

Alert Fatigue Across Three Generations of Security Operations Centre: Traditional, Automated, and Autonomous

Smruti Suman Badjena, Department of CSE, GIET University, Gunupur, Odisha, India,
smrutisumanbadjena@giet.edu

Sanket Nayak, Department of CSE, GIET University, Gunupur, Odisha, India, sanketnayak@giet.edu

Kalikesh Pradhan, Department of CSE, GIET University, Gunupur, Odisha, India, kalikeshpradhan@giet.edu

Abstract: Alert fatigue is the cause of both analyst fatigue and the equally expensive consequences of undetected threats. Prior work has described this challenge mainly on a qualitative level, identifying problems such as rule decay, SOC tool sprawl, lack of context, and poor prioritization of alerts. Yet, scarce empirical assessments exist that reveal how changes in SOC design may affect those drivers. In this paper we introduce a framework to analyzing alert to disposition event data under three common architectures: Traditional (SIEM), Automated (SOAR), and Agentic (Autonomous). We track the progress of our fatigue measure – Alert Volume Reduction Rate (ARR), False Positive Density (FPD), Tool Overlap Duplication Index (TODI), and Mean Time-to-Disposition (MTTD) – to characterize shifts across this regime. We find that although automation reduces alarm volumes, it generates new sources for systemic fatigue that can only be mitigated by truly Autonomous Systems

Keywords—Alert Fatigue, SOC, SIEM Tools, SOAR, False positive rate, Alert Triage, Empirical Security analytics

1. Introduction

Modern enterprise systems produce security telemetry data too quickly for SOC analysts to handle in an adequate manner, exacerbated by factors such as multi-cloud deployments, remote work environments, and the proliferation of IoT devices. Industry research puts the number of alerts per analyst at about 4,484 per day, with 67% of them ignored, and 83% of the alerts that receive attention being false positives [1]. Personal experience supports this with the average number of alerts per analyst being up to 174 in one day, but only one out of every four of those is worth any further action [2].

False-positive rates can be as high as 99% in some environments, and a peer-reviewed USENIX Security study of working analysts independently confirms the same pattern: false positives, not detection misses, account for >90% of data seconds spent by analysts [9]. This cumulative effect — now known as Alert Fatigue — is brutally reflected in workforce statistics, with burnout rates of 67–71% and annual turnover approaching 28% across the industry [1], [2].

Foundational surveys have done useful work cataloguing the qualitative root causes—rule misconfiguration, tool sprawl, context deficit, poor prioritisation—most notably the 2025 ACM Computing Surveys taxonomy of alert fatigue [3]. However, these accounts are also united in using both literature synthesis and analyst self-report, not direct measurement through telemetry. Furthermore, these accounts typically also treat the ‘SOC’ as a monolith and it is unclear how the drivers discussed change when transitioning from a classic human SOC, to an automated SOC with SOAR and then to a contemporary SOC in the automated or even autonomous stage.

This paper addresses that gap with three contributions:

1. A Cross-Generational Taxonomy that formally distinguishes between Traditional (SIEM-Centric), Automated (SOAR-integrated), and Autonomous (Agentic AI) SOC architectures;
2. A Framework for Alert Telemetry that defines four measurable alert fatigue drivers: Alert Volume Reduction Rate (ARR), False Positive Density (FPD), Tool Overlap Duplication Index (TODI) and Mean Time-to-Disposition (MTTD);
3. A data-driven evaluation framework that explicitly correlates these four metrics to the root causes identified in [3], enabling one to rank the dominance of root-cause from alert-to-disposition logs as opposed to simply assuming it based on the literature.

2. Evolutionary Taxonomy of the Three Generations

Before we calculate the level of alert fatigue present, we need to strip out the architectural variables that differentiate one security operations generation from another. In the traditional SOC, a SIEM solution is the heart of the matter ingesting unstructured raw logs, event logs, Net Flows etc. Triage is a manual function by design where T1 analysts write fixed correlation rules to be applied manually and the resulting alerts are examined by human operators. As a result – a deluge of unstructured data with thin contextual intelligence and significant false positives given the static thresholds don't adapt as network traffic naturally drifts.

The Automated SOC brings a SOAR component to this infrastructure and transforms triage to semi-automated using code to implement playbooks. The playbooks collect contextual data where required, block IPs or hosts as defined by rules, or close down known false positives automatically to reduce the overall burden of raw alerts. The trade-off for reducing volume is the added burden of continuously updating those playbooks, including cascading logic failures whenever a 3rd party API or log schema changes. The SOC automation introduces an additional component on top, AI agents, that apply reason across diverse inputs to gather intelligence instead of simply running code to implement predetermined outputs.

Supported by an XDR layer which is the automatically correlating data sources that used to live in disparate security products these agents then enrich, validate intent, and even self-optimize the rules as time goes by. Here, triage is automated to the extent that an entire end-to-end attack scenario is validated before it's seen by a human, rather than individual anomalies generating the event on its own. Table 1 captures the telemetry profile of each generation using published SOC benchmark figures [1], [2], [4], not hypothetical estimates.

Table 1: Generational Telemetry Matrix

Vector / Metric	Gen 1: Traditional	Gen 2: Automated	Gen 3: Autonomous
Primary Triage Engine	Human Analyst (Tier-1)	Deterministic SOAR Playbook	Multi-Agent AI Frameworks
Daily Alert Volume	~ 4,484 alerts/day in queue [1]	120-150 alerts/day post-tiering [4]	multi-stage attack graphs (< 5:1 escalation metric)
Dominant Root Cause	Rules Misconfiguration and Context Deficit	Cross-Tool Overlap and Playbook Logic Fractures	Prompt/Context Drift and Agent Reasoning Errors

False Positive Distribution	Maxima approaching 83% of the human queue [1]	Suppressed for static patterns; high in edge cases [2]	Minimal; shifts to verification of agent reasoning
Mean Time-to-Disposition (MTTD)	Prolonged; ~ 2.7 hours/day spent in triage [1]	Bimodal: sub-second or hours on logic failure	Optimised: seconds to minutes per escalations

3. Empirical Methodology and Mathematical Metrics

Four mathematical functions are defined here, each one ingesting real security logs to calculate the actual distribution of alert fatigue root causes.

3.1 Alert Volume Reduction Rate (ARR): ARR measures how effectively a SOC generation filters raw event noise before any human involvement:

$$ARR = 1 - (A_{escalated} / A_{raw})$$

Here, A_{raw} is the total security events entering the correlation layer, and, $A_{escalated}$ is the subset forwarded to an analyst's queue.

3.2 False Positive Density (FPD): Misconfigured rules continue to be a primary source of fatigue, which FPD models on a per-rule basis. RNN-based classifiers incorporated into SIEM processes have shown accuracies of more than 99% [8] on benchmark sets, thereby proving that this is indeed measurable:

$$FPD(R_i) = AFP(R_i) / A_{Total}(R_i)$$

$AFP(R_i)$ counts alerts from rule R_i disposed as "False Positive" or "Benign." A value approaching 1.0 identifies an overly broad or misconfigured signature.

3.3 Tool Overlap Duplication Index (TODI): Tool sprawl manifests when multiple security tools independently flag the same underlying event, each generating its own discrete alert. We capture this with a temporal clustering window Δt :

$$TODI = [\sum_{m=1}^M (D_m - 1)] / A_{Total}$$

D_m counts the unique security-tool alerts grouped within that window; M is the total number of actual security intrusions across the period. A high TODI value signals that tool sprawl, rather than genuine threat volume, is what's multiplying the analyst's workload.

3.4 Mean Time-to-Disposition (MTTD): The size of that delay is captured by tracking the time delta of manual validation:

$$MTTD = (1/N) \sum_{k=1}^N (T_{disposition,k} - T_{assignment,k})$$

Here, $T_{assignment}$ is the timestamp when an alert lands in an analyst's queue, and $T_{disposition}$ is the timestamp the analyst closes the alert, averaged across N alerts.

4. Data Collection and Root-Cause Mapping

When implementing the SOC model in practice, one uses a pipeline consuming directly from the telemetry pipeline of the SOC data layer and only retrieves logs relevant to what is needed. Security raw logs are ingested through the time series parse layer for normalization and further down in the four main metrics produced, fed into the root-cause rank matrix. For our purposes, as for all three generations of the SOC, we use a common disposition scheme: 'True Positive

Actionable, True Positive Non-Actionable, False Positive Tuning Required, False Positive Benign Anomaly. The logs collected are from a baseline of 90 days of alert metadata (Alert ID, time-stamp, initiating control, rule id, alert severity) and analyst action logs (Analyst ID, ticket ID, time-stamp, ticket state, note length).

Instead of assuming which root cause plays the biggest part, this framework maps each metric directly to structural failures in the literature that have already been given a proper definition. If average FPD of the 10% highest-volume rules is high, then Rule Misconfiguration dominance holds. If TODI is over 30% of total alert-volume variation, then Tool Sprawl dominance also holds (If MTTD grows non-linearly with the number of external assets an analyst must pivot across: Context Deficit dominance holds. Poor Prioritisation dominance is indicated when MTTD for high-severity alerts statistically similar to that of low-severity alerts, an indication that severity labels are no longer serving a purpose to guide analyst attention.

5. Results and Discussion

In the Traditional SOC, Rule Misconfiguration and Context Deficit sit at the top of the fatigue distribution. In the absence of an orchestration layer that would filter the raw data before the analysts start working on it, the analysts will need to manually go through each of the anomalies that come through the static thresholds, which industry experts say this process alone takes up to 2.7 hours of their time per day with more than half of it going to false positives [1], [2]. In the Automated SOC, the fatigue curve changes rather than vanishes. SOAR playbooks work great at suppressing low fidelity alerts known before, and bring FPD numbers down, but due to the duplication of tools and the failures of the playbooks, there is an increase in the TODI curve. This means analysts need to deal with a new cognitive load problem, which comes from the failure in the execution of the playbooks and having to switch context among tools in the middle of the process. According to the SANS 2024 SOC survey, the problem of maintenance of playbooks and automation and not headcount shortages is the limitation of SOAR operations [5]. The Autonomous SOC has the lowest level of raw alert volume among the three, mainly due to the fact that full attack narratives are confirmed by AI agents prior to escalation to a human operator, which is characteristic of agentic SOC's valuing context-based verification above event-driven triggers [6]. However, such an approach involves a downside called Verification Inertia. In the case of escalated alerts that are rare and pre-validated, the effort required of the analyst becomes different: instead of handling huge amounts of simple data, he/she is supposed to process dense, AI-created attack narratives in a time-pressed manner. At that, such changes happen right when the dwell time of the attackers gets significantly reduced: according to Mandiant's 2026 M-Trends report, there is only 22 seconds' gap between first access and lateral movement for modern attacks, whereas in 2022 there were more than eight hours.

Deployment experience demonstrates the reality of this trade-off. The CORTEX multi-agent LLM triage system, for example, lowered false-positive rates from 24.9% to 14.2% — but achieved it with an approximately 3.4 higher latency per alert processed, indicating that the autonomy-based architecture has cost associated, not savings [11]. Similarly, AgentSOC reports that the overall time to reason is less than a second (~506 ms), but of that the large majority (~95%) is due to the LLM inference [10]. The operational challenge, accordingly, has moved from alert queues to the reasoning and supervision costs. An across-system study of LLMs- and agent-based SOC automation also reaches the same conclusion in a different way: today's systems operate as advisory copilots, with no containment decisions possible without human approval [12].

6. Conclusion

In this paper, SOC alert fatigue is moved beyond being studied purely on a qualitative basis to one where it can be quantitatively measured via an examination of the three different generations of security operations and their alert fatigue via log-level telemetry. These four metrics, namely, ARR, FPD, TODI, and MTTD, have been developed and together offer a diagnostic framework that is not possible through any practitioner surveys since the former are based on the data available at the alert-to-disposition level rather than the latter. The important result is that alert fatigue is not solved by automation but rather that automation changes the form of alert fatigue. For the Traditional SOC, the issues revolve around rule misconfiguration and context problems, while the Automated SOC replaces them with tool redundancy and playbook fragility. The next step in line would be to calibrate the threshold – to test out ARR, FPD, TODI, and MTTD against anonymized production logs on an industrial scale, across different verticals, to set up a defensible benchmark. In terms of design of the experiment, a longitudinal design in which we follow the progress of a single organization from conventional security to automation to autonomous security would yield much more convincing causal evidence than a comparative study because organizational differences would remain constant, making it much easier to detect generational differences.

References

1. Vectra AI, “2023 State of Threat Detection Report,” Vectra AI Research, White Paper, 2023.
2. UnderDefense, “Alert Fatigue in Cybersecurity: the SOC playbook to eliminate it,” *UnderDefense Blog*, Apr. 2026.
3. S. Tariq, M. Baruwat Chhetri, S. Nepal, and C. Paris, “Alert Fatigue in Security Operations Centres: Research Challenges and Opportunities,” *ACM Comput. Surv.*, vol. 57, no. 9, Art. 224, pp. 1-38, Apr. 2025, doi: 10.1145/3723158.
4. R. Perkins, “Why More Analysts Won’t Solve your SOC’s Alert Problem,” *Prophet Security/BleepingComputer*, May 2026.
5. C. Crowley, B. Filkins, and J. Pescatore, “SANS 2024 SOC Survey,” SANS Institute White Paper, 2024.
6. Intezer, “Autonomous SOC: Components, Maturity Model, and Adoption Guidelines,” *Intezer Guides*, Feb. 2026.
7. Google Cloud/ Mandiant, “M-Trends 2026: Data, Insights, and Strategies from the Frontlines,” Google Cloud/ Mandiant, Tech. Rep., Mar. 2026.
8. M. N. Aktar, M. N. Sakib, A. Hossain, A. Ullah, K. H. Robin, and K.M. Rahman, “Enhancing False Positive Alert Detection in Security Information and Event Management System Using Recurrent Neural Network,” in *Proc. 27th Int. Conf. Computer and Information Technology (ICCIT)*, 2024, doi: 10.1109/ICCIT64611.2024.11022066.
9. B. A. Alahmadi, L. Axon, and I. Martinovic, “99% False Positives: A Qualitative Study of SOC Analysts’ Perspectives on Security Alarms,” in *Proc. 31st USENIX Security Symp.*, 2022, pp. 2783-2800, arXiv:2210.01649.
10. J. Roy and S. K. Singh, “AgentSOC: A Multi-Layer Agentic AI Framework for Security Operations Automation,” *arXiv preprint arXiv:2604.20134*, 2026.
11. B. Wei, Y. S. Tay, H. Liu, J. Pan, K. Luo, Z. Zhu, and C. Jordan, “CORTEX: Collaborative LLM Agents for High-Stakes Alert Triage,” *arXiv preprint arXiv:2510.00311*, 2025.
12. S. Srinivas et al., “AI-Augmented SOC: A Survey of LLMs and Agents for Security Automation,” *J. Cybersecurity and Privacy*, vol. 5, no. 4, p. 95, Nov. 2025.