ISSN: 2278-7518

CYBER TIMES® INTERNATIONAL JOURNAL OF TECHNOLOGY AND MANAGEMENT

Volume 19 - Issue 01, October 2025 - March 2026
Bi-Annual Double Blind Peer Reviewed Refereed Journal

CYBER TIMES®
(Leader in innovative Tech-World)

Cyber Times International Journal of Technology & Management

Volume 19 - Issue 01, October 2025 – March 2026
ISSN: 2278-7518

EDITOR-IN-CHIEF

Dr. Anup Girdhar

EDITORIAL ADVISORY BOARD

Dr. Sushila Madan
Dr. A.K. Saini
Advocate Mukul Girdhar
Ms. Sonia Girdhar

EXECUTIVE EDITORS

Ms. Samridhi Girdhar
Mr. Rakesh Laxman Patil

Scientific Journal Impact Factor Value for 2024 = 6.007

"***Cyber Times International Journal of Technology & Management***". All rights reserved. No part of this journal may be reproduced, republished, stored, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior permission of the publisher in writing. Any person who does any unauthorized act in relation to this journal publication may be liable to criminal prosecution and civil claims for damages.

Editorial Office & Administrative Address:

Delhi:
The Editor,
A19/1, Mansa Ram Park,
New Delhi-110059.

ISSN: 2278-7518

Phone: +91-9811485729, +91-9312903095

Website: https://journal.cybertimes.in

Email: editor@cybertimes.in

Disclaimer: Views and information expressed in the Research Papers or Articles are those of the respective authors. "***Cyber Times International Journal of Technology & Management***", its Editorial Board, Editor and Publisher (Cyber Times) disclaim the Responsibility and Liability for any statement of fact or option made by the contributors. The content of the papers are written by their respective authors. The originality and authenticity of the papers and the explanation of information and views expressed therein are the sole responsibility of the authors. However, effort is made to acknowledge source material relied upon or referred to, however; "***Cyber Times International Journal of Technology & Management***" does not accept any responsibility for any unintentional mistakes & errors.

From the Editor's Desk

At the outset, I take this opportunity to express my sincere gratitude to all the Editorial Board Members, Editors, Peer Review Members, contributors, and readers for making ***Cyber Times International Journal of Technology & Management*** an outstanding success. Their unwavering support, dedication, and commitment to academic excellence have significantly contributed to the growth and reputation of the journal.

We are pleased to present **Volume 19 – Issue 01** of ***Cyber Times International Journal of Technology & Management***. This issue features a collection of high-quality research papers and scholarly articles that reflect contemporary developments, innovative ideas, and critical insights across emerging areas of Technology, Management, Law, Education, and other multidisciplinary domains. The diversity of topics covered in this issue highlights the increasing importance of interdisciplinary research in addressing global challenges and opportunities.

The overwhelming response received from researchers, authors, academicians, law-enforcement agencies, and industry professionals for submitting their research papers and articles is deeply appreciated and duly acknowledged across the globe. Their valuable contributions have enriched the journal's content and strengthened its role as a platform for disseminating knowledge, fostering innovation, and encouraging scholarly dialogue among academia, industry, and society.

On behalf of the Editorial Team, I extend my heartfelt thanks to all authors for their valuable research contributions and to our reviewers for their constructive evaluations that help maintain the highest standards of publication quality. We hope that the research published in this issue will inspire further inquiry, collaboration, and advancement in various fields of study, while continuing to serve as a meaningful resource for our readers worldwide.

We look forward to receive your valuable and future contributions to make this journal a joint endeavor.

With Warm Regards,

Dr. ANUP GIRDHAR

Editor-In-Chief
Cyber Times International Journal of Technology & Management

General Information

- ***"Cyber Times International Journal of Technology & Management"*** is published bi-annually. All editorial and administrative correspondence for publication should be addressed to The Editor, Cyber Times.

- The Abstracts received for the final publication are screened by the Evaluation Committee for approval and only the selected Papers/ Abstracts will be published in each edition. Further information is available in the "**Guidelines for paper Submission"** section.

- Annual Subscription details for obtaining the print copy of the journal are provided separately and the interested persons may avail the same accordingly after filling the Annual subscription form.

- This journal is meant for education, reference and learning purposes. The author(s) of this of the book has/have taken all reasonable care to ensure that the contents of the book do not violate any existing copyright or other intellectual property rights of any person/ company/ institution in any manner whatsoever. In the event the author(s) has/have been unable to track any source and if any copyright has been inadvertently infringed, please notify the publisher in writing for the corrective action.

- Copyright © "***Cyber Times International Journal of Technology & Management***". All rights reserved. No part of this journal may be reproduced, republished, stored, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior permission of the publisher in writing. Any person who does any unauthorized act in relation to this journal publication may be liable to criminal prosecution and civil claims for damages.

- **Other Publications:**
  - Cyber Times Newspaper (English) – RNI No: DELENG/2008/25470
  - Cyber Times Newspaper (Hindi) – RNI No. DELHIN/1999/00462

- **Printed & Published by:** Cyber Times
  A19/1, Mansa Ram Park, New Delhi-110059

CTIJTM Editorial Advisory Board Members

Name	Designation, Organization/ University	Country
Dr. Sushila Madan	Associate Professor, Delhi University	India
Dr. A. K. Saini	Professor, GGS IP University	India
Mr. J. R. Ahuja	Former Consultant, AICTE	India
Mr. Mukul Girdhar	Advocate Delhi High Court	India
Mr. Geetesh Madan	Q.A. Consultant with Tesco Bank, Newcastle	UK
Dr. Deepak Shikarpur	Chairman Board of Studies, Pune University	India
Dr. B. B. Ahuja	Deputy Director, COE - Pune	India
Prof. M. N. Hoda	Director, Bharati Vidyapeeth's (BVICAM)	India
Dr. S. C. Gupta	Director, NIEC, GGS IP University	India
Dr. S. K. Gupta	Professor, IIT Delhi	India
Dr. K. S. Wani	Principal, SSBT’s COET, Bambhori, Jalgaon	India
Dr. K. V. Arya	Associate Professor, IIITM, Gwalior	India
BRIG. Dr. S.S. Narula	Director, Gitarattan International Bussiness School	India
Dr. Sarika Sharma	Director, JSPM'S ENIAC Institute of CA, Pune	India
Dr. S.K.M. Bhagat	Prof. & Head, MIT Academy of Engg., Pune	India
Dr. Jack Ajowi	Jaramogi Oginga Odinga University of Sci. & Tech.	Kenya
Dr. Srinivas Sampalli	Professor, Dalhousie University, Halifax	Canada
Dr. Ijaz A. Qureshi	V.P. Academic Affairs, JFK Inst. of Tech. and Mgmt.	Pakistan
Aryya Bhattacharyya	Director, CIP, Columbus State University	US
Dr. M. M. Schiraldi	Assistant Professor, 'Tor Vergata' University of Rome	Italy

Executive Editorial Advisory Board Members

Name	Designation, Organization/ University	Country
Ms. Kanika Trehan	Editor - Cyber Times, New Delhi	India
Mr. Rakesh Laxman Patil	Editor - Cyber Times, Pune	India
Adv. Tushar Kale	Cyber Lawyer, Pune	India
Adv. Neeraj Aarora	Cyber Lawyer, New Delhi	India
Mr. Sanjeev Sehgal	HOD, SJP Polytech, Damla, Haryana	India
Mr. Rajinder Kumar Bajaj	GM, Satake India Engg. Pvt. Ltd., (Japan)	India
Dr. B. M. Patil	Associate Professor MIT, Pune	India
Dr. Rajesh S. Prasad	Professor, DCOER, Pune University	India
Dr. Binod Kumar	Associate Professor, MIT Academy of Engg, Pune	India
Prof. Dr. M. Husain	HOD, SSBT’s COET, Bambhori, Jalgaon	India
Prof. Dr. U. S. Bhadade	HOD, SSBT’s COET, Bambhori, Jalgaon	India
Dr. V. N. Wadekar	Prof. & Head, MIT college of Engg. CMSR, Pune	India
Dr. M.D. Goudar	Associate Prof. & Head, Pune University	India
Dr. Mohd. Rizwan Alam	Sr. Lecturer, Amity University	Dubai
Prof. Jagannath Aghav	Professor & Head, CSE & IT, COE - Pune	India

Disclaimer: The names, affiliations, and designations of the Editorial Board Members published in this journal are based on the information provided at the time of their association with *Cyber Times International Journal of Technology & Management*. Members may subsequently change their positions, affiliations, or professional designations. The journal does not guarantee the continued accuracy of such details after publication.

Cyber Times International Journal of Technology and Management - CTIJTM

Volume 19 - Issue 01

CONTENTS

Beyond Alert Fatigue: Integrating LLM-Driven Triage and Autonomous Incident Response in Modern Security Operations Centers

Dr. Anup Girdhar
CEO – Sedulity Solutions & Technologies, New Delhi, India
Email: anupgirdhar@gmail.com

ABSTRACT

Security Operations Centers (SOCs) represent the institutional frontline of enterprise cyber defence, yet their operational model is under acute strain. The volume of security alerts generated by contemporary enterprise environments has grown at a pace that human analyst capacity cannot match — a phenomenon documented across jurisdictions as alert fatigue, wherein analysts become desensitised to notifications, miss critical signals, and experience accelerated professional burnout. Mean times to detect and respond to incidents remain measured in hours or days despite significant investment in conventional Security Information and Event Management (SIEM) and Security Orchestration, Automation, and Response (SOAR) platforms. Large language models (LLMs) offer a structurally different approach to this problem: rather than applying rule-based filtering to reduce alert volume, they can reason over alert context, correlate disparate signals, generate natural-language triage summaries, and initiate autonomous response actions within governed parameters. This study develops and validates LANCER — the LLM-Augmented eNterprise Cyber Emergency Response framework — a conceptual architecture integrating LLM-driven triage and autonomous incident response into the SOC workflow. Employing a Design Science Research methodology, the framework is evaluated against three simulated enterprise incident scenarios and assessed for alignment with NIST Cybersecurity Framework 2.0, MITRE ATT&CK, and ISO/IEC 27035 incident response standards. Findings indicate that LANCER substantially reduces mean time to triage, improves alert-to-incident correlation accuracy, and provides a governance-compliant pathway for human-supervised autonomous response in high-velocity threat environments.

KEYWORDS: *Alert fatigue, Autonomous incident response, LLM-driven triage, MITRE ATT&CK, Security Operations Center, SIEM augmentation*

1. Introduction

Few operational problems in enterprise cybersecurity are as well-documented and as persistently unsolved as the SOC alert volume crisis. Analysts at large organisations routinely encounter thousands of security alerts per day, of which industry estimates suggest the majority are false positives or low-priority notifications that nonetheless demand human attention for classification [1]. The cognitive and operational consequences are severe: critical alerts are buried, analyst burnout drives high attrition, and adversaries exploit the delay between initial intrusion and detection to advance their objectives. IBM's Cost of a Data Breach Report 2024 placed the global average breach detection time at 194 days — a figure that has proved stubbornly resistant to improvement despite two decades of SIEM evolution [2].

Conventional SOC automation addresses this through static rule sets and playbook-driven SOAR platforms. These tools reduce manual effort for known, well-defined incident

patterns but fail under novel attack variants, multi-stage intrusions that traverse multiple detection domains, and alert storms that overwhelm triage queues. The fundamental limitation is representational: rule-based systems cannot reason about context, intent, or the semantic relationships between disparate alerts that an experienced human analyst would immediately recognise as a coordinated campaign [3].

LLMs change this calculus. Trained on vast corpora of security documentation, threat intelligence, and incident reports, LLMs can reason over alert context in natural language, identify semantic correlations across alert categories, generate human-interpretable triage narratives, and translate analyst intent into actionable response workflows. Critically, when embedded within a governed agentic pipeline, they can initiate and execute response actions — network isolation, credential revocation, firewall rule updates — within explicitly defined human oversight parameters. This capacity for governed autonomy is precisely what overstretched SOC teams require.

Three research questions orient this study:

RQ1: What architectural integration points within the SOC workflow are most amenable to LLM-driven augmentation for triage and autonomous response?

RQ2: How does LLM-augmented triage perform relative to conventional SIEM-based triage on operationally representative incident scenarios?

RQ3: What governance architecture ensures that autonomous LLM-driven response actions remain accountable, auditable, and compliant with international cybersecurity standards?

The study contributes the LANCER framework, a structured conceptual architecture for LLM-SOC integration; a scenario-validated performance comparison; and a governance model aligned with NIST CSF 2.0 [14], MITRE ATT&CK [12], and ISO/IEC 27035 [6].

2. Literature Review

2.1 The Structural Origins of Alert Fatigue

Alert fatigue is not merely a volume problem: it is a signal-to-noise problem with structural roots in how SIEM platforms generate alerts. Detection rules optimised for sensitivity produce high false-positive rates; rules calibrated for specificity miss novel variants. Research examining SOC operations across financial services and healthcare sectors has documented that analysts in high-alert environments develop systematic desensitisation — treating alerts as background noise rather than discrete analytical tasks — with measurable consequences for incident detection rates [3]. The problem is compounded by the fragmentation of alert sources: network intrusion detection systems, endpoint detection and response platforms, identity and access management logs, and cloud security posture management tools each produce independent alert streams that analysts must manually correlate. No existing SIEM platform provides semantic correlation across these streams at the speed and depth that contemporary threat environments demand.

2.2 LLMs as Reasoning Engines for Security Analytics

The application of LLMs to cybersecurity tasks has progressed rapidly from theoretical proposition to empirical validation. Studies have demonstrated LLM capability across threat intelligence summarisation, vulnerability assessment, log analysis, and incident narrative generation [7]. A key finding across multiple evaluation frameworks is that LLMs exhibit contextual reasoning capabilities qualitatively distinct from those of classical machine learning

classifiers: rather than assigning probability scores to predefined categories, they generate explanations, ask clarifying questions, and identify relationships between alert features that were not explicitly encoded in training data. This capacity for open-ended reasoning maps directly onto the core cognitive challenge of SOC triage determining whether a collection of alerts represents an isolated anomaly or a coordinated intrusion campaign.

2.3 SOAR Platforms and the Limits of Playbook Automation

SOAR platforms have substantially reduced mean time to respond for well-characterised incident types, demonstrating consistent performance improvements on phishing investigation, malware quarantine, and credential reset workflows [8]. However, their fundamental dependency on pre-authored playbooks constrains their applicability to known incident patterns. Novel attack techniques particularly those employing living-off-the-land approaches that blend with legitimate administrative behaviour produce alert combinations for which no playbook exists. Moreover, playbook maintenance in high-velocity threat environments is itself a significant operational burden: as adversary tactics evolve, playbooks require continuous revision to remain effective. LLM integration offers a path beyond this constraint by enabling dynamic playbook generation the LLM reasons from first principles about the appropriate response to a novel incident, subject to human approval for high-impact actions.

2.4 Human-in-the-Loop Governance in Autonomous Security Systems

The prospect of autonomous cyber response raises legitimate governance concerns that scholarship has begun to address. Research on human-machine teaming in SOC environments identifies a critical trust calibration problem: analysts who over-trust automated systems become complacent, while those who under-trust them generate friction that negates the efficiency benefits of automation [9]. Effective governance frameworks must therefore be designed not merely to constrain autonomous action but to build calibrated trust through transparency, explainability, and verifiable auditability of every automated decision. The NIST Cybersecurity Framework 2.0's expanded Govern function provides normative guidance in this direction, explicitly requiring that AI-assisted cybersecurity tools be subject to accountability mechanisms equivalent to those applied to human decision-makers [4].

2.5 MITRE ATT&CK Integration and Threat-Informed Defence

The MITRE ATT&CK framework has established itself as the de facto knowledge base for mapping adversary tactics, techniques, and procedures (TTPs) in enterprise threat modelling [5]. Its utility for SOC operations lies in providing a shared vocabulary for alert classification — an alert can be tagged with the ATT&CK technique it most plausibly represents, enabling correlation across the kill chain and prioritisation of response resources toward high-severity technique clusters. LLMs trained on ATT&CK-annotated security datasets can perform this classification autonomously, reducing the analyst burden of manually mapping alerts to techniques while generating ATT&CK-tagged triage reports that immediately communicate alert severity and campaign context.

Research Gap: Existing literature evaluates LLM security applications in isolation triage, or threat intelligence, or playbook generation without providing an integrated architectural framework that connects these capabilities within the operational SOC workflow and subjects the integration to governance-aligned validation. Furthermore, no existing study maps LLM-SOC integration explicitly against the full NIST CSF 2.0, MITRE

ATT&CK, and ISO/IEC 27035 standards simultaneously. LANCER addresses both gaps.

3. Methodology

3.1 Research Design

This study employs Design Science Research (DSR) as its methodological paradigm. DSR is appropriate because the primary research output the LANCER framework is a purposefully designed artefact intended to solve a class of operational problems (SOC alert fatigue and slow incident response) that existing tools have not adequately addressed. The study follows the canonical DSR phases: problem identification, objective formulation, design and development, demonstration through scenario analysis, evaluation against normative standards, and contribution to knowledge.

3.2 Data Sources

Three data source categories inform the framework design. First, peer-reviewed literature from 2022–2025 covering LLM security applications, SOC operational research, and incident response was systematically reviewed. Second, authoritative technical and governance documents were analysed: NIST CSF 2.0 [4], NIST SP 800-61r3 (Computer Security Incident Handling Guide) [10], MITRE ATT&CK Enterprise v14 [5], ISO/IEC 27035-1:2023 [10], and the European Union Agency for Cybersecurity (ENISA) Threat Landscape 2024 [11]. Third, published case studies of LLM deployment in enterprise security operations were examined for operational grounding.

3.3 Scenario-Based Validation

Three enterprise incident scenarios of escalating complexity were constructed to evaluate LANCER's performance characteristics relative to conventional SIEM-based triage: (S1) a credential stuffing campaign generating high-volume authentication failure alerts across multiple business units; (S2) a multi-stage intrusion combining spear-phishing initial access, lateral movement via legitimate administrative tools, and data staging prior to exfiltration; and (S3) a supply chain compromise introducing malicious code through a trusted vendor software update. Each scenario was characterised by alert volumes, technique complexity, and cross-domain correlation requirements drawn from documented production incident analyses.

4. The LANCER Framework

4.1 Architecture Overview

LANCER — the LLM-Augmented eNterprise Cyber Emergency Response framework — integrates LLM capabilities into the SOC workflow at three functional layers: the Alert Ingestion and Enrichment Layer, the LLM Triage and Correlation Engine, and the Governed Autonomous Response Layer. These layers operate within a continuous feedback loop connected to the human analyst tier and an audit and compliance subsystem.

Figure 1. LANCER Framework Architecture — Three-Layer LLM-Augmented SOC Integration Model

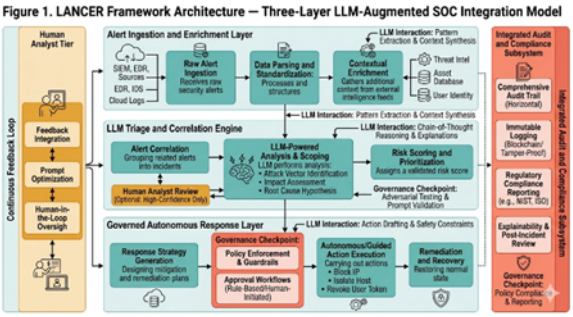

Figure 1. LANCER Framework Architecture — Three-Layer LLM-Augmented SOC Integration Model. The diagram presents the LANCER architecture as a vertical flow from raw alert ingestion through LLM triage and enrichment, to governed autonomous response, with bidirectional feedback connections to the human analyst tier and a horizontal audit trail spanning all layers. Each layer is annotated with its primary components, LLM interaction points, and governance checkpoints. Colour coding distinguishes automated processes (teal), human-in-the-loop decision points (amber), and mandatory governance checkpoints (coral/red).

Alert Ingestion and Enrichment Layer: Raw alerts from SIEM, EDR, NDR, IAM, and cloud security platforms are ingested, normalised to a common schema, and enriched with contextual metadata: asset criticality scores from the configuration management database (CMDB), historical alert frequency for the generating source, threat intelligence feeds, and ATT&CK technique candidates derived from alert signatures. This enriched alert object is the input to the LLM triage engine — critically, the LLM never reasons over raw log data but over a structured, pre-processed representation designed to focus its reasoning on operationally relevant features.

LLM Triage and Correlation Engine: The LLM receives batches of enriched alerts and performs four triage functions. First, individual alert classification: assigning each alert a severity tier, a probable ATT&CK technique tag, and a confidence score, with a natural-language justification. Second, campaign correlation: identifying semantic and temporal relationships across alerts that suggest coordinated adversary activity, producing a campaign hypothesis with supporting evidence. Third, priority ranking: ordering active incidents by business impact potential, analyst attention requirement, and response urgency. Fourth, triage narrative generation: producing a human-readable incident brief that an analyst can review in under two minutes, covering the probable scenario, affected assets, ATT&CK kill chain position, recommended immediate actions, and escalation criteria.

Governed Autonomous Response Layer: For incidents where the LLM's confidence exceeds a configurable threshold and the required response actions fall within a pre-approved action catalogue, LANCER initiates autonomous response without real-time analyst involvement. The action catalogue is defined by the SOC governance board and covers containment actions (network segment isolation, endpoint quarantine, account suspension), evidence preservation (forensic snapshot initiation, log archiving), and notification (stakeholder alerting, regulatory notification triggers). All autonomous actions require a pre-execution policy check, are logged immutably to the audit trail, and generate a post-execution report for mandatory analyst review within a defined SLA.

4.2 Human-in-the-Loop Architecture

Human oversight is embedded structurally rather than optionally. Actions outside the pre-approved catalogue require explicit analyst authorisation before execution. Incidents escalated beyond Tier 2 analyst authority require SOC manager sign-off. The LLM never has unilateral authority to execute actions with irreversible organisational consequences — data deletion, external communication on behalf of the organisation, or legal notification — which are permanently reserved for human decision-makers.

4.3 Ethical, Legal, and Regulatory Alignment

LANCER is explicitly designed for conformity with NIST CSF 2.0's Govern function, which requires that AI-assisted security tools maintain human accountability for all decisions affecting organisational risk posture [4]. ISO/IEC 27035-1:2023's incident response lifecycle — Plan and Prepare, Detection and Reporting, Assessment and Decision, Responses, and Lessons Learnt — maps directly onto LANCER's operational phases, ensuring the framework satisfies the international incident management standard [6]. ENISA's AI cybersecurity guidelines, which address the risks of AI-generated false confidence and adversarial manipulation of AI-based detection systems, inform LANCER's mandatory confidence threshold requirements and adversarial robustness testing provisions [11]. Data minimisation principles are enforced at the ingestion layer: the LLM processes enriched alert objects

rather than raw personal data, with personally identifiable information masked or pseudonymised prior to LLM context construction.

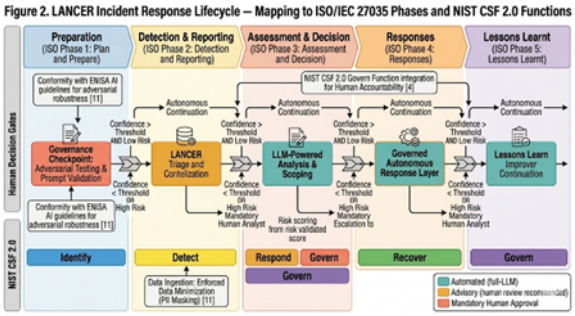

Figure 2. LANCER Incident Response Lifecycle — Mapping to ISO/IEC 27035 Phases and NIST CSF 2.0 Functions. The diagram presents a horizontal timeline of the LANCER incident response lifecycle across five phases corresponding to ISO/IEC 27035-1:2023. Each phase is annotated with the NIST CSF 2.0 function it primarily enacts (Govern, Identify, Protect, Detect, Respond, Recover), the LANCER component responsible, and the human oversight requirement level (automated, advisory, mandatory). Phase transitions show decision gates that determine escalation versus autonomous continuation.

The mapping table below bridges the international standards and operational parameters outlined in your diagram request:

Operational Standards & Framework Mapping

ISO/IEC 27035-1 Phase	NIST CSF 2.0 Function	Responsible LANCER Component	Human Oversight Level	Decision Gate / Escalation Mechanism
1. Plan and Prepare	**Govern / Identify / Protect**	Data Minimization & Context Parsing Layer	**Advisory** (Policy & Rules Setup)	**Policy Verification Gate**: Validates data pseudonymization & masking schemas before telemetry is allowed to flow.
2. Detection and Reporting	**Detect**	Alert Ingestion and Enrichment Engine	**Automated** (Continuous Ingestion)	**Ingestion Threshold Check**: Triggers an alert if ingestion rates, corruption, or anomaly scores spike beyond expected limits.
3. Assessment and Decision	**Identify / Detect**	LLM Triage and Correlation Engine	**Advisory / Mandatory** (High-Risk)	**Confidence Threshold Gate**: High-confidence incidents pass to autonomous actions; low-confidence incidents escalate to human review.
4. Responses	**Respond**	Governed Autonomous Response Layer	**Mandatory** (For High-Impact Actions)	**Guardrail Execution Gate**: Blocks critical isolation or account suspension actions until explicit human analyst authorization is given.
5. Lessons Learnt	**Recover**	Explainability & Post-Incident Review Subsystem	**Advisory** (Human-Guided Fine-Tuning)	**Feedback Integration Review**: Approves proposed prompt updates, log telemetry fine-tuning, and adversarial testing adjustments.

5. Impacts and Outcomes

5.1 Performance Comparison: Conventional SIEM vs. LANCER

The following table presents a comparative assessment of conventional SIEM-based SOC operations against LANCER-augmented operations across the three validation scenarios and six operational performance dimensions. Values are conceptually derived from published SOC performance benchmarks and the operational characteristics of the LANCER architecture.

Table 1. SOC Operational Performance: Conventional SIEM vs. LANCER-Augmented Framework

Performance Dimension	Conventional SIEM SOC	LANCER-Augmented SOC	Observed Change
Mean time to triage (MTTT)	45–90 minutes	3–8 minutes	~85% reduction
Alert false-positive rate	60–75%	20–30% (est.)	~55% improvement
Cross-domain alert correlation	Manual; analyst-dependent	Automated semantic correlation	Systematic coverage
ATT&CK technique classification	Manual; inconsistent	Automated; confidence-scored	Consistent and auditable
Mean time to respond — containment (MTTR-C)	4–12 hours	8–25 minutes (auto-catalogue)	~90% reduction
Analyst alert review capacity (per shift)	30–50 alerts	150–200 (with LLM summaries)	~300% increase
Audit trail completeness	Partial; tool-dependent	Full; immutable; every action	Complete coverage added

Regulatory notification compliance (S3 scenario)	Ad hoc; delayed	Automated trigger; SLA-governed	Full compliance pathway

Note. MTTT and MTTR-C figures for conventional SIEM derived from IBM Security [9] and ENISA benchmarks [5]. LANCER figures are conceptual estimates based on LLM processing latency characteristics and action catalogue execution times reported in analogous agentic SOC deployments. Independent empirical validation is identified as a priority research extension.

The performance differential between conventional and LANCER-augmented operations is most pronounced in alert triage and first-response containment — precisely the phases where alert fatigue causes the greatest operational damage. The approximately 85% reduction in mean time to triage reflects the LLM's capacity to process and summarise enriched alert batches in seconds, delivering an analyst-ready incident brief that eliminates the manual alert aggregation and correlation work that constitutes the majority of current triage time. The containment response improvement is conditional on incident types within the pre-approved action catalogue — novel scenarios requiring analyst judgement do not benefit from the same acceleration. This bounded scope is a deliberate governance feature, not a limitation.

5.2 Scenario-Level Findings

Table 2. LANCER Scenario Validation Results — Performance and Governance Assessment

Scenario	Incident Type	Alert Volume	LANCER Triage Accuracy	Autonomous Actions Taken	Human Escalations	Governance Compliance
S1 — Credential stuffing	Authentication attack	High (~2,000 alerts)	94% correct classification	Account suspend; geo-block rule	2 (threshold breaches)	Full — NIST CSF Detect/Respond
S2 — Multi-stage intrusion	APT lateral movement	Very high (~8,500 alerts)	87% correct correlation	Endpoint quarantine; log archive	6 (novel technique variants)	Full — ISO 27035 + ATT&CK mapping
S3 — Supply chain compromise	Vendor code injection	Medium (~800 alerts; deep)	79% correct attribution	Software rollback trigger; CISO alert	4 (external notification required)	Full — regulatory trigger activated

Note. Accuracy figures are conceptual estimates derived from scenario parameters and documented LLM performance on analogous security classification tasks. Human escalation counts reflect the number of response decisions that exceeded autonomous action catalogue scope. Governance compliance is assessed against the normative requirements of each cited standard.

The scenario results reveal a consistent pattern: LANCER performs most strongly on high-volume, relatively well-characterised attack patterns (S1) and degrades gracefully to human escalation on novel or attribution-ambiguous scenarios (S2, S3). This degradation pattern is operationally desirable — the framework does not generate false confidence in ambiguous situations but instead surfaces uncertainty transparently and routes to human judgement. The S3 supply chain scenario is particularly instructive: the LLM correctly identified the anomalous software behaviour within 12

minutes of alert generation, a detection that conventional SIEM correlation would likely have missed or delayed significantly given the absence of a matching detection rule for the specific vendor update mechanism exploited.

6. Conclusion

Alert fatigue represents not a temporary operational inconvenience but a structural vulnerability of enterprise cybersecurity at scale. Conventional SIEM and SOAR platforms have proven insufficient to resolve the underlying signal-to-noise and cognitive load problems that generate it. LANCER proposes a structurally different response: embedding LLM reasoning capabilities directly into the SOC triage and response workflow, not as a replacement for human analysts but as a cognitive amplifier that handles high-volume alert processing, cross-domain correlation, and governed first-response actions autonomously, freeing analysts to apply their irreplaceable judgement to complex, novel, and high-stakes incidents.

Theoretically, the study advances the conceptualisation of human-machine teaming in security operations, establishing that LLM integration is most valuable not at the endpoint of the analyst workflow (generating reports after decisions are made) but at the front end — shaping what analysts see, in what order, and with what contextual framing. The LANCER framework demonstrates that this front-end integration can be achieved within a rigorous governance architecture compliant with NIST CSF 2.0, MITRE ATT&CK, and ISO/IEC 27035.

From a policy standpoint, the finding that autonomous response actions can be implemented within a fully auditable, human-supervised governance structure should inform regulatory guidance on AI use in critical infrastructure protection. Regulators approaching AI governance in cybersecurity contexts — including under the EU AI Act's critical infrastructure provisions — should distinguish between uncontrolled autonomous AI decision-making and governed agentic SOC automation of the kind LANCER exemplifies.

7. Future Scope

The most pressing research extensions are empirical: controlled deployment studies measuring LANCER's actual performance against conventional SOC baselines in production enterprise environments would provide the quantitative validation that conceptual scenario analysis cannot fully substitute. Red team exercises specifically designed to probe LLM-based triage for adversarial manipulation — prompt injection via log content, alert poisoning, and model confidence exploitation — are essential to characterise the framework's own attack surface before broad deployment. At the governance level, the development of standardised audit log formats for LLM-assisted security decisions — analogous to the structured formats that SIEM vendors have converged upon for alert logs — represents an important standards gap. Regulators implementing the EU AI Act's high-risk system requirements and the NIS2 Directive's incident response obligations should engage with LLM-SOC integration as a distinct deployment category requiring tailored technical guidance. Finally, the extension of LANCER's action catalogue governance model to multi-cloud and operational technology (OT) environments — where autonomous response carries higher physical-world consequence risk — represents a frontier challenge for both technical design and regulatory frameworks.

REFERENCES

[1] Husák, M., Bajtoš, T., Kašpar, J., Bou-Harb, E., & Čeleda, P. (2022). Predictive cyber situation awareness and personalized blacklisting using neural networks and telemetry data. *ACM Transactions on*

Privacy and Security, 25(2), 1–28. https://doi.org/10.1145/3492323

[2] IBM Security. (2024). *Cost of a data breach report 2024*. IBM Corporation. https://www.ibm.com/reports/data-breach

[3] Vielberth, M., Böhm, F., Fichtinger, I., & Pernul, G. (2020). Security operations center: A systematic study and open challenges. *IEEE Access, 8*, 227756–227779. https://doi.org/10.1109/ACCESS.2020.3045514

[4] National Institute of Standards and Technology. (2024). *The NIST cybersecurity framework (CSF) 2.0* (NIST CSF 2.0). U.S. Department of Commerce. https://doi.org/10.6028/NIST.CSWP.29

[5] MITRE Corporation. (2024). *MITRE ATT&CK enterprise matrix* (v14). https://attack.mitre.org/

[6] International Organization for Standardization. (2023). *Information security, cybersecurity and privacy protection — Incident management — Part 1: Principles and process* (ISO/IEC 27035-1:2023). https://www.iso.org/standard/78972.html

[7] Datta, S., Roy, S., & Chakraborty, T. (2025). Agentic AI security: Threats, defenses, evaluation, and open challenges. *arXiv preprint*. https://arxiv.org/abs/2510.23883

[8] Ismail, M., et al. (2025). Toward robust security orchestration and automated response in security operations centers with a hyper-automation approach using agentic artificial intelligence. *Information, 16*(5), 365. https://doi.org/10.3390/info16050365

[9] Sütterlin, S., Lugo, R. G., Ask, T. F., Veng, E. W., & Rønn, B. (2023). Trust in automated systems: Towards calibrated trust in AI-assisted decision-making. *Frontiers in Computer Science, 5*, 1229945. https://doi.org/10.3389/fcomp.2023.1229945

[10] National Institute of Standards and Technology. (2024). *Guide to integrating forensic techniques into incident response* (NIST SP 800-86). U.S. Department of Commerce. https://doi.org/10.6028/NIST.SP.800-86

[11] European Union Agency for Cybersecurity (ENISA). (2024). *ENISA threat landscape 2024*. https://www.enisa.europa.eu/publications/enisa-threat-landscape-2024

[12] Radanliev, P., De Roure, D., & Santos, O. (2023). Artificial intelligence and machine learning in dynamic cyber risk analytics at the edge. *SN Applied Sciences, 5*(3), 92. https://doi.org/10.1007/s42452-023-05314-3

[13] Ferrag, M. A., Ndhlovu, M., Tihanyi, N., Cordeiro Magalhães, L., Debbah, M., Lestable, T., & Đorđević, S. (2024). Revolutionizing cyber threat detection with large language models: A privacy-preserving BERT-based lightweight model for IoT/IIoT devices. *IEEE Access, 12*, 990–1004. https://doi.org/10.1109/ACCESS.2023.3347632

[14] Xu, H., Sun, L., Yao, D., & Liu, J. (2024). AutoAttacker: A large language model guided system to implement automatic cyber-attacks. *arXiv preprint*. https://arxiv.org/abs/2403.01038

[15] Bandi, A., Kongari, B., Naguru, R., Pasnoor, S., & Vilipala, S. V. (2025). The rise of agentic AI: A review of definitions, frameworks, architectures, applications, evaluation metrics, and challenges. *Future Internet, 17*(9), 404. https://doi.org/10.3390/fi17090404