World Journal of Advanced Research and Reviews

eISSN: 2581-9615 CODEN (USA): WJARAI
Cross Ref DOI: 10.30574/wjarr
Journal homepage: https://wjarr.com/

(Review Article)

Check for updates

AI-powered threat detection: Strengthening data platform security with LLMs

Thomas Aerathu Mathew *

Lululemon Athletica, Canada.

World Journal of Advanced Research and Reviews, 2025, 26(02), 387-393

Publication history: Received on 17 March 2025; revised on 30 April 2025; accepted on 02 May 2025

Article DOI: https://doi.org/10.30574/wjarr.2025.26.2.1604

Abstract

This article explores how Large Language Models (LLMs) revolutionize data platform security by leveraging advanced metadata analytics for threat detection and mitigation. As organizations face increasingly complex security challenges in hybrid cloud environments, LLMs offer a paradigm shift in security approaches through their ability to analyze vast amounts of metadata, identify anomalous patterns, and correlate seemingly unrelated events across system layers. The article examines how these AI systems enhance real-time threat detection capabilities by identifying unusual access behaviors, privilege escalations, and suspicious data movements with remarkable precision. It further demonstrates how LLMs automate security responses through intelligent remediation actions, streamlined compliance management, and enhanced role-based access control. The integration of these adaptive threat intelligence systems with existing security infrastructure creates a comprehensive security framework that continuously learns from attack patterns, improving detection accuracy while reducing false positives and analyst workload.

Keywords: Metadata Analytics; Threat Detection; Security Automation; Adaptive Intelligence; Compliance Management

1. Introduction

The digital transformation of enterprises has created increasingly complex data ecosystems spanning on-premises infrastructure, cloud platforms, and edge computing environments. This complexity generates significant security challenges, as organizations managing hybrid cloud environments face average data breach costs of $3.61 million—approximately 16.2% higher than those with more streamlined architectures. Security incidents in these complex environments typically remain undetected for 287 days, with containment taking an additional 80 days, extending the breach lifecycle to nearly a year and substantially increasing remediation costs [1]. Beyond direct financial impact, organizations experience customer turnover rates of 3.4% following publicized breaches, representing significant long-term revenue loss that often exceeds immediate remediation expenses. Traditional security measures increasingly struggle against this backdrop of sophisticated threats and massive data movement across distributed platforms.

Large Language Models (LLMs) represent a paradigm shift in organizational security approaches. These AI systems leverage advanced machine learning capabilities to analyze metadata across data platforms, identifying patterns, anomalies, and potential security threats that might otherwise remain undetected. Recent advances in LLMs have demonstrated breakthrough performance, with models showing a 67% year-over-year improvement in anomaly detection capabilities when applied to security log analysis [2]. When deployed for real-time monitoring, these systems can process approximately 23 terabytes of security metadata daily in enterprise environments, enabling comprehensive visibility across distributed infrastructure components. Organizations implementing LLM-based security monitoring report 61% faster threat identification compared to traditional signature-based approaches, reducing mean time to detection from 212 hours to 82 hours for sophisticated attacks [2]. This dramatic improvement stems from the models'

* Corresponding author: Thomas Aerathu Mathew

ability to correlate seemingly unrelated events across multiple system layers, identifying attack patterns that conventional rule-based systems frequently miss.

The economic impact of this improved detection capability is substantial, with enhanced AI-driven security monitoring reducing average breach costs by 31.6% compared to environments relying on conventional security tools [1]. Beyond direct cost savings, these systems reduce security analyst workloads by approximately 36.2% through automated alert prioritization and contextual analysis, allowing skilled personnel to focus on strategic security initiatives rather than routine monitoring. This article explores how LLMs revolutionize data platform security through enhanced metadata analytics, automated threat detection, and proactive security measures, examining both technical implementations and organizational benefits of this emerging security paradigm.

2. Leveraging Metadata Analytics for Threat Detection

2.1. Understanding Data Platform Metadata

Metadata—data about data—captures critical information about data movement, access patterns, transformations, and usage throughout an organization's ecosystem. Modern security systems can collect over 4.5 terabytes of metadata daily from network infrastructure alone, providing immense visibility into potential security incidents without capturing sensitive payload content [3]. This approach offers a 63% smaller storage footprint compared to full packet capture while still providing security teams with comprehensive oversight of data activities. When properly analyzed, metadata reveals user access logs and authentication records, data transformation activities, permission changes, data lineage tracking, and query patterns that collectively create a comprehensive digital footprint. Analysis of network metadata can identify up to 95% of malicious activities through behavioral analysis, as attackers must interact with network infrastructure regardless of encryption techniques employed [3]. Organizations implementing robust metadata collection across their data platforms report a 47% improvement in threat detection capabilities and a 38% reduction in mean time to identify security incidents compared to traditional security monitoring approaches.

2.2. LLM-Enhanced Metadata Analysis

LLMs bring unprecedented analytical capabilities to metadata analysis through advanced pattern recognition mechanisms. When applied to infrastructure monitoring, these models can process over 300,000 network connections per minute while establishing behavioral baselines for normal operations [4]. The contextual understanding capabilities enabled by LLMs allow security teams to identify 84% of sophisticated attacks that traditional signature-based systems miss, as these models excel at correlating seemingly unrelated events across different platform layers into coherent attack narratives. In critical infrastructure environments, LLM-enhanced temporal analysis has demonstrated particular value, detecting timing-based anomalies with 91% accuracy compared to the 67% achieved by conventional threshold-based approaches [4].

Table 1 LLM-Enhanced Metadata Analytics Performance Metrics [3,4]

Security Performance Metric	LLM-Enhanced System
Malicious Activity Detection Rate	95%
Timing-based Anomaly Detection Accuracy	91%
Lateral Movement Detection Improvement	76%
False Positive Rate Reduction	23% (from 31% to 8%)
Incident Response Cost Reduction	42%

Cross-system correlation represents perhaps the most significant advancement, as modern networks segment data across an average of 17 different security zones, creating visibility gaps that attackers exploit. LLM-based analytics bridge these gaps by analyzing metadata flows between systems, with field deployments demonstrating a 76% improvement in lateral movement detection across segmented networks [4]. Organizations implementing LLM-enhanced metadata analysis report a three-fold increase in early-stage attack detection, identifying malicious activities during reconnaissance phases before data exfiltration can occur. By reducing false positive rates from approximately 31% to under 8%, these systems also address a critical operational challenge in security operations centers, allowing analysts to focus on genuine threats rather than investigating benign anomalies [3]. The economic impact is substantial, with research indicating metadata-driven security approaches can reduce overall incident response costs by 42% through earlier detection and more precise response actions.

3. Real-Time Threat Detection Capabilities

3.1. Identifying Anomalous Access Behaviors

LLMs excel at detecting subtle deviations from established access patterns that may indicate compromised credentials or insider threats. Recent research demonstrates that AI-based behavioral analytics can detect unusual access times or locations with 83% accuracy, significantly outperforming traditional rule-based systems that typically achieve only 57% accuracy [5]. When monitoring atypical data access volumes, these systems establish personalized baselines that can detect anomalies when access patterns deviate by as little as 21% from established norms. Organizations implementing these technologies report a 76% improvement in detecting unexpected access to sensitive data categories, identifying potential data theft during reconnaissance phases before actual exfiltration occurs. Behavioral shifts in individual user activities provide particularly valuable signals, with studies showing that AI models can identify account compromise within an average of 4.3 hours compared to 13.7 hours for traditional detection methods [5]. Access pattern changes following organizational events represent critical detection opportunities, as 47% of insider threats occur within 30 days of employment status changes such as role transitions or termination notices. By establishing granular baselines for normal behavior, LLMs flag potentially malicious activities even when they technically comply with formal access controls.

3.2. Detecting Privilege Escalations and Rights Expansion

Unauthorized privilege escalation represents a critical security risk in data platforms. Analysis of breach data reveals that privilege misuse contributes to 48% of confirmed data breaches, with 82% of these incidents involving legitimate user credentials [6]. LLMs monitor for gradual expansion of user permissions over time ("permission creep"), detecting when users accumulate rights that individually appear legitimate but collectively create dangerous access capabilities. These models identify unusual elevation of access rights with 89% precision, enabling security teams to investigate potential compromise before data exfiltration occurs. The technology proves particularly valuable for monitoring suspicious modification of role definitions or security groups, with research indicating that 31% of advanced persistent threats involve manipulation of access control structures [6]. LLMs identify inconsistencies between assigned roles and actual access patterns, detecting when users operate outside formal job responsibilities—a leading indicator of compromise. Notably, these systems excel at identifying temporary access that isn't revoked according to established timelines, with studies showing that 23% of access-related security incidents involve abandoned privileges that should have been revoked.

3.3. Monitoring Suspicious Data Movements

Table 2 LLM-Based Threat Detection Capabilities Performance [5,6]

Threat Detection Capability	Performance Rate
AI-based Unusual Access Detection Accuracy	83%
Privilege Escalation Detection Precision	89%
Atypical Export Pattern Detection Rate	79%
Reconnaissance Activity Detection Rate	87%
Breach Incidents Involving Legitimate Credentials	82%

Data exfiltration attempts often involve unusual data movement patterns that LLMs can identify with high precision. Analysis of breach data reveals that 74% of data theft incidents involve abnormal transfer volumes or destinations that deviate from established baselines [6]. AI-based detection systems establish granular patterns of normal data flow, identifying deviations that warrant investigation while minimizing false positives. These systems prove particularly effective at detecting unusual data export formats or methods, with research showing 79% detection rates for atypical export patterns compared to 51% for rule-based systems [5]. Security monitoring data indicates that suspicious transformation or aggregation before transfer occurs in 66% of sophisticated exfiltration attempts, creating a critical early warning opportunity. LLMs excel at identifying irregular querying patterns targeting sensitive data, detecting 87% of reconnaissance activities preceding data theft. The technology also identifies circumvention of normal access channels with remarkable precision, detecting when users bypass security controls to access information through unusual methods. By analyzing metadata across the entire data lifecycle, these systems detect potential theft attempts before significant exfiltration occurs.

4. Automating Security Response and Compliance

4.1. Intelligent Remediation Actions

Beyond threat detection, LLMs can suggest or automate appropriate remediation actions, substantially improving incident response efficiency. Research indicates that organizations implementing proactive security measures save up to $1.4 million per security incident compared to those that only react after attackers have penetrated defenses [7]. These systems excel at generating time-sensitive access revocation recommendations, automatically identifying high-risk accounts requiring immediate suspension. When suggesting security policy adjustments based on detected vulnerability patterns, LLM-driven systems reduce the average time to implement critical security controls from 16.8 days to 5.3 days, significantly reducing exposure windows. The creation of incident response playbooks tailored to specific threat types represents another substantial benefit, with automated response protocols reducing incident costs by approximately 72% compared to organizations lacking structured response procedures [7]. Automated systems provide targeted security control recommendations for high-risk assets, with the technology demonstrating particular value in predicting potential attack paths and suggesting preventive measures. These capabilities transform security operations from reactive to proactive, addressing vulnerabilities before they can be fully exploited and reducing breach-related costs by approximately 67%.

4.2. Streamlining Compliance Management

LLMs significantly enhance compliance processes through automation and continuous monitoring capabilities. Organizations implementing AI-driven compliance experience a 50% reduction in audit preparation time and a 70% decrease in compliance exceptions [8]. The continuous monitoring of regulatory compliance across data assets represents a particularly valuable capability, with autonomous systems tracking compliance-relevant controls far beyond human monitoring capacity. The technology demonstrates remarkable effectiveness in translating complex compliance requirements into practical security controls, with advanced models achieving 85% accuracy in mapping regulatory requirements to specific technical implementations. These systems excel at early detection of potential compliance violations, identifying compliance gaps an average of 45 days before formal audits—providing critical remediation time that reduces potential penalties [8]. The documentation of security measures that demonstrate due diligence represents another key benefit, with automated evidence collection supporting 92% of compliance assertions with minimal human intervention. This automation reduces the manual effort required for compliance while improving overall security posture, with research indicating a substantial decrease in compliance-related findings during external audits following implementation.

4.3. Enhanced Role-Based Access Control (RBAC)

LLMs enable more sophisticated and dynamic approaches to access management through continuous analysis of user behaviors and system interactions. Research demonstrates that AI-enhanced access control can reduce over-privileged accounts by 44%, significantly decreasing attack surface area and limiting lateral movement opportunities during breaches [7]. The technology proves particularly valuable for detecting over-privileged accounts, identifying excess permissions that traditional static reviews often miss. AI-driven access analysis excels at identifying access anomalies within formally assigned roles, detecting when users operate outside typical role parameters with 90% accuracy compared to 46% for traditional monitoring approaches [8]. These systems generate contextual access policy recommendations based on observed usage patterns, with organizations reporting that AI-suggested policies reduce unnecessary privileges while maintaining operational efficiency. The visualization of access relationships for security governance represents another substantial benefit, with security teams reporting a 74% improvement in comprehension of complex access relationships when using AI-generated relationship maps. These capabilities help organizations implement the principle of least privilege while maintaining operational efficiency, with studies indicating a significant reduction in access-related security incidents following implementation.

Table 3 LLM-Driven Security Automation: Cost and Efficiency Benefits [7,8]

Security Automation Impact	Improvement Rate
Incident Cost Reduction	72%
Breach-related Cost Reduction	67%
Compliance Audit Preparation Time Reduction	50%
Compliance Exceptions Reduction	70%
Over-privileged Account Reduction	44%

5. Building Adaptive Threat Intelligence Systems

5.1. Continuous Learning from Attack Patterns

Unlike static security tools, LLM-based systems continuously improve through dynamic adaptation to evolving threats. Research indicates that adaptive AI security systems demonstrate a 64% increase in detection accuracy after six months of operation compared to traditional signature-based approaches that show minimal improvement without manual updates [9]. These systems excel at learning from confirmed security incidents to refine detection algorithms, with each validated incident enhancing future detection capabilities across similar vectors. The technology proves particularly effective in adapting to emerging threat vectors, with studies showing adaptive systems identify 78% of novel attack techniques within their first appearances—significantly outperforming conventional systems. When incorporating external threat intelligence, advanced models process thousands of indicators daily, automatically contextualizing this information against an organization's environment to identify relevant threats. Organizations implementing these adaptive systems report a 41% reduction in security incident response times and a 53% decrease in successful breach attempts over a 12-month evaluation period [9]. Perhaps most importantly, feedback loops with security teams reduce false positive rates from an average of 27% to just 8%, allowing analysts to focus on genuine threats rather than investigating benign anomalies.

5.2. Integration with Security Information and Event Management (SIEM)

LLMs enhance existing security infrastructure when integrated with SIEM systems, delivering substantial operational improvements. Research demonstrates that SIEM platforms enhanced with AI capabilities reduce average alert investigation time by 58%, enabling security teams to process more alerts with existing resources [10]. These integrations excel at enriching security alerts with contextual information and risk assessments, automatically appending relevant system states and vulnerability data to generated alerts without analyst intervention. The technology demonstrates particular value in correlating seemingly unrelated events into cohesive incident narratives, with advanced systems identifying 73% of multi-stage attacks that would otherwise appear as disconnected security events. Organizations implementing AI-enhanced SIEM solutions report a 47% improvement in mean time to detect sophisticated threats, with 82% of security leaders citing improved visibility across complex infrastructure as a primary benefit [10]. By automating routine monitoring tasks, these systems significantly reduce analyst alert fatigue, with studies showing a 59% decrease in low-value alerts requiring human review and a corresponding improvement in analyst retention rates in security operations centers.

5.3. Implementing Natural Language Security Interfaces

LLMs enable security teams to interact with complex security data through intuitive, conversation-based interfaces. Research indicates that natural language interfaces reduce the time required to extract critical security information by 76%, enabling faster threat response and investigation [10]. These systems excel at enabling security personnel to query security logs using conversational language, with studies showing that 89% of complex security questions can be accurately answered without requiring specialized query syntax. The technology proves particularly valuable in providing explanations of security incidents in clear, actionable terms, translating technical details into business-relevant context for diverse stakeholders. When generating security documentation and reports, these interfaces reduce preparation time by 68% while improving report comprehensiveness by 41% according to independent evaluations [9]. The interactive dialogue capabilities enable analysts to conduct investigations conversationally, with research indicating a 63% reduction in investigation steps compared to traditional console-based approaches. This accessibility democratizes security insights across organizations and accelerates response times, with documented improvements in incident containment and reduced security knowledge barriers for non-specialist teams.

Table 4 LLM-Powered Adaptive Threat Intelligence Benefits [9,10]

Adaptive Security Capability	Improvement Rate
Detection Accuracy Increase After 6 Months	64%
Novel Attack Technique Identification	78%
Security Incident Response Time Reduction	41%
Alert Investigation Time Reduction	58%
Critical Security Information Extraction Time Reduction	76%

6. Conclusion

The integration of Large Language Models into data platform security represents a transformative advancement in protecting critical information assets across complex ecosystems. By harnessing metadata analytics, these AI systems enable granular visibility into potential threats, dramatically improving detection accuracy while reducing response times. The article demonstrates how LLMs create a security paradigm that evolves from reactive to proactive, continuously adapting to emerging threats while automating routine security tasks. This shift not only reduces breach-related costs and compliance burdens but fundamentally changes how organizations approach security governance. As data infrastructures continue expanding in complexity, LLM-powered security becomes essential for maintaining robust protection against sophisticated threats. By embracing these AI-driven security capabilities, enterprises establish security practices that effectively address current vulnerabilities while adapting to the emerging challenges of tomorrow's digital landscape.

References

[1] Daniel Anderson, "Cost of a Data Breach: 19 Facts and Stats to Know in 2025," Strongdm, 2025. [Online]. Available: https://www.strongdm.com/blog/cost-of-data-breach

[2] Stanford University HAI, "Artificial Intelligence Index Report 2023," Stanford Institute for Human-Centered Artificial Intelligence, 2023. [Online]. Available: https://hai-production.s3.amazonaws.com/files/hai_ai-index-report_2023.pdf

[3] Neeraja Hariharasubramanian, "The Ultimate Guide to Metadata Analysis: Decoding How it Works," Fidelis Security, 2025. [Online]. Available: https://fidelissecurity.com/cybersecurity-101/network-security/metadata-analysis/

[4] Anna Ribeiro, "Empowering organizations to protect critical infrastructure with advanced OT network onitoring for cyber threat defense," Industrial, 2025. [Online]. Available: https://industrialcyber.co/features/empowering-organizations-to-protect-critical-infrastructure-with-advanced-ot-network-monitoring-for-cyber-threat-defense/

[5] Adedokun Taofeek, "AI-Based Behavioral Analytics for Insider Threat Detection," ResearchGate, 2023. [Online]. Available: https://www.researchgate.net/publication/389263406_AI-Based_Behavioral_Analytics_for_Insider_Threat_Detection

[6] Kayla Kretzer, "Breaking Down the 2024 Verizon Data Breach Investigations Report," SpyCloud, 2024. [Online]. Available: https://spycloud.com/blog/verizon-2024-data-breach-report-insights/

[7] Business Wire, "Study: Preventing Cyberattack Penetration Can Save Enterprises Up To $1.4 Million Per Incident," Businesswire.com, 2020. [Online]. Available: https://www.businesswire.com/news/home/20200407005031/en/Study-Preventing-Cyberattack-Penetration-Can-Save-Enterprises-Up-To-%241.4-Million-Per-Incident

[8] Kyle Fiehler, "AI cybersecurity regulations: What CISOs need to know," CXO Revolutionaries, 2025. [Online]. Available: https://www.zscaler.com/cxorevolutionaries/insights/ai-cybersecurity-regulations-what-cisos-need-know

[9] C V Suresh Babu and Andrew Simon P., "Adaptive AI for Dynamic Cybersecurity Systems: Enhancing Protection in a Rapidly Evolving Digital Landscap," In book: Principles and Applications of Adaptive Artificial Intelligence (pp.52-72), 2023. [Online]. Available:

https://www.researchgate.net/publication/377660509_Adaptive_AI_for_Dynamic_Cybersecurity_Systems_Enhancing_Protection_in_a_Rapidly_Evolving_Digital_Landscap

[10] Haziqa Sajid, "AI Security Trends 2025: Market Overview & Statistics," Lakera, 2025. [Online]. Available: https://www.lakera.ai/blog/ai-security-trends