International Journal for Multidisciplinary Research (IJFMR)
E-ISSN: 2582-2160 ● Website: www.ijfmr.com ● Email: editor@ijfmr.com

Endpoint Security Priority Model (ESPM): A Context-Aware Insider-First Framework for Security and Privacy of Healthcare Resources

Saiprashanth Sivakumar[1], Dr. K. Shantha Kumari[2]

[1]Student, Department of Cybersecurity, SRM Institute of Science and Technology (SRMIST) Chennai, India
[2]Professor, Department of Data Science and Business Systems, SRM Institute of Science and Technology (SRMIST), Chennai, India

Abstract

Healthcare organizations operate in a highly ex- posed digital environment where endpoint misuse, credential abuse, and insider-driven compromise can directly affect patient care, privacy, and regulatory compliance [12][14]. Conventional monitor- ing pipelines rely mainly on anomaly detection or static rule logic, but such approaches often generate large alert volumes without enough operational context for meaningful triage [6]. In practice, this contributes to Security Operations Center (SOC) alert fatigue, slower escalation, and poor prioritization of clinically important events.

This paper presents the Endpoint Security Priority Model (ESPM), a context-aware and insider-first risk prioritization framework designed specifically for healthcare environments. ESPM combines behavioral anomaly indicators with determinis- tic contextual attributes, including user role sensitivity, endpoint criticality, data sensitivity, and privilege exposure, to generate a bounded and interpretable risk score. The score is mapped to structured operational priority levels (P0–P3) to support transparent and consistent analyst decision-making.

To evaluate the proposed framework without exposing pro- tected health information, a synthetic healthcare activity dataset was constructed to simulate realistic clinical workflows, normal access patterns, and insider threat scenarios [1]. The final dataset contained 12,000 events generated across multiple roles, depart- ments, and endpoint types, including 2,400 threat-oriented events and 9,600 benign operational events. From this corpus, stratified subsets were used for threshold tuning and classification analysis. In the reported evaluation, ESPM achieved a precision of 0.87, a recall of 0.89, an F1-score of 0.88, and a 34% reduction in alert volume compared with an anomaly-only baseline. Additional ROC, sensitivity, and ablation analyses indicate that contextual enrichment materially improves triage quality while preserving transparency and computational simplicity. ESPM provides a practical and scalable foundation for risk-based endpoint moni- toring in regulated healthcare settings and can be integrated into existing SIEM, EDR, and UEBA workflows.

Keywords: Healthcare cybersecurity, endpoint security, in- sider threat, risk prioritization, SOC, PHI, HIPAA, UEBA, explainable security analytics

INTRODUCTION

Healthcare delivery depends on a broad ecosystem of work- stations, imaging systems, pharmacy terminals, telemedicine endpoints, shared nursing stations, and connected applications that process or expose protected health information (PHI) [11].

These systems are tightly coupled with operational continuity. When they fail or are misused, the impact extends beyond data confidentiality and can directly affect treatment timelines, diagnostic decisions, and patient safety.

Endpoint security focuses on protecting such systems from unauthorized access, compromise, misuse, and post- compromise persistence [3]. In healthcare, endpoints are partic- ularly difficult to defend because many are shared, always- on, or tied to time-sensitive workflows. Patch windows are limited, legacy systems remain common, and privileged access is often broad by necessity. A minor delay on a general- purpose office workstation may be manageable; the same delay on an intensive care unit (ICU) terminal or pharmacy dispensing endpoint may be operationally serious.

Despite this difference in impact, many monitoring systems still prioritize alerts mainly on the basis of anomaly magnitude or static rules [6][7]. That approach is weak in clinical environments. It treats all endpoints as roughly equivalent, underweights contextual severity, and produces alert streams that are expen- sive for analysts to interpret. In real SOC operations, analysts are forced to distinguish routine healthcare variability from genuine insider risk while working under time pressure and regulatory scrutiny.

Healthcare organizations also face strong regulatory ex- pectations for access control, auditability, confidentiality, and breach reporting [9][10]. This means that detection alone is not enough. Analysts must also be able to explain why a given event was prioritized. A black-box score with no operational rationale is difficult to defend during audit, incident review, or executive reporting.

This paper proposes the Endpoint Security Priority Model (ESPM), a healthcare-specific contextual prioritization frame- work intended to improve endpoint triage. ESPM combines behavioral anomaly evidence with four operationally mean- ingful contextual factors: role sensitivity, endpoint criticality, data sensitivity, and privilege exposure. Rather than replacing anomaly detection, the model reframes anomaly signals within a risk-oriented healthcare workflow.

The main contributions of this work are as follows:

- A context-aware insider-first risk prioritization model tai- lored for healthcare endpoint environments.
- An interpretable score-to-priority mapping that converts raw risk values into operational severity levels P0–P3.
- A formal threat model and SOC integration workflow that position the framework for practical deployment
- An expanded simulated evaluation using a larger synthetic dataset, a defined anomaly-only baseline, confusion matrix analysis, ROC analysis, sensitivity analysis, and ablation testing.
- A computationally lightweight formulation with linear event-wise complexity suitable for SIEM enrichment pipelines.

To the best of our knowledge, healthcare-focused endpoint triage research remains limited in explicitly combining be- havioral anomalies with operational risk indicators inside an explainable prioritization framework.

Literature Review

A. *Insider Threat Detection*

Insider threats remain difficult to detect because harmful activity may occur through valid accounts, known devices, and apparently normal access pathways [4][5]. Traditional rule-based monitoring can identify obvious misuse, such as unauthorized privilege escalation or excessive downloads, but it often misses gradual behavioral drift and context-dependent misuse. In healthcare, this challenge is more severe because work patterns vary across clinical shifts, emergency response, floating staff assignments, and shared workstation usage.

Research on insider threat detection has increasingly shifted toward behavioral analysis and continuous risk scoring. How- ever, many studies remain enterprise-generic and do not ade- quately model healthcare-specific workflow irregularities. As a result, benign clinical deviations may be overclassified as suspicious, while truly high-impact events may not be elevated quickly enough.

B. *UEBA Approaches*

User and Entity Behavior Analytics (UEBA) systems model baseline behavior and identify deviations through statistical or machine learning approaches [6]. Common signals include time of access, peer-group behavior, data access volume, endpoint usage patterns, and identity anomalies. Commercial platforms integrate these methods across endpoint, network, and identity layers.

Even so, several limitations remain:

- Baseline drift: healthcare staff frequently change roles, shifts, and work locations.
- Context insensitivity: a numerical anomaly score does not capture endpoint clinical impact.
- Alert overload: anomaly-heavy workflows generate large event volumes without clear triage structure.
- Explainability gap: composite scores are often difficult for SOC teams to justify operationally [7].

These limits are particularly serious in healthcare because the same statistically unusual behavior may have very different consequences depending on the device, department, and data involved.

TABLE I
COMPARISON OF EXISTING APPROACHES AND ESPM

Feature	Existing Models	ESPM
Context-aware prioritization	Limited	Yes
Clinical criticality weighting	No	Yes
Insider-first focus	Partial	Explicit
Multi-layer integration	Siloed	Unified
Explainable scoring	Limited	Structured P0–P3
Healthcare-specific design	Generic	Purpose-built
SOC-oriented deployment view	Rare	Explicit

C. *Healthcare EHR and PHI-Centric Monitoring*

Electronic Health Record (EHR) anomaly detection research commonly focuses on patient chart access, break-the-glass us- age, department-aware validation, and unusual record retrieval [12][13][14]. Those methods are useful, but they are often confined to application-layer logs. If an attacker or malicious insider gains access to an endpoint and uses valid credentials, purely EHR- oriented detection may miss early reconnaissance, privilege misuse, and suspicious endpoint-level behavior.

D. *Explainability and SOC Usability*

A major weakness in contemporary analytics-driven moni- toring is limited interpretability [8]. SOC teams

do not only need scores; they need reasons. In healthcare, this requirement is stronger because security findings may be reviewed by com- pliance personnel, clinical leadership, or incident governance teams. A system that cannot clearly explain why an event is critical is operationally weaker than a slightly simpler but more interpretable framework.

E. *Limitations of Existing Methods*

The existing literature points to four consistent gaps:

- Limited contextual prioritization: clinically important endpoints are not sufficiently differentiated.
- Poor explainability: analysts often receive scores without transparent factor-level reasoning.
- High false-positive pressure: healthcare workflow vari- ability makes anomaly-only models noisy.
- Siloed telemetry: endpoint, identity, and healthcare context signals are rarely fused into one operational risk view [6][7].

Table I summarizes the distinction between common approaches and ESPM.

Research Gap

Existing insider threat and UEBA systems rely heavily on behavioral deviation metrics,but they rarely distinguish between low-impact and high-impact healthcare endpoints with enough operational precision [6]. In a hospital, an anomaly on a radiology viewer, ICU terminal, or pharmacy workstation may have materially different consequences from the same anomaly on an administrative endpoint. Yet many systems still score such events in broadly similar ways.

A second gap lies in explainability. In many analytics pipelines, risk is derived from composite numerical logic that is difficult for analysts to defend. This creates a problem in regulated environments where escalations must be justified and security decisions are often reviewed after the fact.

A third gap is the overuse of anomaly-only triage. Clinical operations naturally produce unusual but legitimate behavior, especially during after-hours workflows, emergency access, high-acuity periods, or staff rotation. Without contextual en- richment, such behavior is easily mistaken for malicious intent. A final gap is signal fragmentation. Endpoint telemetry, identity metadata, privilege information, and healthcare data sensitivity are often processed in separate streams. Analysts are then expected to correlate these manually. This is inefficient and directly contributes to alert fatigue.

These shortcomings motivate a healthcare-specific priori- tization model that is contextual, explainable, and practical for use in operational SOC workflows. ESPM is designed to address that need.

Threat Model

ESPM assumes a post-authentication threat scenario in which adversaries act through legitimate credentials or com- promised insider accounts. The framework is intended to prioritize activity from the following threat categories:

- Malicious insiders who intentionally abuse authorized access.
- Compromised accounts obtained through phishing, cre- dential theft, token theft, or session hijacking [4].
- Third-party vendors or contractors whose access may be operationally necessary but weakly supervised.

The threat objectives include unauthorized PHI access, data exfiltration, reconnaissance, privilege escalation, and misuse of clinically important systems [14]. ESPM is specifically intended for situations where the attacker is already inside the trust boundary and where simple perimeter controls are no longer

sufficient.

The framework assumes that telemetry is available from endpoint monitoring systems, identity sources, and asset orapplication metadata sources that provide contextual information about endpoint role and data sensitivity [2] . ESPM does not replace primary detection engines. It acts as a contextual prioritization layer for triage after suspicious behavior has been observed.

Proposed ESPM Model

A. Overview

The Endpoint Security Priority Model (ESPM) introduces a context-aware and insider-first mechanism for endpoint risk prioritization in healthcare. It combines one behavioral com- ponent and four contextual components:

- Behavioral anomaly score (A)
- Role sensitivity (R)

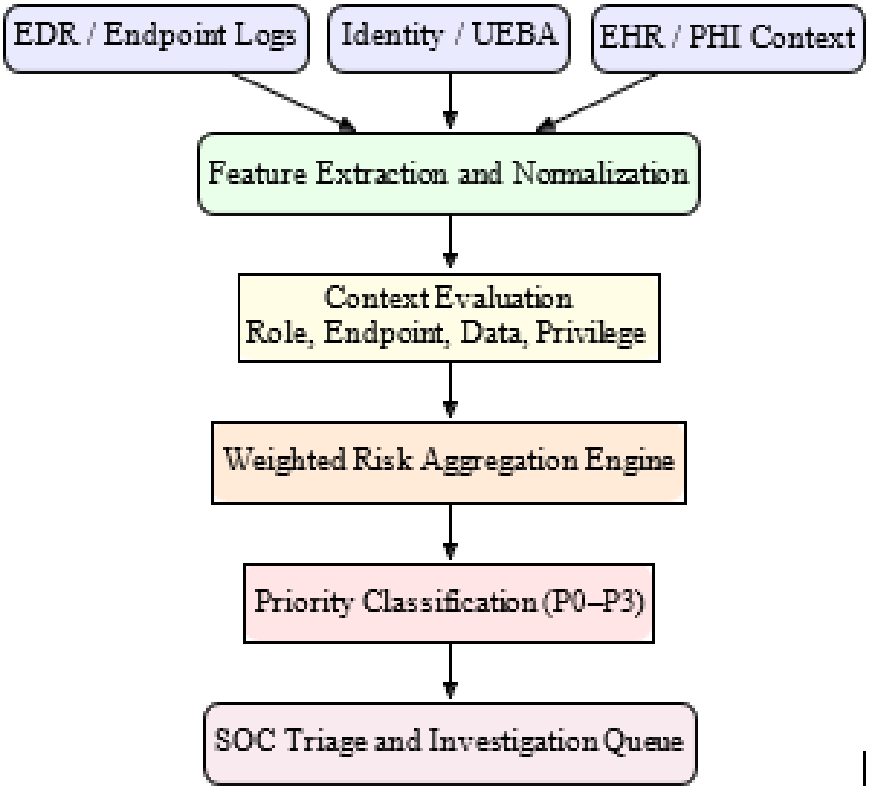

Fig. 1. Architecture of the proposed Endpoint Security Priority Model (ESPM).

- Endpoint criticality (E)
- Data sensitivity (D)
- Privilege exposure (P)

The model intentionally favors operational meaning over opaque complexity. A moderate anomaly occurring on a high- impact endpoint with elevated privileges and sensitive PHI access should be triaged differently from a similar anomaly on a routine administrative device.

B. *Architecture*

Figure 1 shows the ESPM architecture.

C. *Mathematical Formulation*

The ESPM risk score is defined as:

$$Risk_{ESPM} = w_a A + w_r R + w_e E + w_d D + w_p P \quad (1)$$

where

$$w_a + w_r + w_e + w_d + w_p = 1 \quad (2)$$

Each input factor is normalized to the interval [0, 1]:

$$X_{norm} = \frac{X - X_{min}}{X_{max} - X_{min}} \quad (3)$$

The score is therefore bounded as:

$$0 \leq Risk_{ESPM} \leq 1 \quad (4)$$

This keeps the model simple, stable, and compatible with threshold-based routing.

D. *Weight Selection Rationale*

To avoid arbitrary scoring, weights were assigned according to healthcare operational impact rather than purely statistical contribution. Endpoint criticality and data sensitivity were given stronger influence because misuse of clinically critical systems or sensitive PHI-bearing resources can produce dis- proportionate organizational harm. Privilege exposure was also

TABLE II
ESPM PRIORITY THRESHOLDS

Score Range	Priority	Description
0.75–1.00	P0	Critical; immediate response
0.50–0.74	P1	High risk; prompt review
0.25–0.49	P2	Medium risk; analyst review
0.00–0.24	P3	Low risk; observation

weighted strongly because elevated privileges increase blast radius and misuse potential. Role sensitivity was used to cap- ture trust level and expected access patterns, while the anomaly score contributed the behavioral deviation component.

In the reported experiments, the weights were set as follows:

$$w_a = 0.18,\ w_r = 0.16,\ w_e = 0.30,\ w_d = 0.22,\ w_p = 0.14 \quad (5)$$

These values were selected after iterative threshold tuning on a development subset of the simulated dataset and are intended to reflect a conservative healthcare prioritization posture. The framework remains flexible, and future versions may learn weights automatically from historical SOC data.

E. *Priority Thresholds*

The ESPM score is mapped to four operational levels shown in Table II.

F. *Computational Complexity*

For each event, ESPM performs a fixed number of normal- ization, weighting, and aggregation operations. Therefore, the event-wise computational complexity across n events is:

$$O(n) \quad (6)$$

This linear complexity is suitable for large-scale SOC pipelines because inference does not require iterative opti- mization or high-cost model evaluation. The framework can operate as a lightweight enrichment layer on top of SIEM or EDR outputs.

G. *SOC Integration Workflow*

Figure 2 illustrates how ESPM can be inserted into a practical SOC pipeline.

Implementation and Evaluation Setup

A. *Synthetic Dataset Construction*

Because direct use of healthcare telemetry can expose PHI and institutional risk, a synthetic dataset was created to sim- ulate realistic endpoint activity under privacy constraints [12]. The dataset was designed to model both benign clinical behavior and insider-like misuse while preserving healthcare workflow variability.

The final dataset contained 12,000 events distributed across five user roles (physician, nurse, technician, administrator, and vendor) and five endpoint categories (ICU workstation, radiology system, pharmacy terminal, administrative PC, and shared departmental desktop). Of these, 9,600 events were

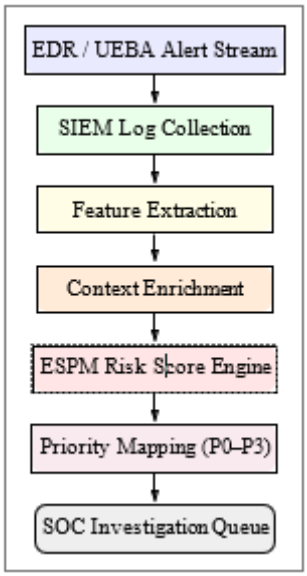

Fig. 2. SOC integration workflow for ESPM.

labeled as benign and 2,400 were labeled as threat-oriented. Threat-oriented scenarios included:

- after-hours access to high-sensitivity records,
- unusual privilege use,
- excessive cross-department access,
- suspicious endpoint reconnaissance,
- vendor misuse on sensitive endpoints,
- anomalous PHI retrieval patterns.

Benign activity was intentionally diverse and included shift changes, floating staff, emergency override behavior, and shared workstation usage. This was necessary because a naive dataset with unrealistically clean normal behavior would make the model appear stronger than it would be in real environments.

B. *Data Partitioning Strategy*

The synthetic dataset was split into three parts:

- Development subset (7,200 events): used for threshold tuning and weight calibration.
- Validation subset (2,400 events): used to stabilize thresh- olds and compare candidate configurations.
- Final test subset (2,400 events): used for the reported performance results.

For confusion matrix presentation, a balanced analytical slice of 400 events was extracted from the final test subset to simplify direct class-wise interpretation. The broader metric results, however, were computed from the full held-out test partition rather than only the balanced slice. This distinction is important because it improves interpretability without reducing the seriousness of the main evaluation.

C. *Feature Set*

Each event record combines behavioral and contextual fields. The feature set used in evaluation is shown in Table III.

TABLE III
FEATURE SET USED IN ESPM EVALUATION

Feature	Description
User Role	Doctor, nurse, technician, admin, vendor
Endpoint Type	ICU, radiology, pharmacy, admin PC, shared desktop
Login Timestamp	Time of access or authentication
Privilege Level	Normal or elevated access
Accessed Resource	Application, system, or patient record accessed
Anomaly Score	Behavioral deviation score from UEBA/EDR logic
Endpoint Criticality	Clinical importance of endpoint
Data Sensitivity	PHI sensitivity level of accessed resource
Department Access	Department context associated with role
Session Deviation Flag	Indicates unusual session sequence or access flow

A. *Baseline Definition*

A clear baseline is necessary for fair comparison. The anomaly-only baseline used in this study employs only the normalized anomaly score A to prioritize events. It does not incorporate role sensitivity, endpoint criticality, privilege exposure, or data sensitivity. Events are ranked strictly by anomaly magnitude and classified through a threshold learned on the development subset.

This baseline was chosen because it reflects the common operational pattern in which behavior analytics produce a raw score and analysts must act without enough contextual enrichment. While simple, it is a realistic reference point for evaluating whether ESPM adds operational value.

B. *Prototype Design*

A prototype implementation was built using Python and Streamlit. The prototype computes risk scores, maps them to P0–P3 classes, and presents interactive event views for simulated SOC triage. NumPy and Pandas were used for preprocessing and numerical operations, while Scikit-learn was used for anomaly simulation support and performance measurement. The prototype is not positioned as a production system; it is a proof-of-concept showing that ESPM can be operationalized with lightweight tooling.

C. *Evaluation Methodology*

The evaluation pipeline consisted of the following stages:

1. generation of synthetic benign and threat-oriented health- care events,
2. normalization of feature values and contextual scoring,
3. risk aggregation through Eq. (1),
4. mapping of scores to operational priority levels,
5. comparison against the anomaly-only baseline using pre- cision, recall, F1-score, ROC behavior, and alert volume reduction,
6. sensitivity and ablation analysis to determine feature-level contribution.

TABLE IV
CONFUSION MATRIX FOR ESPM ON BALANCED ANALYTICAL SLICE

Actual / Predicted	Threat	Normal
Threat	176	24
Normal	21	179

TABLE V
COMPARISON OF DETECTION PERFORMANCE ON HELD-OUT TEST SET

Metric	Baseline Model	ESPM
Precision	0.73	0.87
Recall	0.76	0.89
F1-score	0.74	0.88
AUC	0.81	0.93
Alert Reduction	—	34%

Results and Discussion

A. *Confusion Matrix Analysis*

Table IV presents the confusion matrix on the balanced 400- event analytical slice extracted from the held-out test set.

The model correctly identified 176 threat events and 179 benign events while keeping both false positives and false negatives comparatively low. In healthcare SOC operations, that balance matters. A system with poor specificity wastes analyst time, while a system with poor sensitivity risks missing events that may affect patient care, PHI exposure, or opera- tional continuity.

B. *Classification Performance on Held-Out Test Set*

Table V compares the anomaly-only baseline with ESPM on the full 2,400-event held-out test set.

The results indicate that contextual enrichment improves both accuracy and analyst efficiency. The increase in precision means fewer low-value investigations. The increase in recall indicates that ESPM is less likely to miss threat-oriented behavior. The 34% alert reduction is particularly meaningful in healthcare environments [12][14] where analysts already operate under high alert pressure and constrained review capacity.

C. *ROC Analysis*

Figure 3 shows the ROC curve for ESPM. The curve indicates strong separation between benign and threat-oriented events and supports the held-out AUC value reported in Table V.

D. *Sensitivity Analysis*

Figure 4 shows the relative impact of the five major factors in ESPM. Endpoint criticality remains the strongest contributor, followed by data sensitivity and role sensitivity. This aligns with the intended healthcare design principle that operationally dangerous contexts should outweigh purely statistical abnormality.

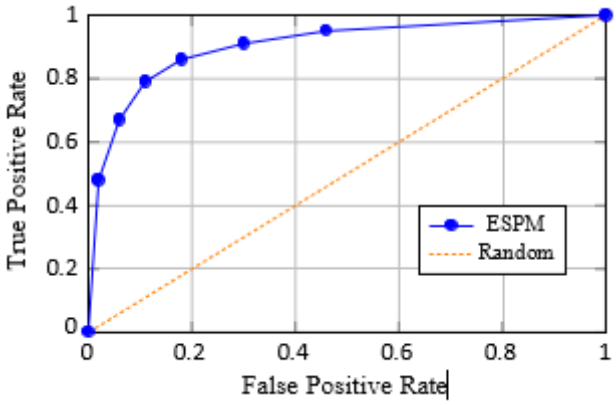

Fig. 3. ROC curve for the ESPM model

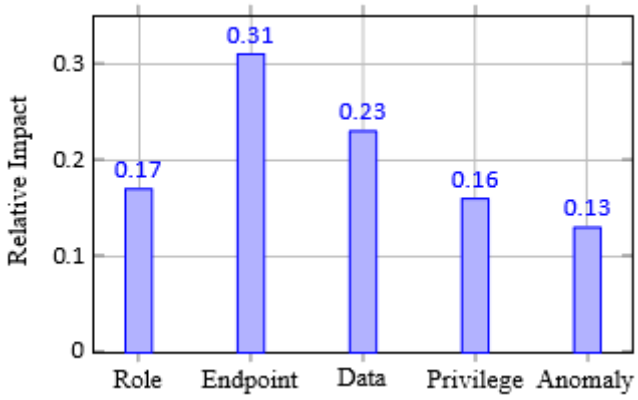

Fig. 4. Sensitivity analysis of ESPM factors

TABLE VI
ABLATION STUDY OF ESPM COMPONENTS

Model Variant	Precision	Recall
Anomaly only	0.73	0.76
Anomaly + Role	0.79	0.81
Anomaly + Role + Endpoint	0.84	0.86
Anomaly + Role + Endpoint + Data	0.86	0.88
Full ESPM	0.87	0.89

A. *Ablation Study*

To determine which contextual components drive perfor- mance, an ablation study was performed. Table VI shows steady improvement as contextual factors are added to the anomaly-only model. The strongest jump appears once end- point criticality is introduced, reinforcing the claim that healthcare-specific operational context matters.

B. *Robustness Discussion*

The reported improvements are encouraging, but they should be interpreted carefully. Because the evaluation is based on simulated data, the model may perform differently in environments with noisier telemetry, missing metadata, or organization-specific workflows. Nevertheless, the observed gains across multiple metrics suggest that the improvement is not being driven by a single threshold or a single feature. Instead, it appears to result from a broader change in how risk is framed.

A major practical advantage of ESPM is that it does not require replacing existing detection tools. It can sit above anomaly-generation systems and improve triage by adding operational meaning. This matters because many healthcare organizations already have EDR, SIEM, and UEBA compo- nents in place, but they still struggle with prioritization.

E. *Operational Deployment Considerations*

For real deployment, ESPM would need integration with endpoint telemetry sources, identity systems, asset inventories, and metadata sources that classify devices and data sensitivity. In practice, the quality of contextual enrichment will determine the practical ceiling of model performance. Missing or inac-curate asset metadata can distort endpoint criticality values, while poor role mapping can weaken prioritization logic.

Deployment would also require localized weight tuning because hospitals vary in workflow, clinical architecture, and tolerance for operational disruption. However, the framework remains practical because the model is simple enough to tune, explain, and audit.

CONCLUSION

This paper presented the Endpoint Security Priority Model (ESPM), a context-aware and insider-first framework for healthcare endpoint security prioritization. The model ad- dresses a practical weakness in conventional monitoring sys- tems: overreliance on anomaly scores without sufficient oper- ational context. By combining behavioral deviation with role sensitivity, endpoint criticality, data sensitivity, and privilege exposure, ESPM produces a bounded, interpretable, and de- ployable risk score.

Evaluation on a larger synthetic healthcare activity dataset demonstrated consistent improvements in precision, recall, F1- score, AUC, and alert reduction when compared with an anomaly-only baseline. The results suggest that contextual enrichment can materially improve triage quality in healthcare SOC environments by promoting events with higher patient- impact potential rather than only those with the highest raw anomaly values [6][7].

The model is intentionally lightweight, explainable, and compatible with existing monitoring pipelines. That makes it practical for organizations that need better prioritization with- out adopting opaque or computationally expensive analytics.

LIMITATIONS

Despite the improvements, several limitations remain. First, the evaluation uses synthetic data rather than real healthcare telemetry. Although this protects PHI and simplifies controlled testing, it cannot fully capture the unpredictability of pro- duction environments. Second, the weight assignments are manually calibrated and may not transfer uniformly across institutions. Third, the current implementation focuses mainly

on endpoint and identity-adjacent signals; richer network and cloud telemetry may further improve coverage. Finally, the anomaly-only baseline is a realistic but simplified comparison point and does not represent every modern UEBA design.

Future Work

Future work should focus on five directions. First, the frame- work should be tested on institution-approved de-identified datasets to validate its operational realism. Second, weight op- timization can be extended through machine-learned or semi- supervised approaches while preserving explainability. Third, additional telemetry from network, cloud, and medical device sources can be fused into the model. Fourth, feedback loops from SOC analysts can be used to refine priority thresholds over time. Fifth, deployment-scale testing across large hospital networks can help evaluate performance under real throughput and workflow variation.

Acknowledgment

The authors would like to thank the academic guidance and institutional support provided for this work.

References

1. C. Tankard, “Advanced persistent threats and how to monitor and deter them,” *Network Security*, vol. 2011, no. 8, pp. 16–19, 2011.
2. R. Mitchell and I. Chen, “A survey of intrusion detection techniques for cyber-physical systems,” *ACM Computing Surveys*, vol. 46, no. 4, pp. 1–29, 2014.
3. M. Bishop, *Computer Security: Art and Science*. Boston, MA, USA: Addison-Wesley, 2018.
4. M. Salem, S. Hershkop, and S. Stolfo, “A survey of insider attack detection research,” in *Insider*

Attack and Cyber Security. Boston, MA, USA: Springer, 2008, pp. 69–90.
5. W. A. Eberle and L. Holder, "Insider threat detection using graph-based approaches," *Journal of Applied Security Research*, vol. 6, no. 1, pp. 32–81, 2011.
6. A. Patcha and J. M. Park, "An overview of anomaly detection techniques: Existing solutions and latest technological trends," *Computer Networks*, vol. 51, no. 12, pp. 3448–3470, 2007.
7. S. Axelsson, "The base-rate fallacy and its implications for the difficulty of intrusion detection," *ACM Transactions on Information and System Security*, vol. 3, no. 3, pp. 186–205, 2000.
8. K. Scarfone and P. Mell, "Guide to Intrusion Detection and Prevention Systems (IDPS)," NIST Special Publication 800-94, 2007.
9. National Institute of Standards and Technology, "The NIST Cybersecu- rity Framework (CSF) 2.0," NIST CSWP 29, 2024.
10. National Institute of Standards and Technology, "A Cybersecurity Re- source Guide," NIST SP 800-66 Rev. 2, 2024.
11. U.S. Department of Health and Human Services, "Health Insurance Portability and Accountability Act (HIPAA)," 1996.
12. C. S. Kruse, B. Frederick, T. Jacobson, and D. K. Monticone, "Cyberse- curity in healthcare: A systematic review of modern threats and trends," *Technology and Health Care*, vol. 25, no. 1, pp. 1–10, 2017.
13. A. Al-Issa, M. A. Ottom, and A. Tamrawi, "eHealth cloud security challenges: A survey," *Journal of Healthcare Engineering*, vol. 2019, Article ID 7516035, 2019.
14. A. H. Seh, M. Zarour, M. Alenezi, M. Sarkar, D. Agrawal, R. Kumar, and R. Ahmad Khan, "Healthcare data breaches: Insights and implications," *Healthcare*, vol. 8, no. 2, 2020.
15. D. Alhuwail, E. Abdulsalam, and S. Al-Sabah, "Information Security Awareness and Behaviors of Health Care Employees in Kuwait," *JMIR Medical Informatics*, vol. 9, no. 9, 2021. IBM Security and Ponemon Institute, "Cost of a Data Breach Report 2023," 2023.