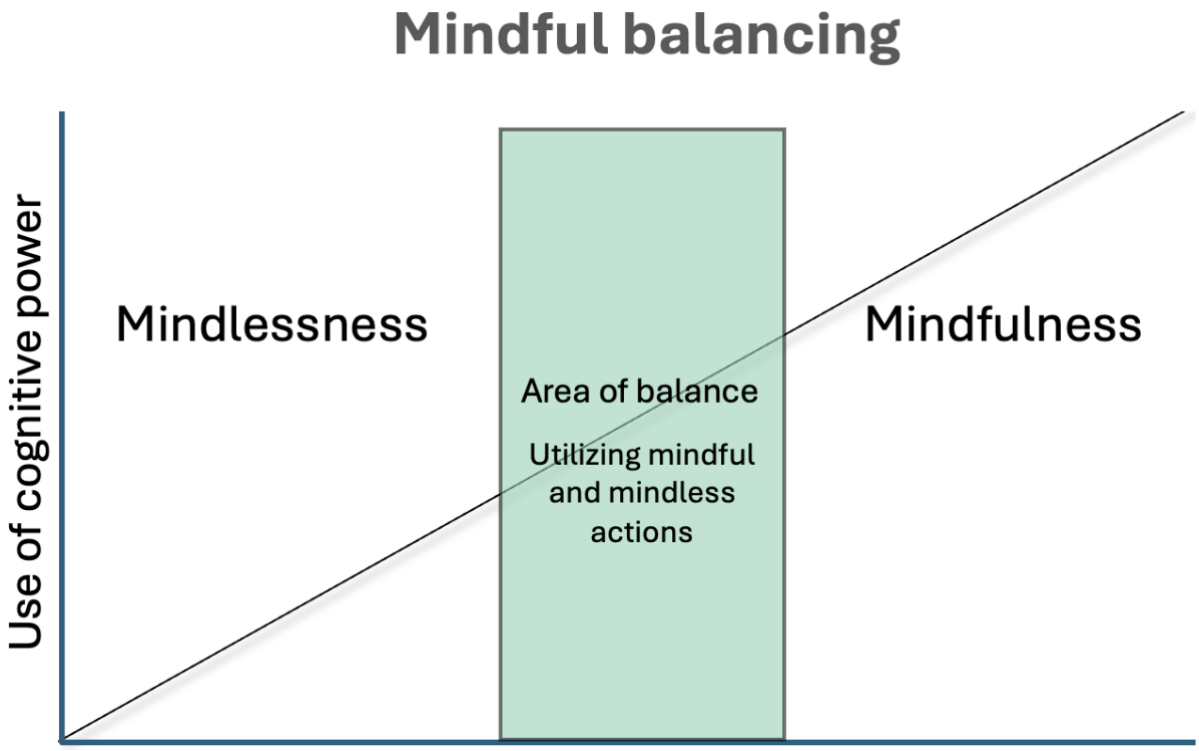

Figure 6: Mindful balancing

Figure 6 depicts the Mindful Balancing. We see on the left that the analysts will preserve too much of their cognitive power and act mindless, thus performing shallow and improper analysis. On the right side, we see that the analyst will use too much of their cognitive power and might become cognitively exhausted. In the middle, we find the sweet spot, where the analysts are in balance with their mindfulness and mindlessness. They are able to save cognitive power by mindlessly performing tasks while staying sufficiently mindful by being aware of the alarm's context. Another way they balance their cognitive power is by taking short or long breaks from analyzing alerts. We suggest that their mindfulness can be replenished to some degree by taking short breaks and, to a larger degree, by taking longer breaks.

This theory is a novel theory that we hope will aid other researchers in investigating mindfulness in SOC or similar environments and to find how individuals in these environments perform mindful balancing in their environments.

6.1.2 The alert fatigue and entrenchment problem

The entrenchment problem is described by Salovaara et al. (2019) as a condition that might affect individuals who are performing mindful operations. They are prone to become too reliant on technical solutions, which leads them to, among other things, fail to anticipate unexpected events, fixate on a single perspective, and dismiss new and conflicting information. These effects suggest that the individual will have an increased chance of acting mindlessly instead of mindfully (Salovaara et al., 2019). These characteristics are similar to the effects of alert fatigue found in our cases. The analysts

will be unable to see the context of the alerts, dismiss information outside their checklist, and sometimes fixate on getting alerts away instead of analyzing them properly. As analysts feel the effects of cognitive overload and get desensitized to alerts, they will lose their ability to act mindful and start to act mindless.

Although the concepts are similar, there are some distinctions between them. The entrenchment problem involves being trapped in group thinking and growing overconfident, which has not been found to be the case for alert fatigue. More importantly, while alert fatigue is caused by a vast number of false positive alerts, the entrenchment problem involves reliance on technology and trust in that the technology will make the right decision (Salovaara et al., 2019). The Venn diagram in Figure 7 visualizes the similarities between the two concepts.

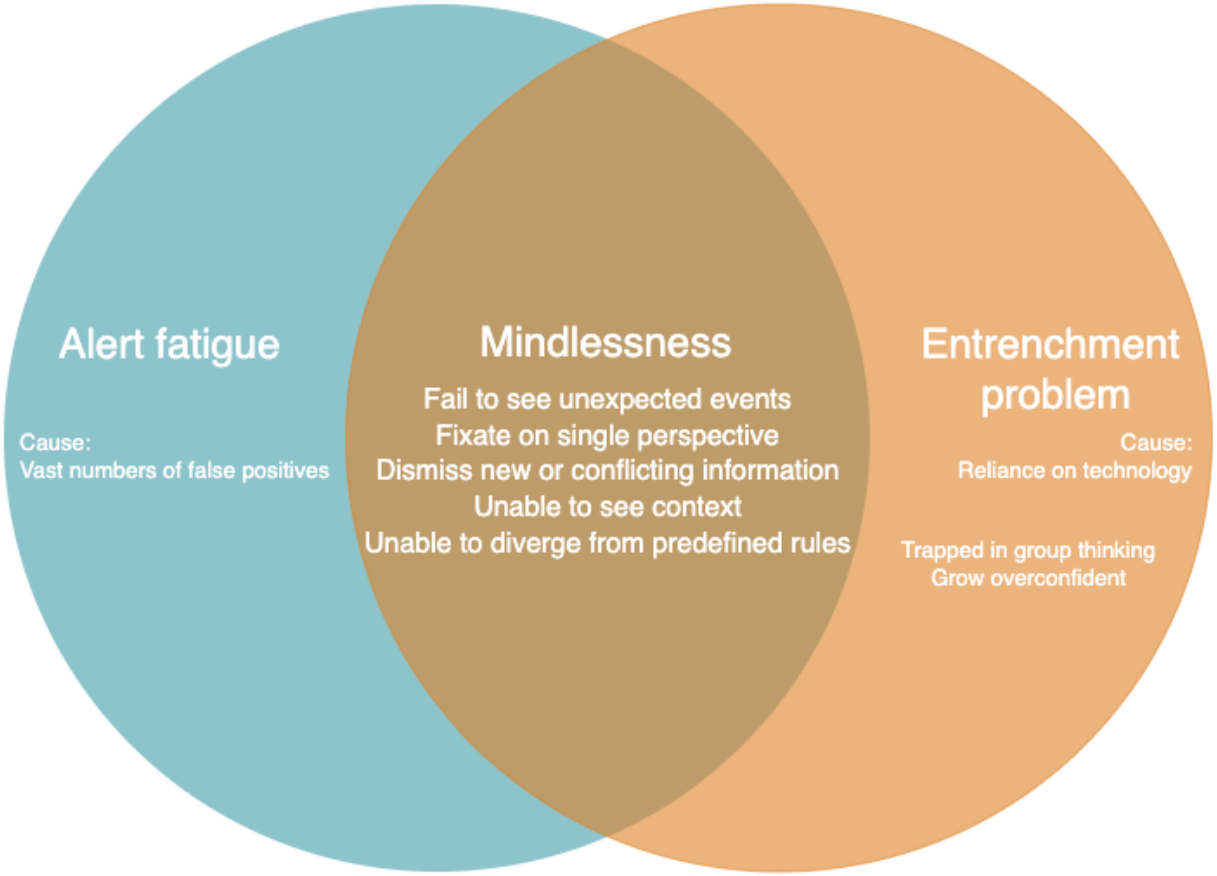

Figure 7: Alert fatigue and Entrenchment problem

Although these two concepts are similar, we do not find any strong indicators of the entrenchment problem in our cases. We have investigated if the analyst would make decisions based on what their information system suggests, which would indicate that they will rely too much on their technology. Their system only suggests the criticality of the alerts and not actions that should be performed based on it. We did not find any cases where the analysts made decisions based on criticality, and therefore no signs of the entrenchment problem. However, they will use the criticality as an indicator of what alert to analyze first when they have multiple alerts in front of them. If the analysts were

to trust that criticality too much, then it could be an indicator of the entrenchment problem.

It could be interesting to investigate the entrenchment problem in SOC further. Since they both cause analysts to act mindlessly, we hypothesize that there could be cases where alert fatigue gets the blame for something caused by the entrenchment problem, and vice versa. The Venn diagram above can aid in finding the cause of the mindlessness and thus help to identify the right mitigations for the problem.

6.1.3 AI or ML

What's interesting to see is that the ML and AI mitigations found on the technical side of the literature are absent when looking at the empirical findings. This is very interesting since most research on the mitigation of alert fatigue in a SOC environment bases itself on this type of technology. One potential reason for this could be the fear of skill erosion, which usually occurs when someone becomes too reliant on tools and systems under the umbrella of cognitive automation. The article by Rinta-Kahila et al. (2023) discusses this phenomenon. When individuals become too reliant on an intelligent system, their skills and expertise will slowly decline with a called "degeneration effect". It could potentially end up in a situation called "complacency". This is a situation where an expert will consider all is well because they have blind trust in the system and have no awareness of the actual circumstances (Rinta-Kahila et al., 2023). This fear of skill erosion or complacency seems to be absent from the findings; rather, there exists a lack of trust in these tools, and currently, they create more work for the analysts.

Interestingly, this supports an argument made by Alahmadi et al. (2022) states that AI is based on statistics and probability. Therefore, analysts should be cautious when putting their trust in these systems completely and additionally, discussing that AI models currently are usually unintelligible to personnel who hold no expertise in it. They state in the findings that the lack of trust in these systems is their inability to correctly detect what should be detected and it generates more false positives.

6.2 Practical Implications

In this part of the thesis, we will discuss the practical implications that this study provides for SOC environments. These practical implications are split into two categories: assisting mindful balancing and facilitation for mindful balancing. These insights are actionable approaches for the SOCs to use when wanting to prevent or reduce alert fatigue. We have identified some practical implications that may assist the SOCs in their venture, such as dividing responsibilities among analysts, managing detection rules and suppression, giving analysts breaks, and strengthening their attitude.

6.2.1 Assisting mindful balancing

These mitigations are seen as strategies that directly assist the analysts in maintaining and utilizing mindful balancing when working. This is either by giving them additional responsibilities to perform or by allowing them to take breaks to recover their mindfulness.

Work rotation

In our findings, we found that some SOCs allow their analysts to have more responsibilities besides only being analysts who analyze alerts. They let them work on various projects that let them divide their time between two or more roles at their SOC. This allows the analysts to temporarily take a break from the environment where alerts appear and reduce their exposure to alerts, which helps to prevent desensitization. As discussed in our Mindful Balancing theory, this break will help the analyst replenish their cognitive power. We, therefore, suggest introducing a work rotation in the SOC where the analyst can get time away from the alerts, which will assist them when it comes to Mindful balancing.

Breaks

We have also found that breaks are frequently used. This allowed the analysts to regain some of the cognitive power to continue working after they had begun feeling alert fatigued. It lets them remove themselves from the environment with all the alerts, which lessens desensitization. Not having to utilize their cognitive power during breaks gives the analysts time to replenish their mindfulness so they can maintain balance. We can, therefore, suggest giving the analysts the ability to take 5-minute or longer breaks occasionally without compromising the SOC operability.

6.2.2 Facilitation of mindful balancing

These are mitigations that facilitate mindful balancing. This can be done by either reducing the frequency of alerts in the short or long term or by adopting a mindset that will fortify them mentally.

Detection and suppression

As discussed earlier, mindfulness is seen as a resource that can be exhausted, and when this occurs, the individual becomes cognitively overloaded. Thus, to prevent this, one can lower the amount of work that needs to be performed. Additionally, reducing the number of false positives will help prevent analysts from becoming desensitized to the alerts. We, therefore, suggest two different mitigation strategies that aim at this specifically. In the findings, the detection engineers mention that inadequate detection rules are the leading cause of alert fatigue since, when done poorly, they will generate a high number of false positives. Our first suggestion is to improve the detection rules.

This will lower the number of false positives which helps the analysts save their mindful resources by lowering the amount of work and lowering the number of false positives which helps to prevent desensitization. Additionally, as a bonus, it could also improve the detection capabilities of the SOC, thus better safekeeping of the system it monitors.

This can be done by investing more in detection engineering, either through more personnel or increasing their proficiency in detection engineering. Additionally, implement testing and periodical samples of rules to see if they are generating a lot of noise and that they detect correctly.

Suppression is the other solution. This allows an analyst to block an alert from reappearing in their view. If the detection rule is triggered, it will temporarily stop making new alerts. This gives the same benefits as improving detection rules, like lowering the workload for the analysts and reducing the number of false positives. Additionally, it can be implemented much faster, thus giving more of an immediate result. However, it's only a temporary solution; it does not improve the detection of the SOC, and If performed by inexperienced analysts, it could potentially blind the SOC. It should, therefore, be used with care and needs authorization from more experienced analysts to be activated.

Attitude

Attitude has also been seen as a way of coping with the number of alerts they have to work with. They do this to lessen the stress and not focus on how much time they spend on each alert so they can perform an adequate analysis. This is done by only focusing on one alert at a time and avoiding trying to multitask on many different alerts, spreading themselves out. We suggest implementing weekly meetings about what to focus on and what attitude they should have while working. This can also be done by sending messages in a team chat or discussing with each other one-on-one. It allows them to potentially maintain better focus in stressful situations.

6.3 Limitations and Further Research

It is essential to acknowledge that every study will have limitations, and thus, it is vital to recognize the limitations that this study has. Time is an aspect that will influence the study as it must be performed in one semester which is around five to six months. This affected the depth of the thesis and the amount of data gathered. The potential population for the research was also limited due to the lack of companies that provide SOC services. This made it difficult to obtain a sample size large enough to represent the whole community.

Another limitation of the study is that the information from the interviews could negatively reflect on both the company and, in some ways, the interviewee as well. This made some of the data gathering more difficult. Some of the interviewees could potentially have withheld important data, as the sensitivity of the subject could lead to participants being unwilling or uncomfortable to share valuable information.

Some of the interviewees also had limited knowledge when it came to a deeper understanding of what was occurring on the backend. This could potentially have led us to miss out on more information on both AI and ML, including other mitigations that could have been used by the organization. The mitigation findings in our research have mainly been based on information collected from security analysts who mainly work directly with the alerts and don't have much insight into all the changes happening behind the scenes. The last limitation is the level of confidentiality promised to the companies. Our research approach is multiple case studies, and one should provide enough information to discern the different cases. However, we promised to provide as little information about the companies as possible to prevent anyone from discovering their identities, thus making our research approach somewhat limited.

Further research should look into interviewing other roles in the SOC; they will most likely have different areas of expertise and could provide better insight into other mitigation strategies that might be in use. One interesting finding in this study was that external SOCs felt alert fatigue, and the internal SOCs did not. We hypothesize that an external SOC is bound to focus on many different systems putting more pressure on the analysts to have a better understanding of multiple systems than an internal SOC that primarily works with the same or similar systems. Examining this could be interesting and further down the line provide insight into the potential benefits of having an internal rather than investing in an external SOC. Further research on the mindful balancing theory should also be conducted to investigate whether these actions occur in all SOCs and whether other HROs also utilize different ways to balance both mindfulness and mindlessness.

7 Conclusion

This thesis sought to answer three research questions, "*How does mindful organizing happen in the SOC environment?*", "*How does alert fatigue influence mindful conduction in SOC?* ", and "*What strategies are being implemented to mitigate alert fatigue?*". To be able to answer these questions we utilized HRO as our theoretical framework, reviewed previous literature with a systematic literature review, and analyzed our empirical data. Our findings suggest that mindful organizing occurs in the SOC and that they perform actions that fall within all the dimensions of the five characteristics of mindful organizing. This includes, among other things, testing rules before they are implemented, having people on standby 24/7, and being ready for unexpected events. Additionally, our findings identify that alert fatigue caused by a large number of false positives influences mindful conduction by causing a shift towards mindless behavior. Many organizations utilize different mitigation strategies, like detection rules or suppression in order to lower the number of false positives, however, there was little variation between the SOC's mitigation strategies.

One of the key findings in our research is that SOC analysts attempt to balance their mindfulness and mindlessness, for which we have proposed the novel theory of Mindful Balancing. This theory proposes that the analysts in SOC environments are saving their cognitive power, which is an exhaustible resource, by performing mindless and mindful actions. This has been identified by their mindless action by following a mental checklist and then performing a mindful action by verifying if there is anything else that might indicate an attack. A model has also been developed to help visualize the theory.

Another interesting finding is that we identified some similarities between the entrenchment problem and alert fatigue and that they differ in their causes. It can be interesting to investigate if one is being confused with the other in other SOCs. However, the entrenchment problem was not identified as an issue in our cases. Lastly, we have discovered how the use of AI or ML as mitigation strategies is almost nonexistent due to mistrust and the fear of it becoming a burden since it can generate more false positives.

We also identified various mitigation strategies, which we divided into two categories regarding mindful balancing: assisting mindful balancing and facilitating mindful balancing. These mitigations can be implemented to aid the analysts in mindful balancing and create an environment where it is less challenging to perform.

This thesis contributes to the literature on mindfulness and mindlessness through the new theory of mindful balancing, including creating a bridge between these concepts and alert fatigue. We hope that this research will provide a better understanding of how SOCs as HROs attempt to maintain their mindfulness when faced with a large number of alerts and that further research can be conducted on this theory.

Bibliography

Aaronson, L. S., Teel, C. S., Cassmeyer, V., Neuberger, G. B., Pallikkathayil, L., Pierce, J., Press, A. N., Williams, P. D. & Wingate, A. (1999). Defining and Measuring Fatigue. *Image: The Journal of Nursing Scholarship*, *31*(1), 45–50. https://doi.org/10.1111/J.1547-5069.1999.TB00420.X

About Google Scholar. (n.d.). Retrieved February 4, 2024, from https://scholar.google.com/intl/en/scholar/about.html

About IEEE Xplore. (n.d.). Retrieved February 4, 2024, from https://ieeexplore.ieee.org/Xplorehelp/overview-of-ieee-xplore/about-ieee-xplore

Adams, W. C. (2015). Conducting Semi-Structured Interviews. *Handbook of Practical Program Evaluation: Fourth Edition*, 492–505. https://doi.org/10.1002/9781119171386.CH19

Åkerstedt, T., Knutsson, A., Westerholm, P., Theorell, T., Alfredsson, L. & Kecklund, G. (2004). Mental fatigue, work and sleep. *Journal of Psychosomatic Research*, *57*(5), 427–433. https://doi.org/10.1016/J.JPSYCHORES.2003.12.001

Alahmadi, B. A., Axon, L. & Martinovic, I. (2022). 99% False Positives: A Qualitative Study of SOC Analysts' Perspectives on Security Alarms. *Usenix The Advanced Computing System Association*. https://www.usenix.org/conference/usenixsecurity22/presentation/alahmadi

Ancker, J. S., Edwards, A., Nosal, S., Hauser, D., Mauer, E. & Kaushal, R. (2017). Effects of workload, work complexity, and repeated alerts on alert fatigue in a clinical decision support system. *BMC Medical Informatics and Decision Making*, *17*(1), 1–9. https://doi.org/10.1186/S12911-017-0430-8

Anderson, R. (2007). *Thematic Content Analysis (TCA) Descriptive Presentation of Qualitative Data*. https://rosemarieanderson.com/wp-content/uploads/2014/08/ThematicContentAnalysis.pdf

Ault, M. & Brandley, B. (2023). The human cost of chronic mindfulness in U.S. law enforcement: toward a more nuanced understanding of HRO theory. *Journal of Applied Communication Research*, *51*(1), 18–36. https://doi.org/10.1080/00909882.2022.2118547

Ban, T., Takahashi, T., Ndichu, S. & Inoue, D. (2023). Breaking Alert Fatigue: AI-Assisted SIEM Framework for Effective Incident Response. *Applied Sciences 2023, Vol. 13, Page 6610*, *13*(11), 6610. https://doi.org/10.3390/APP13116610

Bans-Akutey, A. & Tiimub, B. M. (2021). Triangulation in Research. *Academia Letters*. https://doi.org/10.20935/AL3392

Carter, N., Bryant-Lukosius, D., Dicenso, A., Blythe, J. & Neville, A. J. (2014). The Use of Triangulation in Qualitative Research. *Number 5 / September 2014*, *41*(5), 545–547. https://doi.org/10.1188/14.ONF.545-547

Casey, S., Avalos, G. & Dowling, M. (2018). Critical care nurses' knowledge of alarm fatigue and practices towards alarms: A multicentre study. *Intensive and Critical Care Nursing*, *48*, 36–41. https://doi.org/10.1016/J.ICCN.2018.05.004

Caulfield, J. (2019). Using thematic analysis in psychology. *Qualitative Research in Psychology*, *3*(2), 77–101. https://doi.org/10.1191/1478088706QP063OA

Claudio, D., Deb, S. & Diegel, E. (2021). A Framework to Assess Alarm Fatigue Indicators in Critical Care Staff. *Critical Care Explorations*, *3*(6), E0464. https://doi.org/10.1097/CCE.0000000000000464

Ding, S., Huang, X., Sun, R., Yang, L., Yang, X., Li, X., Liu, J., Yang, H., Zhou, H., Huang, X., Su, F., Shu, L., Zheng, X. & Wang, X. (2023). The relationship between alarm fatigue and burnout among critical care nurses: A cross-sectional study. *Nursing in Critical Care*, *28*(6), 940–947. https://doi.org/10.1111/NICC.12899

Etikan, I., Musa, S. A. & Alkassim, R. S. (2015). Comparison of Convenience Sampling and Purposive Sampling. *American Journal of Theoretical and Applied Statistics 2016, Volume 5, Page 1*, *5*(1), 1–4. https://doi.org/10.11648/J.AJTAS.20160501.11

Hravnak, M., Pellathy, T., Chen, L., Dubrawski, A., Wertz, A., Clermont, G. & Pinsky, M. R. (2018). A call to alarms: Current state and future directions in the battle against alarm fatigue. *Journal of Electrocardiology*, *51*(6), S44–S48. https://doi.org/10.1016/J.JELECTROCARD.2018.07.024

Introduction - Web of Science platform - LibGuides at Clarivate Analytics. (n.d.). Retrieved February 4, 2024, from https://clarivate.libguides.com/webofscienceplatform

Kirsh, D., Maglio, P., Elkan, C., Elman, J., Flor, N., Hendler, J., Hutchins, E. & Matlock, T. (1994). On Distinguishing Epistemic from Pragmatic Action. *COGNITIVE SCIENCE*, *18*, 513–549. https://doi.org/10.1207/s15516709cog1804_1

Kokulu, F. B., Shoshitaishvili, Y., Soneji, A., Zhao, Z., Ahn, G. J., Bao, T. & Doupé, A. (2019). Matched and mismatched SOCs: A qualitative study on security operations center issues. *Proceedings of the ACM Conference on Computer and Communications Security*, 1955–1970. https://doi.org/10.1145/3319535.3354239

Lerdal, A., Moum, T., Wahl, A. K., Rustøen, T. & Hanestad, B. R. (2005). Fatigue in the general population: A translation and test of the psychometric properties of the Norwegian version of the fatigue severity scale. *Https://Doi.Org/10.1080/14034940410028406*, *33*(2), 123–130. https://doi.org/10.1080/14034940410028406

Levintha, D. & Rerup, C. (2006). *Crossing an Apparent Chasm: Bridging Mindful and Less-Mindful Perspectives on Organizational Learning on JSTOR*. Organization Science, Vol. 17, No. 4. https://www.jstor.org/stable/25146053

Makri, C. & Neely, A. (2021). Grounded Theory: A Guide for Exploratory Studies in Management Research. *International Journal of Qualitative Methods*, *20*. https://doi.org/10.1177/16094069211013654

Mayer, I. (2016). Qualitative research with a focus on qualitative data analysis. *International Journal of Sales, Retailing and Marketing*, *4*(9), 53–67. https://www.circleinternational.co.uk/wp-content/uploads/2021/01/IJSRM4-9.pdf#page=57

McElwee, S., Heaton, J., Fraley, J. & Cannady, J. (2017). Deep learning for prioritizing and responding to intrusion detection alerts. *Proceedings - IEEE Military Communications Conference MILCOM*, *2017-October*, 1–5. https://doi.org/10.1109/MILCOM.2017.8170757

McRee, G. R. (2022). Improved Detection and Response via Optimized Alerts: Usability Study. *Journal of Cybersecurity and Privacy 2022, Vol. 2, Pages 379-401*, *2*(2), 379–401. https://doi.org/10.3390/JCP2020020

Microsoft. (n.d.). *What is a security operations center (SOC)? | Microsoft Security*. Retrieved February 29, 2024, from https://www.microsoft.com/en-us/security/business/security-101/what-is-a-security-operations-center-soc

Mirza, H., Bellalem, F. & Mirza, C. (2023, 17. May). *(PDF) Ethical Considerations in Qualitative Research: Summary Guidelines for Novice Social Science Researchers*. https://www.researchgate.net/publication/370838199_Ethical_Considerations_in_Qualitative_Research_Summary_Guidelines_for_Novice_Social_Science_Researchers

Movahedi, A., Sadooghiasl, A., Ahmadi, F. & Vaismoradi, M. (2023). A grounded theory study of alarm fatigue among nurses in intensive care units. *Australian Critical Care*, *36*(6), 980–988. https://doi.org/10.1016/J.AUCC.2022.12.004

Myers, M. D. & Newman, M. (2007). The qualitative interview in IS research: Examining the craft. *Information and Organization*, *17*(1), 2–26. https://doi.org/10.1016/J.INFOANDORG.2006.11.001

Ndichu, S., Ban, T., Takahashi, T. & Inoue, D. (2021). A Machine Learning Approach to Detection of Critical Alerts from Imbalanced Multi-Appliance Threat Alert Logs. *Proceedings - 2021 IEEE International Conference on Big Data, Big Data 2021*, 2119–2127. https://doi.org/10.1109/BIGDATA52589.2021.9671956

Niederman, F. & March, S. (2019). The "Theoretical Lens" Concept: We All Know What it Means, but do We All Know the Same Thing? *Communications of the Association for Information Systems*, *44*(1), 1. https://doi.org/10.17705/1CAIS.04401

NSM. (2023). *Risiko 2023 Økt uforutsigbarhet krever høyere beredskap*. https://nsm.no/getfile.php/1312547-1676548301/NSM/Filer/Dokumenter/Rapporter/Risiko%202023%20-%20Nasjonal%20sikkerhetsmyndighet.pdf

Okoli, C. (2015). A Guide to Conducting a Standalone Systematic Literature Review. *Communications of the Association for Information Systems*, *37*(1), 43. https://doi.org/10.17705/1CAIS.03743

Priya, A. (2020). Case Study Methodology of Qualitative Research: Key Attributes and Navigating the Conundrums in Its Application.

Https://Doi.Org/10.1177/0038022920970318, *70*(1), 94–110. https://doi.org/10.1177/0038022920970318

PSNet. (2019, 7. September). *Alert Fatigue | PSNet*. https://psnet.ahrq.gov/primer/alert-fatigue

Raff, E., Filar, B. & Holt, J. (2020). Getting Passive Aggressive about False Positives: Patching Deployed Malware Detectors. *IEEE International Conference on Data Mining Workshops, ICDMW*, *2020-November*, 506–515. https://doi.org/10.1109/ICDMW51313.2020.00074

Recker, J. (2021). *Planning Your Research*. Springer, Cham. https://doi.org/10.1007/978-3-030-85436-2_3

Rinta-Kahila, T., Penttinen, E., Salovaara, A., Soliman, W. & Ruissalo, J. (2023). The Vicious Circles of Skill Erosion: A Case Study of Cognitive Automation. *Journal of the Association for Information Systems*, *24*(5), 1378–1412. https://doi.org/10.17705/1jais.00829

Robinson, N. (2023). HUMAN FACTORS SECURITY ENGINEERING: THE FUTURE OF CYBERSECURITY TEAMS. *EDPACS*, *67*(5), 1–17. https://doi.org/10.1080/07366981.2023.2211429

Salovaara, A., Lyytinen, K. & Penttinen, E. (2019). High reliability in digital organizing. *MIS Quarterly*, *43*(2), 555–578. https://doi.org/10.25300/MISQ/2019/14577

Seok, Y., Cho, Y., Kim, N. & Suh, E. E. (2023). Degree of Alarm Fatigue and Mental Workload of Hospital Nurses in Intensive Care Units. *Nursing Reports 2023, Vol. 13, Pages 946-955*, *13*(3), 946–955. https://doi.org/10.3390/NURSREP13030083

Wang, X., Yang, X., Liang, X., Zhang, X., Zhang, W. & Gong, X. (2024). Combating alert fatigue with AlertPro: Context-aware alert prioritization using reinforcement learning for multi-step attack detection. *Computers and Security*, *137*. https://doi.org/10.1016/J.COSE.2023.103583

Weick, K. E. & Sutcliffe, K. M. (2006). Mindfulness and the Quality of Organizational Attention. *Https://Doi.Org/10.1287/Orsc.1060.0196*, *17*(4), 514–524. https://doi.org/10.1287/ORSC.1060.0196

Wilken, M., Hüske-Kraus, D., Klausen, A., Koch, C., Schlauch, W. & Röhrig, R. (2017). Alarm Fatigue: Causes and Effects. *Studies in Health Technology and Informatics*, *243*, 107–111. https://doi.org/10.3233/978-1-61499-808-2-107

Wunderlich, M. M., Amende-Wolf, S., Krampe, H., Kruppa, J., Spies, C., Weiß, B., Memmert, B., Balzer, F. & Poncette, A. S. (2023). A brief questionnaire for measuring alarm fatigue in nurses and physicians in intensive care units. *Scientific Reports 2023 13:1*, *13*(1), 1–10. https://doi.org/10.1038/s41598-023-40290-7

Xiao, Y. & Watson, M. (2019). Guidance on Conducting a Systematic Literature Review. *Journal of Planning Education and Research*, *39*(1), 93–112. https://doi.org/10.1177/0739456X17723971/ASSET/IMAGES/LARGE/10.1177_0739456X17723971-FIG2.JPEG

Yin, R. K. (2009). *Case Study Research: Design and Methods*. 219. https://books.google.com/books/about/Case_Study_Research.html?hl=no&id=FzawIAdilHkC

Yuan, S. C., Chou, M. C., Chen, C. J., Lin, Y. J., Chen, M. C., Liu, H. H. & Kuo, H. W. (2011). Influences of shift work on fatigue among nurses. *Journal of Nursing Management*, *19*(3), 339–345. https://doi.org/10.1111/J.1365-2834.2010.01173.X

Appendix A – Overview of literature

Journal	Author	Title	Theme
BACCN	Shenglan Ding MSN et al, (2023)	The relationship between alarm fatigue and burnout among critical care nurses: A cross-sectional study	Alert fatigue, burnout, False positives, Alert management, Health sector.
Proceedings of the ACM Conference on Computer and Communications Security	Faris Bugra Kokulu et al, (2019)	Matched and Mismatched SOCs: A Qualitative Study on Security Operations Center Issues	False positives, SOC, Security Analysts, Cybersecurity.
Computers & Security	Xiaoyu Wang et al, (2024)	Combating alert fatigue with AlertPro: Context-aware alert prioritization using reinforcement learning for multi-step attack detection	Alert fatigue, SOC, Cybersecurity, Mitigation.
Australian Critical Care	Movahedi et al, (2023)	A grounded theory study of alarm fatigue among nurses in intensive care units	Alert Fatigue, Mitigation, Health sector.
Scientific Reports	Maximilian Markus Wunderlich et al, (2023)	A brief questionnaire for measuring alarm fatigue in nurses and physicians in intensive care units	Alert fatigue, Health sector.
Nursing reports	Yoonhee Seok et all, (2023)	Degree of Alarm Fatigue and Mental Workload of	Alert fatigue, Mitigation, Health sector.

		Hospital Nurses in Intensive Care Units	
Critical care explorations	Claudio David, (2021)	A Framework to Assess Alarm Fatigue Indicators in Critical Care Staff	Framework, Alert fatigue, Health sector.
Journal of Electrocardiology	Marilyn Hravnak et al, (2018)	A call to alarms: Current state and future directions in the battle against alarm fatigue	Alert fatigue, false positive Health sector.
Nursning managment	Yuan Su Chuan et al, (2011)	Influences of shift work on fatigue among nurses	Fatigue, Shift work, Health sector.
Australian Critical care	Ali Movahedi et al, (2023)	A grounded theory study of alarm fatigue among nurses in intensive care units	Alert fatigue, False positive, Health sector.
the Journal of Nursing Scholarship	Lauren S. Aaronson RN et al, (1999)	Defining and Measuring Fatigue	Fatigue, Framework.
Journal of Psychosomatic Research	T. Åkerstedt, et al., (2004)	Mental fatigue, work and sleep	Mental fatigue, Sleep, Healthsector.
Scandinavian Journal of Public Health	Anners lerdal et al, (2005)	Fatigue in the general population: A translation and test of the psychometric properties of the Norwegian version of the fatigue severity scale	Fatigue, scale.

BMC	Jessica S. Ancker, et al, (2017)	Effects of workload, work complexity, and repeated alerts on alert fatigue in a clinical decision support system	Alert fatigue, Health sector, Overload.
Applied Sciences	Tao Ban et al, (2023)	Breaking Alert Fatigue: AI-Assisted SIEM Framework for Effective Incident Response	SIEM, Alert Fatigue, Mitigation, Cybersecurity.
Intensive and Critical Care Nursing	Siobhán Casey et al, (2018)	Critical care nurses' knowledge of alarm fatigue and practices towards alarms: A multicentre study	Alert fatigue, Health sector.
IOS Press	Marc Wilken, et al, (2017)	Alarm Fatigue: Causes and Effects	Alert fatigue, Health sector
Journal of Cybersecurity and Privacy	Griffith Russell McRee, (2022)	Improved Detection and Response via Optimized Alers: Usability Study	Alert fatigue, Security Analyst, Cybersecurity. SOC, Mitigation.
IEEE International Conference on Big Data and Smart Computing	Samuel Ndichu, et al, (2022)	A Machine Learning Approach to Detection of Critical Alerts from Imbalanced Multi-Appliance Threat Alert Logs	Security Analyst, SIEM, SOC, false positives, Mitigation, Cybersecurity.
IEEE International Conference on Data Mining Workshops	Edward Raff et al, (2020)	Getting Passive Aggressive About False Positives: Patching Deployed Malware Detectors	False positive, SOC, Mitigation, Cybersecurity.
MILCOM, Military Communications Conference	Steven McElwee, et all, (2017)	Deep learning for prioritizing and responding to	Security Analyst, Cybersecurity, Alert fatigue, Mitigation.

		intrusion detection alerts	
USENIX - The Advanced Computing Systems Association	Bushra A. Alahmadi, et al, (2022)	99% False Positives: A Qualitative Study of SOC Analysts' Perspectives on Security Alarms	False positive, Alert fatigue, Cybersecurity, SIEM, SOC, Security analyst.
EDPACS	Robinson (2023)	Human factor security engineering: The future of cybersecurity teams	Human factor, Burnout, SOC, cognitive overload, Mitigation
Journal of the Association for information systems	Rinta-Kahila et al, (2023)	The Vicious Circles of Skill Erosion: A Case Study of Cognitive Automation	Cognitive automation, AI, Skill erosion, mindful conduction.

Appendix B – Consent form

Vil du delta i forskningsprosjektet

Work practices in SOC

Formålet med prosjektet

Dette er et spørsmål til deg om du vil delta i et forskningsprosjekt hvor formålet er å

- Prosjektet skal analysere om hvordan du opplever alert fatigue og om bedriften har iverksatt tiltak for å minimere dette.
- Forskningsprosjektet er en masteroppgave
- Personopplysningene kommer til å arkiveres til masteroppgaven er innlevert

Hvorfor får du spørsmål om å delta?

Du får denne forespørselen fordi du jobber i et område som vi mener blir påvirket av emnet vi undersøker. Vi ønsker å undersøke om dette er tilfellet eller ikke, og da trenger vi informasjon fra deg.

Hvem er ansvarlig for forskningsprosjektet?

Instituttet for informasjonssystemer ved Universitetet i Agder er ansvarlig for personopplysningene som behandles i prosjektet.

Det er frivillig å delta

Det er frivillig å delta i prosjektet. Det vil ikke ha noen negative konsekvenser for deg hvis du ikke vil delta eller senere velger å trekke deg.

Hva innebærer det for deg å delta?

- Metoden som skal brukes for å samle inn data er semi-strukturerte intervjuer.
- Opplysningene som smales inn er navn og stillingstittel.
- Opplysningene registrerer med elektroniske notater og lyd-/videoopptak.

Kort om personvern

Vi vil bare bruke opplysningene om deg til formålene vi har fortalt om i dette skrivet. Vi behandler personopplysningene konfidensielt og i samsvar med personvernregelverket. Du kan lese mer om personvern under.

Med vennlig hilsen

Wael Anwar Abdel Aziz Soliman Terje Heum Seljåsen

- Du kan lese mer om personvern på neste side.

Utdypende om personvern – hvordan vi oppbevarer og bruker dine opplysninger

Det er bare studentene som er involvert i oppgaven og veileder (Wael Anwar Abdel Aziz Soliman) som får tilgang til personopplysningen som blir delt.

Beskyttelse av data:

- Lydfilene vil ligge på OneDrive som er blitt tildelt gjennom Universitetet i Agder.
- Navn vil bli kodet og ligge på et adskilt ark fra øvrige data.
- Etter oppgaven er levert vil all data bli slettet.
- Ingen personopplysninger vi bli publisert.
- Transkriberingsverktøy som Word Diktering kan bli brukt for å behandle lydfiler.

Hva gir oss rett til å behandle personopplysninger om deg?

Vi behandler opplysninger om deg basert på ditt samtykke.

På oppdrag fra Instituttet for informasjonssystemer ved Universitetet i Agder har personverntjenestene ved Sikt – Kunnskapssektorens tjenesteleverandør, vurdert at behandlingen av personopplysninger i dette prosjektet er i samsvar med personvernregelverket.

Dine rettigheter

Så lenge du kan identifiseres i datamaterialet, har du rett til:

- å be om innsyn i hvilke opplysninger vi behandler om deg, og få utlevert en kopi av opplysningene,
- å få rettet opplysninger om deg som er feil eller misvisende,
- å få slettet personopplysninger om deg,
- å sende klage til Datatilsynet om behandlingen av dine personopplysninger.

Vi vil gi deg en begrunnelse hvis vi mener at du ikke kan identifiseres, eller at rettighetene ikke kan utøves.

Hva skjer med personopplysningene dine når forskningsprosjektet avsluttes?

Prosjektet vil etter planen avsluttes 01.07.2024

Opplysningene vil da slettes.

Spørsmål

Hvis du har spørsmål eller vil utøve dine rettigheter, ta kontakt med:

- Wael Anwar Abdel Aziz Soliman, wael.soliman@uia.no, tlf: 38142756
- Vårt personvernombud: Wael Anwar Abdel Aziz Soliman

Hvis du har spørsmål knyttet til Sikts vurdering av prosjektet, kan du ta kontakt på e-post: personverntjenester@sikt.no, eller på telefon: 73 98 40 40.

Samtykkeerklæring

Jeg har mottatt og forstått informasjon om prosjektet Alert fatigue, og har fått anledning til å stille spørsmål. Jeg samtykker til:

☐ å delta i semi-strukturert intervju

☐ at opplysninger om meg publiseres slik at jeg kan gjenkjennes

☐ at mine personopplysninger slettes etter prosjektslutt

Jeg samtykker til at mine opplysninger behandles frem til prosjektet er avsluttet

---

Appendix C – Interview guide

The questions below were used as guidance when we were conducting the interviews, if something interesting appeared during an interview, then we would often go away from the questions and ask around those interesting topics, and if needed, go back to these questions again if the interview went too far astray. This was done since we were using semi-structured interviews as our data collection method.

Del 1: Litt informasjon om personen som blir intervjuet, Ser etter mindful organizing

1. Fortell om stillingen din og forklar hva oppgaven din er. Hva er en vanlig dag på jobben

2. Hvor krevende er det for deg å klassifisere og identifisere hva alarmen innebærer?

3. Har dere en eskaleringsprosess, og hvordan fungerer den?

4. Hvor automatisert er denne?

5. Stiller du noen gang spørsmål til din egen analyse som er gjort, og hvis ja, hvorfor?
   (Hva er det som motiverer deg eller forhindrer deg i å stille kritiske spørsmål til dine egne analyser.)

6. Hvordan er mulighetene for å kryss-sjekke egne vurderinger?

7. Hvordan er SOC organisert? Er det kunnskaps overføring blant avdelinger?

8. Hvis en kritisk del av systemet i Socen går ned, hva gjør dere?

9. I deres daglige arbeid, føler dere at dere kan sammenligne deres arbeid med andre organisasjoner som driver med svært kritiske operasjoner, som f.eks flytrafikk?

10. Hvis du gjør en feil, kan dette være kritisk for kunden?

Del 2: Få informasjon om hvordan det er å jobbe på SOC og deres alarm mengde, om de føler noe av å kjenne tegnene til alert fatigue.

1. Fortell om stillingen din og forklar hva oppgaven din er. Hva er en vanlig dag på jobben

2. Hvor lang erfaring har du i SOC?

3. Hva gjør du når du responderer på alarmer?
    a. Har dere faste rutiner du følger eller tar du egne valg? Er det gitte rutiner du må følge når det kommer gitte alarmer?

4. Har dere muligheten til å endre eller undertrykke [supresse] alarmer, og hvordan gjør dere det?
    a. Hvor ofte hender det at du må gjøre dette?
    b. Hvilke typer alarmer er dette?

5. Er det lett for deg å estimere alvorlighetsgraden av en alarm basert på informasjonen du får? Eventuelt, hvordan estimere du alvorlighetsgraden av en alarm?

6. Får du nok informasjon av alarmene? hvorfor/hvorfor ikke?

7. Hvor ofte har du følt deg overveldet på jobb*, og hva er hovedgrunnen til det?*

8. *Hva gjør du i situasjoner der du støter på alarmer som ligner på tidligere alarmer?*

9. Hvordan håndterer du situasjoner hvor det dukker opp en alarm som du har undersøkt tidligere? (*Når det dukker opp en alarm som du har undersøkt tidligere, bruker du lengre tid på å reagere på den? (gjør du andre ting i mellomtiden fremfor å respondere?)*

10. Om du får inn store mengder alarmer, hvordan påvirker dette konsentrasjonen din i disse situasjonene?

11. Tar du raskere avgjørelser jo lengre ut i skiftet du er?
    a. Hvis skiftet inneholder masse alarmer, tar du da raskere avgjørelser enn om det
    b. kommer en alarm en gang iblant?

12. Har du, etter et skift, tenkt tilbake på en alarm du undersøkte og tenkt at du burde ha gjort en dypere eller bedre analyse?

13. Hvordan opprettholder du konsentrasjonen over lengre skift?

14. Vet du hva slags teknologi som er brukt for baksystemet?

15. Hvordan er mengden false positives du møter i et vanlig skift.
16. Hva er grunnen til at dere har mange/få FP?

17. Påstand: Noen analytikere kan prøve å spare på energi ved gjøre en raskere analyse av kjente alarmer, for å ha mer energi å bruke senere i skiftet. Kjenner du deg igjen i denne situasjonen?

Del 3: Her forteller vi dem at vi ser etter alert fatigue og prater litt rundt dette

1. Hvor ofte har alert fatigue blitt et tema under diskusjoner/møter.

2. Hva gjør du for å unngå AF?

3. Hvor mange alarmer har dere i løpet av en dag? Hva er mye, og hva er lite?