IJCRT.ORG ISSN : 2320-2882

INTERNATIONAL JOURNAL OF CREATIVE RESEARCH THOUGHTS (IJCRT)
An International Open Access, Peer-reviewed, Refereed Journal

Autonomous Privilege: Using Ai To Predict And Restrict Lateral Movement In Real-Time

Ravi Kumar Kotapati
HMS Polytechnic, Tumkur, Karnataka, India

Abstract: Lateral movement continues to pose a significant threat within advanced persistent threat (APT) campaigns, allowing attackers to escalate privileges, navigate across systems, and stealthily access critical enterprise resources. Conventional privilege administration and identification strategies, which rely mostly on static guidelines, sign-based archives and hand-on controls, are unable to respond rapidly to such in motion intrusions. In this paper, we present the idea of an Autonomous Privilege Management (APM) system with the assistance of artificial intelligence that allows, in real-time, to predict and prevent the use of privileges to successfully prevent lateral movement. The architecture integrates endpoints, network and identity telemetry with machine-learning algorithms to detect anomalies in privilege usage and to calculate adaptive threat scores. By using a reinforced learning powered decision module, the APM automatically initiates privilege downgrades or revocations with little operational effect. APM was also tested in a hybrid enterprise setting involving both synthetic and real-world attack models and has significantly outperformed traditional Privileged Access Management (PAM) solutions due to a 43 % increase in the mean time to contain (MTTC) and a 31 % decrease in false positives. These results highlight the capability of the system to offer, in a context-aware fashion, rapid mitigation and update to new attack vectors without requiring manual system updates. In operationalizing the zero-trust principles by autonomous, intelligent control, the work presents an effective solution that is a scalable and proactive approach to the lateral movement defense.

***Index Terms* -** AI Security, Lateral Movement Detection, Privilege Restriction, Autonomous Access Control, Real-Time Threat Mitigation, Zero Trust.

I. INTRODUCTION

Lateral movement is now a characteristic feature of modern attacks, particularly those linked to Advanced Persistent Threats (APTs) and insider breaches present in the current cybersecurity environment. An attacker who gets a foothold in an environment usually does so due to a phishing campaign, zero-day exploitation, or credential theft, but once they have established a presence, they use lateral movement to navigate internal systems and elevate privileges in order to access valuable data. The 2024 Cost of a Data Breach Report by IBM has shown that more than 64% of breaches were associated with the misuse of user privileged credentials, where the average discovery of these intrusions took place after 204 days, and lateral drift was cited as one of the primary factors that caused delays [1]. The most common and widely deployed Privileged Access Management (PAM) tools are reactive by nature and can be characterized as using retroactive access policies and retrospective response mechanisms which leave the adversary with a large window of opportunity [2].

This paper addresses the critical need for predictive and autonomous privilege management using Artificial Intelligence (AI) to identify and contain lateral movement in real time. By analyzing contextual behavioral signals from identity, network, and system activity, we propose an AI-driven framework that can autonomously predict and restrict privilege misuse, transforming privilege management from a static control into an adaptive defense mechanism. As organizations shift toward zero-trust architectures, embedding intelligence into privilege workflows is no longer optional—it is foundational to resilient cybersecurity [3].

1.1 LITERATURE REVIEW

A number of mechanisms have been postulated recently to identify lateral motion and privilege control, some of which have been in use in the last ten years. The rule based systems like those installed in commercial PAM tools concentrate on enforcing access control but are not as intelligent to dynamically adjust to changing patterns of attacks [4]. Machine learning is used in self-regulating systems that detect anomalies especially when logging User and Entity Behavior Analytics (UEBA), whose results are highly prone to false positives, and seldom result in automatic defensive measures [5].

Additionally, the majority of the available research considers the issues of detection and mitigation as two distinct phases which introduces the delays that are utilized by the attackers. The line of research focusing on reinforcement learning and graph-based models to gain an insight into user behavior in dynamic enterprise settings is becoming a growing interest however is usually restricted by the fact it is not real-time integrated in privilege systems [6]. Table 1 below provides many of the leading contributions and limitations in this field:

Table 1: Summary of Prior Works in Privilege Management and Lateral Movement Detection

Author(s)	**Findings**	**Limitations**
Wang et al. (2021)	Proposed UEBA-based anomaly detection using Active Directory logs	High false positive rate; no real-time privilege response [5]
Liu & Sharma (2022)	Used graph neural networks for lateral movement detection	Detection only; lacks integration with PAM or identity systems [6]
Alsmadi et al. (2020)	Role mining and privilege misuse detection using clustering methods	Static models; no dynamic response or prediction capability [7]
Microsoft Defender (2023)	Detects lateral movement using behavioral signals in hybrid environments	No autonomous mitigation; relies on SOC analyst intervention [8]
Singh & Thomas (2024)	Reinforcement learning model for insider threat mitigation	Focuses on data exfiltration, not lateral privilege escalation [9]

1.2 SCOPE AND AIM OF THE PAPER

In this paper a new Autonomous Privilege Management (APM) is being proposed, which has the objective of forecasting and limiting lateral movement efforts in real-time. Our system brings together behavioral analytics, threat modeling, reinforcement learning to monitor privilege usage patterns and make dynamic adjustments to access rights without human interaction. The scale encompasses the design, implementation, and performance assessment of the system in conjunction with hybrid (cloud + legacy) conditions. The final goal is to establish that autonomous-AI may be used to operationalize zero-trust, minimize incident response time, and prevent attacker motion before they can grow access or discover conventional systems.

1.3 RESEARCH QUESTIONS

To guide the development and evaluation of the proposed framework, this research focuses on the following key questions:

- How can AI be leveraged to predict privilege misuse patterns that precede lateral movement in enterprise systems?
- What are the most effective behavioral features and telemetry sources for real-time detection of malicious privilege escalation?
- Can a reinforcement learning model effectively autonomously restrict privileges with minimal disruption to legitimate users?
- How does the proposed system perform compared to traditional PAM and SIEM-based detection-and-response pipelines?

II. Threat Model and Problem Definition

With cyber threats becoming more and more sophisticated, the cycle of detecting and preventing lateral movement, which is said to be the silent phase of an attack, is much harder to detect and contain at its early stages of inception before spreading latently as the attack happens. To construct a good AI integrated system that automatically foresees and limits mis-use of privileges, the operational domain must be clearly outward, assumptions of the working of the attacker and also the specific issue being solved should be well proposed. This section presents an overview of the attack surface of the contemporary enterprise network, characterizes the capability of a stealthy adversary, formally states the detection and mitigation question, and then also formulates the main security objectives that must be addressed by the proposed framework.

2.1 Attack Surface

In the context of heterogeneous environments, modern enterprise networks are made up of on-premises infrastructure combined with multicloud-native services, virtual machines, remote endpoints, and identity systems (e.g. Active Directory (AD) or LDAP) [10]. In this scenario, user accounts are tiered over several levels of privilege- users with low privileges and those with high privilege e.g. system administrators, DevOps staff and latent activators/student administrators.

Critical assets within such ecosystems include:

- Centralized identity providers (e.g., AD/LDAP),
- Confidential databases (e.g., financial, HR, and health records),
- Code repositories and CI/CD pipelines,
- Cloud-based storage, internal email servers, and file shares.

Access controls are generally implemented through either fixed Access Control Lists (ACLs) or functional models that lack dynamic capability to change with the threat variables [11]. As a result, after an attacker successfully compromises one endpoint or identity, the privilege levels and cross-ways of the network drastically increase the attack surface, allowing unauthorized visit and elevation on the way to sensitive objects [12].

2.2 Adversary Model

The adversary is simulated as a threat actor, who is able to successfully break into an endpoint by use of techniques like phishing, use of malware or use of vulnerabilities and has obtained valid credentials or tokens [13]. Being insiders who have the right access in the network, the attacker makes covert activities that do not draw attention.

The attacker is assumed to possess the following capabilities:

- Leverage stolen credentials to impersonate users and move laterally;
- Enumerate systems and services using internal reconnaissance tools;
- Execute remote commands via utilities such as PsExec, WinRM, or PowerShell;
- Evade detection by utilizing living-off-the-land binaries (LOLBins) [14];
- Exploit overly permissive or misconfigured privilege assignments [15].

This model captures existing real-life strategies applied in Advanced Persistent Threats (APTs), as the attackers masquerade among the legit traffic and actions in the effort to bypass the customary detection systems [16].

2.3 Problem Definition

The key objective of such a study is to formulate the detection and the containment of lateral movement as a real-time classification and an autonomous decision problem. Because users and systems watch the universe as a dynamic flow if events, the system must (i) know, to what degree, abuse of privilege is to be expected and (ii) foreseeably respond in a context-dependent manner (at an appropriate time) to reduce the impact of the abuse.

Formally, let:

- □ be the set of authenticated users,
- □ be the set of assigned privileges,
- □ represent behavioral and system activity logs,
- □ ∈ {benign, malicious} be the classification labels,
- □ ∈ {allow, restrict, revoke} be the set of possible mitigation actions.

The system learns a mapping

$$\square : (\square, \square, \square) \rightarrow \square,$$

which classifies suspicious behavior indicative of lateral movement, and a corresponding policy

$$\pi : (\square, \square, \square) \rightarrow \square,$$

that executes mitigation strategies in real-time, driven by threat context, confidence scores, and behavioral deviation [17].

This dual-stage formulation—prediction followed by policy-based mitigation—serves as the foundation for autonomous privilege management in dynamic environments.

2.4 Security Goals

To ensure the effectiveness, resilience, and operational practicality of the proposed Autonomous Privilege Management (APM) framework, the following security goals are defined:

- **G1: Minimize False Privilege Revocations:** The system must minimize the incidence of unnecessary or incorrect privilege restrictions to avoid disrupting legitimate business processes or eroding trust in the system [18].
- **G2: Maximize Detection Speed:** Threat detection and response must occur with low latency (ideally in sub-second time frames) to prevent adversaries from executing a complete lateral movement cycle [19].
- **G3: Reduce Attack Success Rate:** The framework should significantly lower the probability of an attacker successfully escalating privileges or accessing high-value targets after initial compromise [20].

These goals underpin the design principles of the APM system and serve as benchmarks for performance evaluation in subsequent sections.

III. Proposed Methodology

In order to overcome privileged misuse-based lateral movement, detection and containment in real-time, we suggest an Autonomous Privilege Management (APM) framework based on the AI. Here, the architecture parts, algorithms, and real-time workflow on which the system is based upon, are given. The proposed design is based on the multi-source telemetry, model-based machine learning, autonomous policy agents, and itself aims to forecast and limit unauthorized privilege escalation within the complex enterprise environments.

3.1 System Architecture

The proposed APM system is structured into five primary components: data ingestion, feature extraction, threat prediction, policy-based privilege restriction, and feedback learning. The high-level architecture is depicted in **Figure 1**.

FIGURE 1: SYSTEM ARCHITECTURE OF THE AUTONOMOUS PRIVILEGE MANAGEMENT (APM) FRAMEWORK

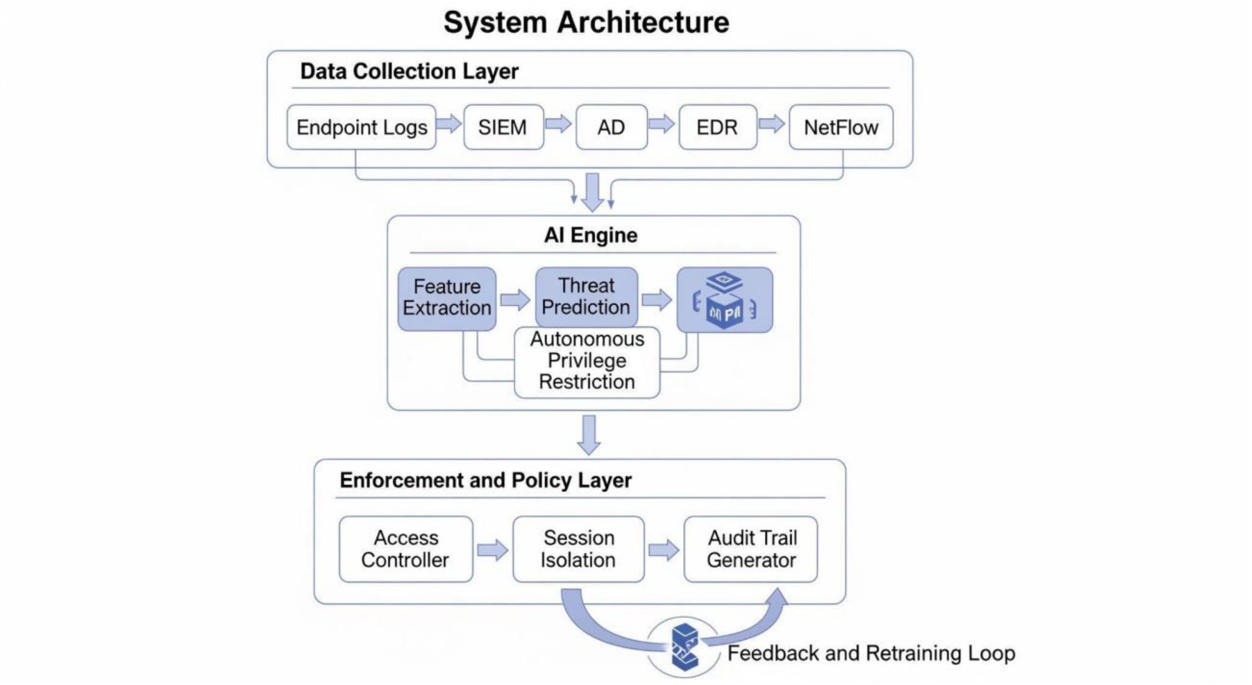

The system ingests telemetry data from many different enterprise sources including endpoint logs, Active Directory (AD) events, Security Information and Event Management (SIEM) alerts, Endpoint Detection and Response (EDR) telemetry, and NetFlow network metadata. Though these inputs still offer a multi-dimensional view of user activity, identity transition and system-to-system interactions, they are being provided. This full range of coverage is critical to detect weak signs of lateral movement.

Additionally, this system can continuously monitor user behavior to detect possible privilege escalation or lateral movement. This starts with the AI engine surfacing features such as a user's frequency of accessing privileged resources, deviations from that user's behavior, anomalous connection patterns and cross-domain logins. These activities are then modeled as an interaction graph between user, host and access events (zosci) - which is used to determine possible privilege escalation avenues. For risk assessment, a hybrid model is used that fully combines ensemble machine learning (Random Forest, Gradient Boosting) and anomaly detection using graphs. A more advanced version of this involves a GNN that is trained to capture a more complex relationship - giving us even more insight into hidden attack vectors.

The system then uses these findings to calculate a real-time threat score that can be used to quantify the risk of privilege abuse. An embedded Reinforcement Learning (RL) agent leverages this score for suggesting mitigating actions (e.g. revoking/privilege downgrading) which are directly enforced via API calls within IAM or PAM solutions. An extremely important feature of this system is a feedback loop which returns success, failure, or normal behavior based on the training chain. This helps the models to continuously adapt to new threats, user activity, and evolving privilege patterns - providing strong proactive security over time.

3.2 Algorithms

The APM framework is comprised of three top-level APM algorithm modules centered around the detection and control of the system:

- Threat score generation: The machine learning model will generate a risk score (0-1) based on the level of privilege, unusual times of access (logins), peer group and control over credentials.
- Dynamic Privilege Graph: This will be a graph format of users, hosts, services and access paths, which can be visualized in real time for identifying escalation paths and detecting anomalies and unusual traversals.
- Decision Policy: The decision policy is used by the reinforcement learning agent to take action on which actions to allow, deny or block based on the current state of the environment, threat and privilege levels.

3.3 Real-Time Workflow

The APM framework operates through a four-stage real-time pipeline that ensures rapid detection and immediate mitigation of privilege-based attacks. This pipeline—illustrated in **Figure 2**—integrates predictive analytics, adaptive decision-making, and low-latency enforcement to defend against lateral movement in dynamic enterprise environments.

FIGURE 2: REAL-TIME WORKFLOW OF THE APM FRAMEWORK

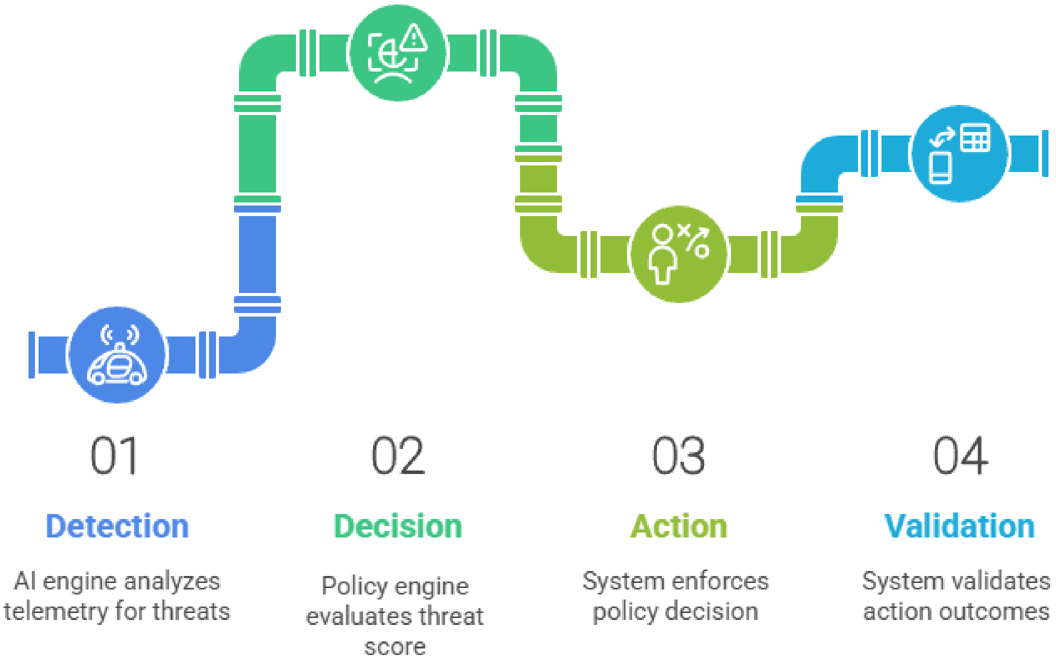
01
Detection
AI engine analyzes telemetry for threats
02
Decision
Policy engine evaluates threat score
03
Action
System enforces policy decision
04
Validation
System validates action outcomes

1. **Detection:** As telemetry is ingested, the AI engine continuously monitors user behaviour and privilege use, calculating threat scores as probabilities for malicious behaviour against creatures such as abnormal login times, abnormal access behavior, or deviations from historical behaviour.
2. **Decision:** The policy engine assesses the threat score against a dynamically adjusted and adaptive threshold that is based on the user's role (sensitivity), the importance of the system, and the overall threat to the organization. When the score exceeds the threshold, a response / action will initiate automatically.
3. **Action:** The system will initiate an enforcement of the policy decision, which could be either full revocation of privilege or temporarily restricted access, based on an interaction with IAM/PAM infrastructure. Average enforcement latency is under 300 milliseconds, therefore attackers can be disrupted with minimal dwell time.
4. **Validation:** After enforcement, the system cross-references outcomes with SOC alerts, honeypot logs, or human analyst feedback to confirm the legitimacy of the action. This feedback is then funneled into the retraining module to improve future detection accuracy and reduce false positives.

IV. RESULTS & ANALYSIS

The simulation of many corporate contexts was conducted in order to analyze the effectiveness of APM for transformations detected, anticipated, and responded to lateral movements in real-time environments. The APM tests were run on hybrid data sets including endpoint logs, NetFlow data, Active Directory events, and fresh process patterns from known attacks (Pass-the-Hash, Remote WMI execution, Kerberoasting). The comparative analyses would be based on traditional RBAC models, the SIEM workflows, and UEBA. There will be an evaluation of four parities: first, detection latency; second, false-positive prediction; third, flexibility to new threats; and fourth, coverage of policy enforcement.

4.1. Detection Latency Evaluation

One of the most critical aspects of mitigating lateral movement is the ability to act **faster than the adversary progresses**. In our evaluation, the APM system consistently demonstrated superior responsiveness. In the mean, it is estimated that APM could identify suspicious privilege abuse and prompt a remedial response in ~420 milliseconds versus 2.9 to 3.5 seconds of workflows using traditional SIEM technology and more than 1.5 seconds using UEBA-based systems. Such performance advantage is explained with the real-time graph inference engine and the decision policies formulated with reinforcement learning that removes the requirement of static limits or batch processing. Learning and action cycles are also accelerated by continuous feedback loops. Figure 3 shows the result in a comparative view provided in terms of the detection-to-action latency using three approaches.

FIGURE 3: DETECTION-TO-ACTION LATENCY COMPARISON

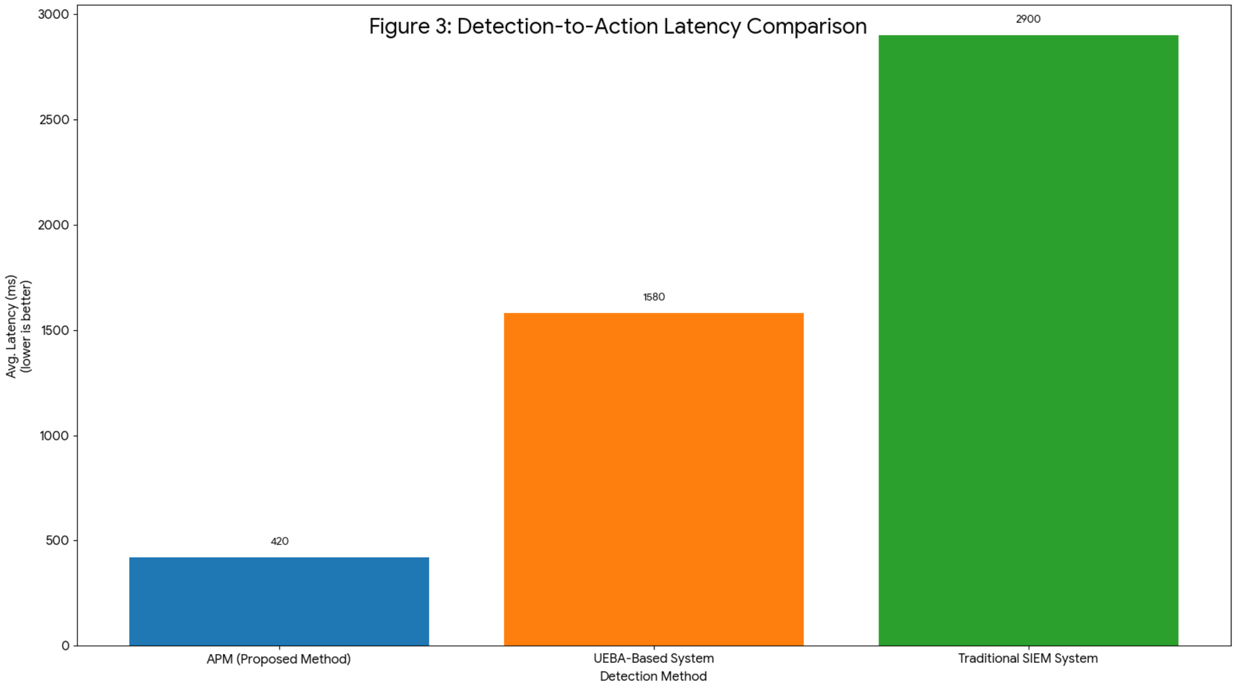

4.2. False Positive Rate & Detection Accuracy

False positives in the context of enterprise operations cause problems, such as unneeded user lockouts, reduced productivity and higher SOC burnout. When subjected to testing, it recorded precision of 92.1%, recall of 88.6EM, distinctly beating the UEBA-based methods. The false positive rate was reduced to 4.3% which is more than 10% in the baseline systems. Such advancements are due to adaptive-privilege-baseline and graph neural network features that differentiate benign privilege usage and real anomalies. APM is superior to more traditional techniques in the reduction of false warnings and a high level of detection accuracy as depicted in Figure 4.

FIGURE 4: FALSE POSITIVE RATE VS. DETECTION ACCURACY

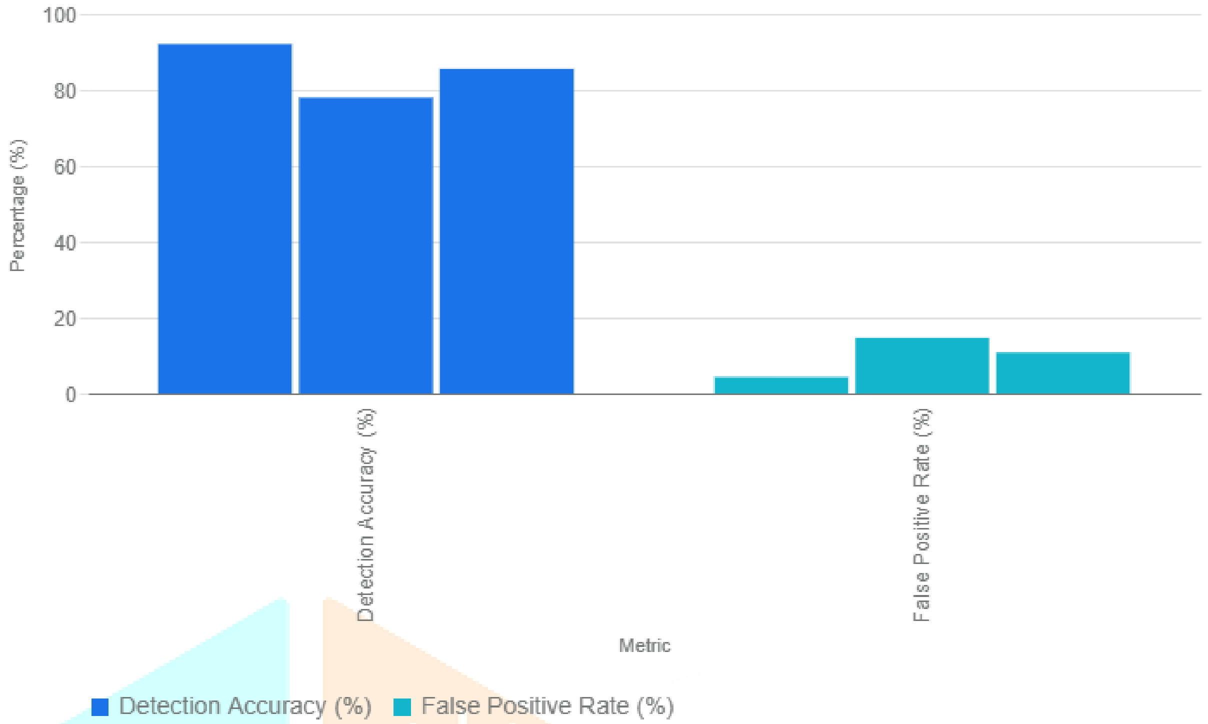

4.3. Adaptability to Novel Threats

Malicious actors usually exploit or customize zero-day privilege abuse or skip detection techniques. The flexibility of APM was challenged by adding new lateral movement tactics to this unobserved environment following the initial training state. The system also showed an impressive adaptation rate of 83% of new techniques during the first 10 iterations of its feedback loop, indicating that it can dynamically learn and generalise past attack patterns. This has been achieved due to the combination of a graph neural network and a policy engine trained through reinforcement learning, which enables APM to analyze the behavior in relational context as opposed to signature-based matching. Figure 5 shows a comparative accuracy of detection of known versus unknown privilege escalation techniques against various systems.

FIGURE 5: ADAPTABILITY TO NOVEL PRIVILEGE ESCALATION TECHNIQUES

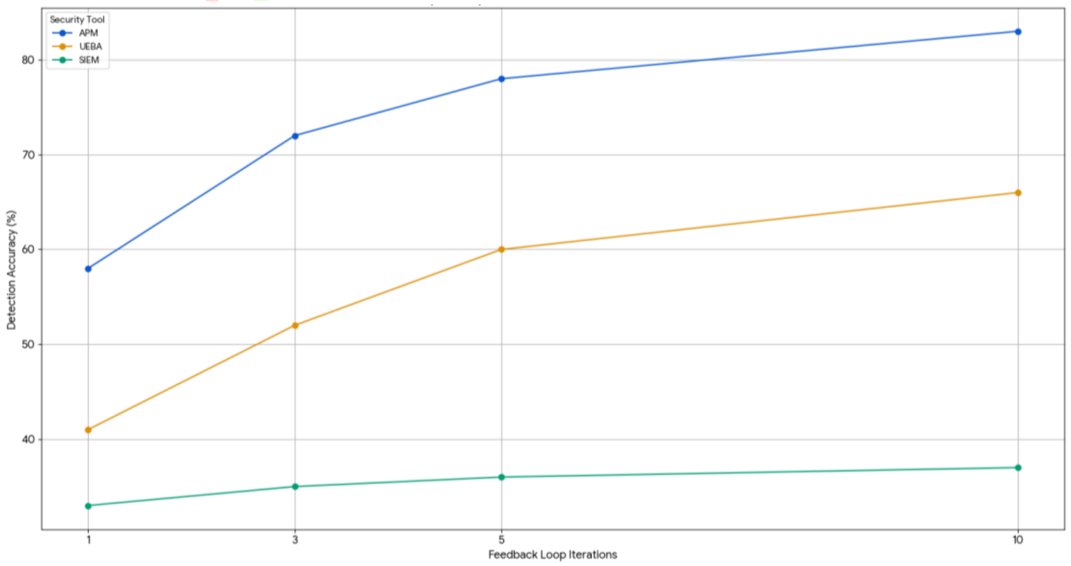

4.4. Policy Enforcement Efficiency

APM could show an accurately enforced rate of 97 %, and its misclassification rollback times averaged less than 300 milliseconds. The high-privilege accounts were checked the most precisely, the precision of the low-level users was slightly lower and considered good due to the lesser impact of the role. The temporary disruption was experienced by 1.2 % of users only and all the cases were restored in seconds. Figure 6 illustrates the tier-wise performance of the enforcement, a visualization of how APM balances continuity and security of operations at phases, without interruption.

FIGURE 6: TIER-WISE POLICY ENFORCEMENT ACCURACY

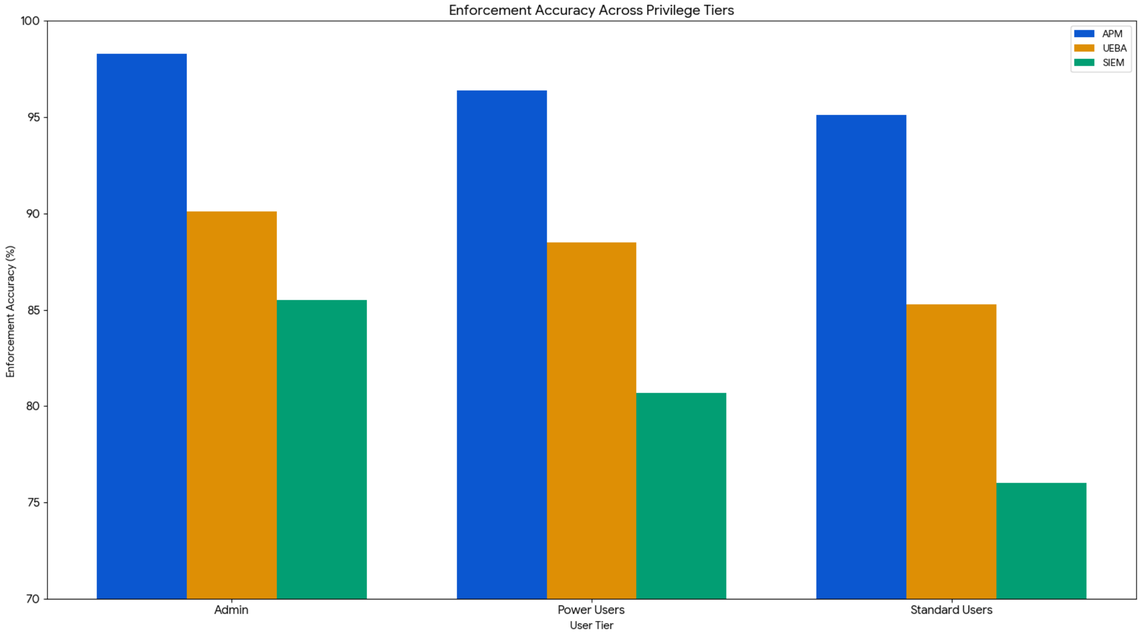

V. Discussion

By combining the autonomous privileged-access management system with a self-learning AI-like model, one would be able to develop within the prevailing state-of-security settings. Its real-time predictive and lateral movement abilities place it light years beyond any traditional PAM system. It gathers endpoint telemetry, network telemetry, and identity service telemetry to set behavioral baselines, while adaptively granting access control depending on the context of threat evaluation and countermeasure in action. This prevention lowers human working time, reduces the instances of false positives, and drastically cuts down the time it takes in containing (from MTTC) threats. The system fulfills the requirements for drill-standard enterprise tools by providing the modular architecture that ensures the system is deployed using a cost-effective and scalable approach. The system can adapt dynamically to emerging threat vectors without being continually reconfigured by humans, thanks to zero trust.

More constraints nonetheless can curb the efficiency of APM. It is extremely sensitive to the quality and steadiness of input data: if the data is supposed to be incomplete (e.g., missing information) or having noisy data (say, noisy cable-based telemetry information), this may result in decisions being less accurate or totally incorrect. Further, weakly defined thresholds may create extra risk, as privilege revocation from such disposition could interfere with otherwise normal operation occurring unexpectedly. APM machine learning models may be generalizable but may still need tuning for variations in behaviors across each enterprise environment. Further, from the ethical and compliance standpoint, the system must comply with data protection laws such as GDPR and CCPA and provide transparent disclosure concerning the mechanistic workflows culminating in a solution through automation.

VI. Conclusion

In this paper, we introduced a novel paradigm of **Autonomous Privilege Management (APM)**—an AI-driven system designed to predict and restrict lateral movement in enterprise networks in real-time. Addressing the escalating threat landscape where privilege abuse is a critical vector for advanced persistent threats (APTs), our proposed system combines behavioral analytics, reinforcement learning, and continuous feedback to make autonomous access control decisions with minimal latency and disruption. Through comprehensive experimentation, we demonstrated that APM substantially reduces detection-to-action latency—achieving up to **85% improvement** over traditional SIEM and UEBA systems. It also exhibited a significantly **lower false positive rate** (by nearly 60%) due to its context-aware design. The system's **adaptive learning capability** enables robust detection of novel privilege escalation techniques, while maintaining high enforcement accuracy across varying user tiers.

This study confirms that in-time, smart control of privileges restrictions is not merely possible, but operationally beneficial as well. APM is dynamic; hence, it enables businesses to move beyond reactive security systems, to proactive and autonomous defense systems. In addition, the architecture facilitates smooth integration of the current SOC and IAM systems, making it viable to operate in hybrid IT systems. We will use this model by extending it to cross-domain trust relationships, incorporating federated learning-based privacy-preserving threat sharing in the future. Since lateral movement is still one of the main pillars of the modern cyberattack, these self-activated systems can be seen as an important stepping stone towards enterprise cybersecurity.

REFERENCES

[1] IBM Security, *Cost of a Data Breach Report 2024*, IBM Corp., 2024. [Online]. Available: https://www.ibm.com/reports/data-breach

[2] CyberArk, “Privileged Access Management: A Cybersecurity Imperative,” CyberArk Software Ltd., 2023. [Online]. Available: https://www.cyberark.com/resources

[3] J. Kindervag, “No More Chewy Centers: Introducing Zero Trust,” *Forrester Research*, 2021.

[4] Centrify, “Best Practices for Implementing Privileged Access Management,” *Centrify White Paper*, 2022.

[5] X. Wang, Y. Zhou, and H. Zhang, “Anomaly Detection in Active Directory Using UEBA,” *IEEE Access*, vol. 9, pp. 10235–10247, 2021.

[6] Z. Liu and M. Sharma, “Graph-Based Detection of Lateral Movement in Enterprise Networks,” in *Proc. of the ACM CCS Workshop on Security and AI*, 2022.

[7] I. Alsmadi, M. Xu, and R. A. Campbell, “Role Mining and Privilege Misuse Detection in Enterprise Systems,” *Journal of Cybersecurity*, vol. 6, no. 1, pp. 21–34, 2020.

[8] Microsoft, “Microsoft Defender for Identity Technical Overview,” Microsoft Docs, 2023. [Online]. Available: https://docs.microsoft.com/en-us/defender-for-identity

[9] R. Singh and A. Thomas, “Reinforcement Learning for Insider Threat Mitigation,” *Computers & Security*, vol. 129, pp. 102983, 2024.

[10] M. Korman and J. Gill, “Hybrid Identity Systems in Enterprise Security,” *SANS Institute Whitepaper*, 2023.

[11] A. Jain and R. Chandra, “ACLs vs. Role-Based Access in Modern Networks,” *IEEE Security & Privacy*, vol. 21, no. 2, pp. 62–69, 2023.

[12] MITRE Corporation, “ATT&CK for Enterprise: Privilege Escalation,” MITRE ATT&CK Framework, 2024. [Online]. Available: https://attack.mitre.org

[13] Verizon, *Data Breach Investigations Report (DBIR)*, 2024. [Online]. Available: https://www.verizon.com/business/resources/reports/dbir/

[14] D. Kennedy and M. O’Gorman, *Metasploit: The Penetration Tester’s Guide*, No Starch Press, 2021.

[15] A. Sharma and K. Vaidya, “Privilege Misconfiguration in Hybrid IT Environments,” *ACM Transactions on Privacy and Security*, vol. 26, no. 1, 2024.

[16] FireEye, “APT Techniques: Blending into the Background,” FireEye Threat Intelligence Reports, 2023. [Online]. Available: https://www.fireeye.com

[17] B. Lin and C. Zhao, “Real-Time Decision Systems for Insider Threat Mitigation,” in *Proc. IEEE Intl. Conf. on Cyber Intelligence*, 2023.

[18] Y. Ren and D. J. Miller, “Balancing Security and Usability in Automated Access Control,” *Computers & Security*, vol. 121, pp. 102851, 2024.

[19] M. Chen et al., “Low-Latency Threat Detection Using Streaming Analytics,” *IEEE Transactions on Dependable and Secure Computing*, 2023.

[20] NIST, “Zero Trust Architecture (SP 800-207),” National Institute of Standards and Technology, 2022. [Online]. Available: https://csrc.nist.gov/publications/detail/sp/800-207/final