Understanding Fatigue in a Modern, AI-Enabled Security Operations Center

Warda Usman, Alycia Carey, Joshua Reynolds, Chris Fennell
Walmart

Abstract

Security Operations Center (SOC) analysts operate under sustained cognitive and time pressure, yet fatigue in modern, AI-enabled SOCs remains under-explored. We present a survey of 95 SOC professionals examining how fatigue varies across roles per NIST Functions, its key drivers, how analysts cope, and their recommended solutions. We find moderate but uneven fatigue, highest in Detect and Govern functions, driven primarily by time pressure, false alarms, limited recovery, and the high cost of errors. While generative AI is often perceived to reduce workload, its impact is mixed. Our findings highlight fatigue as a systemic, role-dependent challenge and suggest opportunities for improving SOC workflows, tooling, and recovery practices in AI-enabled environments.

1 Introduction

Security Operations Centers (SOCs) bring together people, processes, and technologies to manage high-volume security data and respond to cyber incidents [14, 16]. However, SOC work is inherently fatigue-inducing, requiring analysts to operate under sustained cognitive and operational demands [15]. Prior research has identified high alert volumes, false positives, inefficient workflows, turnover, and technical and managerial factors as key contributors to analyst fatigue and operational challenges [11–14].

However, much of this work reflects earlier stages of SOC evolution and does not fully capture today's landscape. The rapid increase in threat volume and sophistication, coupled with the growing integration of AI into both attack and defense, has fundamentally changed SOC work [4, 6]. Empirical evidence on analyst fatigue under these conditions, however, remains limited.

Understanding this is particularly important given the direct link between analyst performance and security outcomes, as well as the broader implications for workforce sustainability [5, 8, 9].

In this study, we aim to address this gap by examining fatigue in the SOC through both quantitative and qualitative lenses. We surveyed 95 SOC professionals, asking about their experiences of fatigue, the factors that contribute to it, how they cope with it, how AI impacts their job, and what they believe could alleviate it.

RQ1: How does the extent of fatigue vary across SOC functions in modern, AI-enabled SOC environments, and what factors contribute to it?

RQ2: How do SOC practitioners cope with fatigue, and what solutions do they perceive as most effective?

2 Methods

We conducted a survey to assess self-reported fatigue, contributing stress factors, perceptions of SOC practices, time pressure, and the impact of AI. Our survey included both structured questions (e.g., Likert-scale and multiple-choice) and open-ended responses to capture coping strategies and potential solutions. See Appendix A.1 for details.

Recruitment. We recruited participants a mature SOC at a large retail enterprise, with 270 personnel across multiple teams processing roughly 6 trillion points of telemetry [2]. Participants were recruited via an email sent to all SOC employees, inviting voluntary participation. This resulted in a self-selected sample (N = 95).

Mapping of Roles. To analyze fatigue across operational responsibilities, we systematically mapped participants' job functions (based on job descriptions or team roles) to the

USENIX Symposium on Usable Privacy and Security (SOUPS) 2025.
August 10–12, 2025, Seattle, WA, United States.

NIST Cybersecurity Framework (CSF) categories [10]: *Identify*, *Protect*, *Detect*, *Respond*, and *Govern*. This allowed us to examine how fatigue manifests across different types of security activities. No participants mapped to the *Recover* function in our sample.

Analysis. We conducted quantitative analyses to examine variation in fatigue across roles and its association with contributing factors. This included descriptive statistics, group comparisons, and regression modeling to assess relationships between workload characteristics and fatigue. In parallel, we analyzed open-ended responses using an inductive thematic approach, where we applied open coding and refined codes into broader themes.

Ethical Considerations. While our institution does not have a formal IRB, we followed similarly strict standards to ensure participant protection. Participation was voluntary, and all respondents provided informed consent before taking part. Participants were informed of the purpose of the study, how their data would be used, and that they could skip questions or withdraw at any time. All results are reported in aggregate to preserve confidentiality.

3 Findings

Fatigue Levels: We observed moderate fatigue overall, with variation across functional roles. Analysts in Detect (M = 3.06, SD = 0.85, n = 32) and Govern (M = 3.00, SD = 0.63, n = 11) report the highest levels, with 29% of Detect employees falling into the high-fatigue range. In contrast, Protect, Respond, and Identify report lower average fatigue and fewer high-fatigue responses.

Factors Contributing to Fatigue: The regression results suggest that fatigue at this SOC is driven primarily by false alarms ($\beta = +0.56$, $p < .05$) and time pressure ($\beta = +0.41$, $p < .05$), with lack of downtime ($\beta = +0.40$) also showing a notable but non-significant effect. Other factors, such as fear of errors, high cost of mistakes, and high-value targets, have smaller positive associations. In contrast to prior research highlighting task switching, staffing gaps, and tooling inefficiencies [14], our findings suggest fatigue is more concentrated around alert load, urgency, and limited recovery, along with the high-stakes nature of the work.

Time pressure: As shown in Table 2, participants in Detect reported the highest levels of time pressure (M = 3.39, median = 3.0), with 41.9% indicating that their work is largely driven by urgent, reactive tasks. In contrast, roles such as Identify reported comparatively lower time pressure (M = 2.14, with no respondents in the highest range). Across the sample, time pressure is positively associated with fatigue ($\beta = 0.23$, $p = 0.002$), suggesting that more reactive, interruption-driven work is linked to higher post-shift fatigue (see Table 4).

AI Impact: We also examined how participants perceive the impact of generative AI on their work. Most respondents (67%) reported that AI slightly reduces their workload, while about 16% reported an increase. Despite this generally positive perception, we found a significant association between perceived AI-related workload increases and fatigue ($\beta = 0.18$, $p = 0.044$). Analysts who report the highest increases in workload due to AI also report higher fatigue (M = 3.67) compared to those who report workload reductions (M = 2.56). This suggests that while AI can alleviate effort, it may also introduce new cognitive or operational burdens that contribute to fatigue when poorly integrated.

Coping Strategies: Participants described coping strategies that fall into two broad categories: active and passive recovery. Active strategies involve intentionally engaging the mind or body to reset after high-pressure work, most commonly through physical activity, time with family, friends, or coworkers, and hobbies or personal projects which provide fulfillment outside of work. In contrast, passive strategies emphasize low-effort mental disengagement, including "playing video games," "sleep," or simply "sit and decompress." Many participants also described stepping away from work environments as a way to create distance from work.

Barriers to Recovery: Participants highlighted notable barriers to recovery, including persistent high-priority workloads and limited opportunities to step away, reflected in statements such as "often it is not possible" and "there is always another high priority task." End-of-day fatigue further limits engagement in more active strategies (e.g., "need to exercise, but too tired"), and in some cases participants reported no meaningful disengagement at all.

Recommended Solutions: Participants emphasized reducing time pressure through better prioritization and "realistic deadlines," especially since "everything is urgent and important." They also highlighted improving signal quality by "reducing false positives" and adding automation. Addressing lack of downtime came up through calls for more staffing and the ability to step away after high-intensity work. To ease pressure around mistakes, participants noted the need for "clear SOPs." More broadly, repeated mentions of "more people," "less meetings during high priority tasks," and "protecting uninterrupted, deep-focus time," along with quieter spaces and WFH flexibility, reinforce the importance of reducing interruptions and increasing capacity.

4 Conclusion

The preliminary findings suggest that reducing fatigue in modern SOCs requires shifting focus from general workflow inefficiencies to the core drivers of alert burden, time pressure, and limited recovery. This points to the need for improving signal quality (e.g., reducing false positives [1,3,7]), better prioritization of alerts and tasks, and creating protected recovery time following high-intensity work.

The mixed impact of AI also highlights that simply introducing automation is insufficient; tools must be carefully integrated to reduce, rather than redistribute, cognitive load.

References

[1] Bushra A Alahmadi, Louise Axon, and Ivan Martinovic. 99% false positives: A qualitative study of {SOC} analysts' perspectives on security alarms. In *31st USENIX Security Symposium (USENIX Security 22)*, pages 2783–2800, 2022.

[2] Bradley Barth. Inside walmart global tech: Where cybersecurity isn't discounted. *SC World*, February 2023.

[3] Mohan Baruwal Chhetri, Shahroz Tariq, Ronal Singh, Fatemeh Jalalvand, Cecile Paris, and Surya Nepal. Towards human-ai teaming to mitigate alert fatigue in security operations centres. *ACM Transactions on Internet Technology*, 24(3):1–22, 2024.

[4] Check Point Software Technologies. 2025 cyber security predictions: The rise of ai-driven attacks, quantum threats, and social media exploitation, 2024. Published October 28, 2024. Accessed: 2026-04-21.

[5] Josiah Dykstra and Celeste Lyn Paul. Cyber operations stress survey COSS: Studying fatigue, frustration, and cognitive workload in cybersecurity operations. In *11th USENIX Workshop on Cyber Security Experimentation and Test (CSET 18)*, 2018.

[6] IBM. How ai-driven soc co-pilots will change security center operations, 2025. Accessed: 2026-04-21.

[7] Fatemeh Jalalvand, Mohan Baruwal Chhetri, Surya Nepal, and Cecile Paris. Alert prioritisation in security operations centres: A systematic survey on criteria and methods. *ACM Computing Surveys*, 57(2):1–36, 2024.

[8] Paul Kearney, Mohammed Abdelsamea, Xavier Schmoor, Fayyaz Shah, and Ian Vickers. Combating alert fatigue in the security operations centre. *Available at SSRN 4633965*, 2023.

[9] Filiz Mizrak, Hatice Gökçe Demirel, Okan Yaşar, and Turhan Karakaya. Digital detox: exploring the impact of cybersecurity fatigue on employee productivity and mental health. *Discover Mental Health*, 5(1):25, 2025.

[10] National Institute of Standards and Technology. The NIST Cybersecurity Framework (CSF) 2.0, 2024.

[11] Calvin Nobles. Stress, burnout, and security fatigue in cybersecurity: A human factors problem. *HOLISTICA–Journal of Business and Public Administration*, 13(1):49–72, 2022.

[12] Cassidy Norton, Zainab Ruhwanya, and Jacques Ophoff. Factors contributing to cybersecurity fatigue. In *International Symposium on Human Aspects of Information Security and Assurance*, pages 198–211. Springer, 2025.

[13] Sathya Chandran Sundaramurthy, Alexandru G. Bardas, Jacob Case, Xinming Ou, Michael Wesch, John McHugh, and S. Raj Rajagopalan. A human capital model for mitigating security analyst burnout. In *Proceedings of the Eleventh USENIX Conference on Usable Privacy and Security*, SOUPS '15, page 347–359, USA, 2015. USENIX Association.

[14] Shahroz Tariq, Mohan Baruwal Chhetri, Surya Nepal, and Cecile Paris. Alert fatigue in security operations centres: Research challenges and opportunities. *ACM Computing Surveys*, 57(9):1–38, 2025.

[15] Tines. State of mental health in cybersecurity, 2022. Report based on survey of 1,027 cybersecurity professionals. Accessed: 2026-04-21.

[16] Manfred Vielberth, Fabian Böhm, Ines Fichtinger, and Günther Pernul. Security operations center: A systematic study and open challenges. *Ieee Access*, 8:227756–227779, 2020.

A Appendix

A.1 Survey Instrument

1. How would you rate your usual fatigue level after a shift?
   - Not at all fatigued
   - Slightly fatigued
   - Moderately fatigued
   - Very fatigued
   - Extremely fatigued
2. Which of the following issues regularly contribute to stress? (Select all that apply)
   - High false-alarm rate
   - Disconnected or overloaded dashboards
   - Slow access to logs/telemetry
   - Manual repetitive tasks
   - Lack of automation for routine operations
   - Inefficient or unclear SOPs
   - Staff and skills shortages (insufficient expertise or bandwidth)
   - The fear of missing something and/or the fear of getting something wrong
   - Ineffective or unclear communication with client team or CISO
   - Managing expectations from multiple stakeholders
   - The need to act quickly/time pressure

- Constant switching between tasks
- Lack of downtime/rest
- High cost of a mistake
- High value associated to the project
- Other: [open response]

3. Rate each of the following statements (from Strongly Disagree to Strongly Agree)
   - Urgent tasks I receive are relevant
   - Our dashboards and tools are well-integrated and easy to use
   - Our SOPs are clear and practical in high-pressure situations
   - Automation exists for repetitive tasks
   - Time pressure forces trade-offs that may reduce work quality
4. Which statement best describes how you spend most of your time at work?
   - I consistently have time for strategic or proactive work
   - I often have time for strategic or proactive work
   - I sometimes have time for strategic or proactive work
   - I rarely have time for strategic or proactive work
   - My work is almost entirely driven by urgent or reactive tasks
5. Rate the following statement: I don't have enough time to recover my full energy between shifts (from Strongly Disagree to Strongly Agree)
6. Overall, how has generative AI impacted your job?
   - Made my job much easier
   - Made my job somewhat easier
   - No noticeable impact
   - Made my job somewhat harder
   - Made my job much harder
7. Have you experienced any of the following due to high priority tasks? (Select all that apply)
   - Stress/Anxiety
   - Insomnia
   - Burnout
   - Negative impact on social life/relationships
   - Physical reactions (e.g. weight gain or loss)
   - Panic attacks
   - Other: [open response]
8. What do you do to unwind after high priority tasks at work?
9. What single change would most reduce your fatigue during high priority work tasks?
10. Total number of years working in current role
11. Which team are you on?

A.2 Quantitative Analysis Details

A.2.1 Fatigue Levels

Table 1 presents the descriptive statistics for reported fatigue levels grouped by NIST functions.

Table 1: Fatigue by NIST Function. % High = proportion rating fatigue 4 or 5 out of 5.

NIST Function	N	M	Mdn	SD	SE	% High
Identify	14	2.57	2.0	0.73	0.20	14.3%
Protect	11	2.73	3.0	0.75	0.23	9.1%
Detect	32	3.06	3.0	0.84	0.15	29.0%
Respond	27	2.64	3.0	0.77	0.15	7.1%
Govern	11	3.00	3.0	0.60	0.18	18.2%

A.2.2 Factors Contributing to Fatigue

Table 2 presents OLS regression results examining which self-reported stress factors independently predict post-shift fatigue.

Table 2: OLS Regression: Stress Factors Predicting Post-Shift Fatigue. $N = 96$, $R^2 = 0.37$, $F(16, 79) = 2.84$, $p = 0.001$.

Factor	β	SE	p	95% CI
False alarms	0.56	0.25	.025*	[0.073, 1.057]
Time pressure	0.41	0.18	.021*	[0.063, 0.765]
Lack of downtime/rest	0.40	0.20	.051	[−0.003, 0.809]
Fear of missing/error	0.23	0.16	.164	[−0.096, 0.558]
High cost of a mistake	0.22	0.18	.216	[−0.133, 0.579]
High value targets	0.20	0.24	.423	[−0.288, 0.679]
Task switching	0.08	0.17	.665	[−0.270, 0.422]
Slow access to logs	0.02	0.22	.930	[−0.411, 0.449]
Managing expectations	0.01	0.18	.966	[−0.348, 0.364]
Manual repetitive tasks	−0.00	0.26	.995	[−0.528, 0.525]
Disconnected dashboards	−0.03	0.30	.933	[−0.625, 0.574]

* $p < .05$.

A.2.3 Time Pressure

Table 3 presents descriptive statistics for time pressure scores across participants grouped by NIST CSF function.

Table 3: Time Pressure by NIST Function. % High = proportion indicating largely urgent/reactive work.

NIST Function	N	M	Mdn	SD	SE	% High
Identify	14	2.14	2.0	0.74	0.20	0.0%
Protect	11	3.00	3.0	1.21	0.36	27.3%
Detect	32	3.39	3.0	1.01	0.18	41.9%
Respond	27	2.89	3.0	1.20	0.23	22.2%
Govern	11	2.91	3.0	0.51	0.16	9.1%

A.2.4 Time Pressure and Fatigue

Table 4 presents the OLS regression results testing whether time pressure predicts self-reported fatigue.

Table 4: OLS Regression: Time Pressure predicting self-reported fatigue. $N = 95$, $R^2 = 0.10$, $F(1,93) = 10.43$, $p = 0.002$.

Term	β	SE	95% CI	p
Intercept	2.14	0.23	[1.69, 2.59]	<.001
Time Pressure (Q25)	0.23	0.07	[0.09, 0.37]	.002

A.2.5 AI Impact

Table 5 presents the OLS regression results testing whether perceived AI-related workload change predicts self-reported fatigue.

Table 5: OLS Regression: AI Impact predicting self-reported fatigue. $N = 95$, $R^2 = 0.04$, $F(1,94) = 4.19$, $p = 0.044$.

Term	β	SE	95% CI	p
Intercept	2.39	0.22	[1.942, 2.829]	<.001
AI Impact (Q24)	0.18	0.09	[0.005, 0.347]	.044*

* $p < .05$.