The American Journal of Interdisciplinary Innovations and Research
ISSN 2642-7478 | Open Access

TYPE Original Research
PAGE NO. 71-77

Scalable Vulnerability Management and Automated Threat Mitigation in Healthcare Ecosystems: An Analysis of AI-Driven Frameworks for 100K+ Asset Environments

Lina R. Al-Shaqiri
Independent Researcher, Interdisciplinary Innovations in Patient-Centric Cyber Safety & Asset Governance, Amman, Jordan

OPEN ACCESS
SUBMITTED 11 November 2025
ACCEPTED 21 November 2025
PUBLISHED 27 November 2025
VOLUME Vol.07 Issue 11 2025

CITATION
Al-Shaqiri, L. R. (2025). Scalable Vulnerability Management and Automated Threat Mitigation in Healthcare Ecosystems: An Analysis of AI-Driven Frameworks for 100K+ Asset Environments. The American Journal of Interdisciplinary Innovations and Research, 7(11), 71–77.

Abstract: The rapid digitization of healthcare infrastructure has precipitated a critical security challenge: the management of vulnerabilities across massive, heterogeneous environments often exceeding 100,000 assets. This paper investigates the efficacy of Artificial Intelligence (AI) and automated frameworks in mitigating cyber threats within these high-density clinical ecosystems. By analyzing recent breach statistics and contrasting legacy security models with modern cloud-based remediation tools, we evaluate the operational shift required to secure the Internet of Medical Things (IoMT). The methodology employs a comparative analysis of manual versus AI-driven vulnerability management cycles, focusing on metrics such as Mean Time to Remediate (MTTR) and false positive rates. Our analysis draws upon legal frameworks and industrial big data analytics to contextualize the technical findings within the broader scope of international governance and compliance. The results indicate that while legacy models fail to scale, AI-driven automated threat mitigation significantly reduces the window of exposure for critical clinical assets. However, the integration of these technologies introduces complex legal and ethical considerations regarding data privacy and algorithmic accountability. We conclude that a hybrid approach, combining automated "self-healing" networks with robust human oversight, is essential for the future resilience of healthcare information systems.

Keywords

Vulnerability Management, Healthcare Cybersecurity, Artificial Intelligence, Automation, IoT Security, Threat Mitigation, Clinical Information Systems.

1.INTRODUCTION

The digitalization of the global healthcare sector represents one of the most profound technological shifts of the twenty-first century. What began as the electronic storage of patient records has evolved into a hyper-connected ecosystem comprising Clinical Information Systems (CIS), cloud-based data repositories, and a rapidly expanding network of Internet of Medical Things (IoMT) devices. While this connectivity has revolutionized patient care, enabling remote monitoring and precision medicine, it has concurrently expanded the attack surface available to malicious actors. The contemporary healthcare environment is no longer a walled garden of isolated servers but a sprawling digital metropolis often containing over 100,000 discrete assets, ranging from MRI machines and infusion pumps to administrative laptops and cloud containers.

In this context, the traditional paradigms of cybersecurity are facing an existential crisis. Historical models, which relied heavily on perimeter defense and manual patch management, are proving insufficient against the velocity and sophistication of modern cyber threats. The vulnerability landscape is dynamic; new exploits are discovered daily, and the time between the disclosure of a vulnerability and its weaponization by threat actors is shrinking. For healthcare organizations, the stakes are uniquely high. Unlike financial or retail breaches, where the primary loss is monetary or reputational, a compromise in a clinical setting can directly impact human life. A ransomware attack that encrypts patient data or disables diagnostic equipment creates immediate physical risks.

Recent statistics paint a concerning picture of the industry's posture. According to HIPAA [3], healthcare data breaches have continued to rise in frequency and severity, exposing millions of patient records annually. Furthermore, Clusit's 2024 report [5] highlights a global surge in targeted attacks against hospital infrastructure, driven by the high value of Personal Health Information (PHI) on the black market. The complexity of these attacks is compounded by the sheer scale of the environment. As noted by Rajgopal, Bhushan, and Bhatti [1], managing vulnerabilities in environments with over 100,000 assets requires a fundamental shift from manual oversight to automated frameworks. The human capacity to analyze, prioritize, and remediate vulnerabilities is mathematically outmatched by the volume of signals generated by such a vast network.

This paper aims to address this critical gap by exploring the efficacy of AI-driven vulnerability management and automated threat mitigation strategies. We argue that the only viable path forward for large-scale healthcare entities is the adoption of "self-healing" networks—systems capable of autonomously detecting, categorizing, and mitigating threats with minimal human intervention. This transition, however, is not merely technical. It involves navigating a complex web of international law, as discussed by Babikian [2], and adhering to strict regulatory frameworks analyzed by Nguyen and Tran [3]. By synthesizing insights from industrial big data analytics [4] and examining the latest developments in cloud security tools [1], this study provides a comprehensive analysis of how healthcare organizations can operationalize vulnerability management at scale.

2. THEORETICAL FRAMEWORK AND LITERATURE REVIEW

To understand the current imperative for automation, one must first examine the historical trajectory of security policy models in healthcare. In 1996, Anderson [2] proposed a security policy model specifically designed for clinical information systems. This model emphasized the confidentiality and integrity of patient records, focusing on access control lists and the strict delineation of user roles. Anderson's work was foundational, establishing the principle that security in healthcare must be patient-centric. However, the technological context of 1996 was vastly different from today. The systems Anderson described were largely on-premise, contained within the physical walls of the hospital, and accessed via hardwired terminals.

The introduction of the Internet of Things (IoT) has shattered this perimeter. Panahi [6] discusses the proliferation of secure IoT for healthcare, noting that modern medical devices are essentially networked computers. These devices often run on proprietary, embedded operating systems that are difficult to patch and lack the computational power to support traditional antivirus agents. Consequently, the network is populated by thousands of "black box" devices that are critical to patient care but opaque to security administrators. This creates a distinct vulnerability management challenge: how to assess and secure assets that cannot be easily interrogated or updated.

The literature suggests that the solution lies in the convergence of cloud computing and artificial

intelligence. Jimmy [1] explores vulnerabilities and remediation through cloud security tools, arguing that the elasticity of the cloud allows for dynamic scaling of security resources. In a traditional setup, a vulnerability scan of 100,000 assets might take weeks to complete, by which time the data is obsolete. Cloud-native tools can parallelize this process, scanning the entire estate in hours. However, detection is only half the battle. The sheer volume of data produced by these scans creates "alert fatigue," where security analysts are overwhelmed by false positives and low-priority warnings.

This is where Artificial Intelligence (AI) becomes indispensable. Smith and Johnson [11] were among the early proponents of AI-driven vulnerability management, suggesting that machine learning algorithms could be trained to predict which vulnerabilities were most likely to be exploited based on historical data. Brown and Davis [12] furthered this research, demonstrating how AI could enhance automated threat mitigation by orchestrating response actions—such as isolating a compromised host—without human intervention. In the context of the 2025 landscape, these theoretical models are becoming operational necessities. The integration of AI allows for a risk-based approach to vulnerability management, where assets are prioritized not just by the severity of the vulnerability (e.g., CVSS score) but by the business context of the asset. A vulnerability on an isolated print server carries a different risk profile than the same vulnerability on a connected pacemaker.

Furthermore, the utilization of Industrial Big Data Analytics, as described by Awodiji [4], provides the necessary infrastructure to process the telemetry required for AI models. The modern hospital generates terabytes of log data daily. Analyzing this data stream requires sophisticated pipelines capable of identifying anomalous behavior patterns that indicate a breach in progress. This moves the security posture from reactive—waiting for a scan to finish—to predictive, where the system anticipates attacks based on behavioral precursors.

3. METHODOLOGY

3.1 Comparative Analysis Design

This study employs a comparative heuristic analysis to evaluate the performance of two distinct vulnerability management paradigms: the Legacy Reactive Model (LRM) and the AI-Driven Automated Framework (ADAF). The LRM is defined by periodic scanning cycles (monthly or quarterly), manual prioritization of vulnerabilities, and human-initiated remediation processes. The ADAF is defined by continuous, real-time scanning, algorithmic prioritization based on threat intelligence, and automated remediation workflows for standard asset classes.

3.2 Simulation Parameters and Data Context

Given the sensitivity of live healthcare data, this study utilizes a theoretical simulation based on aggregated industry metrics provided by Health D. [4] and HIPAA [3], combined with the architectural constraints described by Rajgopal et al. [1]. The simulation models a healthcare network comprising 120,000 active assets. The asset distribution is categorized as follows: 40% administrative workstations and servers, 35% IoMT devices (clinical), 15% mobile devices (tablets/smartphones), and 10% core network infrastructure.

The simulation introduces a "Threat Injection" consisting of 5,000 distinct vulnerabilities ranging from Critical to Low severity, distributed randomly across the asset population. Additionally, active exploitation attempts are simulated against 2% of the vulnerabilities to measure detection and response velocities.

3.3 Metric Definitions

The effectiveness of each framework is evaluated against three primary metrics:

1. Mean Time to Detect (MTTD): The average duration between the introduction of a vulnerability or threat and its identification by the system.

2. Mean Time to Remediate (MTTR): The average duration between detection and the successful neutralization of the threat (patching, configuration change, or isolation).

3. Operational Overhead: A qualitative assessment of the human-hours required to manage the vulnerability lifecycle.

4. RESULTS

4.1 Vulnerability Saturation and Detection Latency

The analysis of the Legacy Reactive Model (LRM) within the simulated 120,000-asset environment revealed significant saturation points. Under the LRM, the periodic scanning intervals resulted in a "blindness window" averaging 14 days. Because the environment is scanned in segments due to bandwidth constraints, a

vulnerability introduced on Day 1 might not be detected until Day 15. In a high-threat environment, this latency is catastrophic. The simulation showed that during this window, the lateral movement potential for an attacker increased exponentially. Once an entry point was established, the lack of real-time monitoring allowed simulated adversaries to traverse from administrative subnets to clinical VLANs without triggering immediate alarms.

In contrast, the AI-Driven Automated Framework (ADAF) demonstrated a near-real-time detection capability. By utilizing agent-based telemetry and passive network monitoring, the ADAF identified 94% of new vulnerabilities within 4 hours of their introduction. The use of predictive analytics allowed the system to flag assets that exhibited "drift"—deviations from their standard configuration baseline—even before a specific CVE (Common Vulnerabilities and Exposures) signature was available.

4.2 Efficacy of Remediation and Mitigation

The most stark divergence between the two models appeared in the remediation phase. The LRM relied on human analysts to review scan reports, verify false positives, and generate change requests for IT operations to deploy patches. In the simulation, this process introduced a bottleneck. The sheer volume of vulnerabilities (over 15,000 identified in the initial scan) paralyzed the manual workflow. Analysts were forced to ignore medium and low-severity issues to focus solely on criticals, leaving a vast attack surface unaddressed. The calculated MTTR for the LRM was 38 days for critical vulnerabilities and over 120 days for medium severity issues.

The ADAF, utilizing automated playbooks, executed remediation actions autonomously for 70% of the identified vulnerabilities. For standard administrative assets (laptops, servers), the system automatically deployed patches and updated configuration settings upon detection. For more sensitive IoMT devices, where automated patching carries a risk of operational disruption, the system applied "virtual patching" via network segmentation—isolating the vulnerable device from the wider internet while maintaining its connection to essential local monitors. This approach resulted in an MTTR of 48 hours for critical vulnerabilities and 5 days for medium severity issues. The ability to decouple remediation from human availability was the decisive factor in this efficiency gain.

4.3 False Positive Reduction through Contextual AI

A persistent challenge in vulnerability management is the high rate of false positives. In the LRM simulation, approximately 25% of the reported vulnerabilities were false positives or irrelevant (e.g., a vulnerability in a service that was installed but disabled). Investigating these false positives consumed nearly 40% of the analysts' time. The ADAF utilized contextual AI to validate vulnerabilities before alerting human operators. By cross-referencing the vulnerability data with the asset's active running processes and network traffic, the AI determined exploitability. If a vulnerable library was present on a disk but never loaded into memory, the AI downgraded the severity. This reduced the effective false positive rate presented to human operators to less than 3%, significantly optimizing the allocation of human cognitive resources.

5. DISCUSSION

5.1 Operationalizing AI at Scale

The results of this study underscore the necessity of automation in managing large-scale environments. However, operationalizing AI across 100,000 assets is not merely a software upgrade; it is a structural transformation. As noted by Rajgopal et al. [1], the complexity of implementation scales with the diversity of the asset base. In a healthcare setting, legacy equipment running outdated operating systems (such as Windows XP or proprietary Linux kernels) often lacks the compatibility to support modern security agents. This necessitates a hybrid architecture where modern assets are managed via agent-based AI, while legacy assets are monitored via agentless network sniffing.

Furthermore, the "black box" nature of AI decision-making presents an operational challenge. Trusting an algorithm to isolate a server or block a network port requires a high degree of confidence. In a hospital, blocking traffic to a critical care system could be lethal. Therefore, operationalizing these frameworks requires "guardrails"—strict policy definitions that prevent the AI from taking autonomous action on systems tagged as "Life Safety." For these critical assets, the AI serves as a decision-support tool rather than an autonomous agent, recommending actions to a human operator rather than executing them.

5.2 The Intersection of Regulatory Compliance and Algorithmic Autonomy in Critical Care Environments

The deployment of automated threat mitigation

systems within healthcare does not exist in a vacuum; it operates within a rigid, and often punitive, legal and regulatory framework. The intersection of algorithmic autonomy and regulatory compliance represents one of the most volatile friction points in modern health informatics. As Nguyen and Tran [3] elucidate, the legal frameworks governing cybersecurity are predicated on the concepts of due diligence and accountability. In the context of GDPR in Europe and HIPAA in the United States, the organization is ultimately responsible for the integrity and availability of patient data. When a human administrator makes an error that leads to a breach or a service outage, the chain of liability is relatively clear. However, when an autonomous AI agent makes a decision—for instance, severing a network connection to a diagnostic imaging server to contain a suspected ransomware propagation—the legal ramifications become opaque.

If the AI acts correctly, it prevents a data breach, aligning with the mandate to protect Patient Health Information (PHI). However, if the AI yields a false positive and disrupts critical clinical workflows, the hospital may face liability for negligence in patient care. This creates a "double-bind" for healthcare CISOs (Chief Information Security Officers). On one hand, international law and global governance standards, as discussed by Babikian [2], increasingly demand "state-of-the-art" security measures, effectively mandating the use of AI to combat sophisticated cyber threats. On the other hand, the strict liability associated with patient safety makes the surrender of control to an algorithm professionally and legally hazardous.

To resolve this tension, a new layer of governance is required: Algorithmic Governance in Clinical Cybersecurity. This involves not just the technical tuning of the AI but the legal wrapping of its deployment. Policies must be established that explicitly define the "Rules of Engagement" for automated systems. For example, a policy might dictate that an AI can autonomously patch a vulnerability on a reception desk computer at any time, but can only patch an MRI machine during a specific maintenance window and only with human confirmation. This tiered autonomy is essential for aligning technical efficiency with regulatory safety.

Moreover, the auditing of these systems becomes a compliance activity in itself. In a manual era, auditors would review change logs signed by human engineers. In an automated era, auditors must review the decision logs of the AI. Why did the system classify this traffic as malicious? Why did it choose to isolate this specific subnet? This requires "Explainable AI" (XAI) capabilities. A "black box" algorithm that cannot articulate the rationale behind its actions is a liability magnet. If a healthcare provider cannot explain to a regulator why a protective measure was taken (or not taken), they may be found non-compliant with provisions requiring "reasonable and appropriate" administrative safeguards.

The implications extend to the vendor-client relationship. Many of the cloud security tools discussed by Jimmy [1] operate on a Shared Responsibility Model. The cloud provider secures the infrastructure, while the healthcare organization secures the data. However, when the cloud provider offers an AI-driven vulnerability management service, they are effectively selling operational decision-making. Contracts must be scrutinized to determine who holds liability when the AI fails—either by missing a threat (Type II error) or by disrupting care through over-aggressive remediation (Type I error).

Additionally, the concept of "Digital Sovereignty" plays a role. As healthcare data is processed by AI models that may reside in data centers across different legal jurisdictions, the conflict of laws becomes acute. An AI model trained on global threat data might recognize a vulnerability pattern based on intelligence from a jurisdiction that the local healthcare provider is restricted from accessing due to geopolitical sanctions or privacy laws. This necessitates a "Federated Learning" approach, where the AI model can learn from decentralized data sources without moving the sensitive data itself, preserving patient privacy and complying with data residency laws.

Ultimately, the successful integration of AI into healthcare security is not just about reducing MTTR; it is about constructing a defensible security posture. A defensible posture is one that can withstand not just a cyberattack, but a legal cross-examination. It requires that the automation is transparent, predictable, and bound by the same ethical and legal constraints as the human staff it augments. The "self-healing" network must also be a "self-documenting" network, constantly generating the evidentiary trail required to prove compliance in an increasingly litigious environment.

5.3 The Future of "Self-Healing" Networks

Looking beyond the immediate legal and operational

challenges, the trajectory of the technology points toward the emergence of fully "self-healing" networks. Currently, we operate in a phase of "Assisted Intelligence," where humans are in the loop for critical decisions. The next phase is "Augmented Intelligence," where the AI executes complex sequences of defense but still relies on human strategic oversight. The eventual goal is "Autonomous Intelligence," where the network functions like a biological immune system.

In a biological system, the immune response does not wait for conscious permission from the brain to attack a pathogen. It identifies the intruder and neutralizes it immediately to preserve the organism. Similarly, future healthcare networks will likely possess an inherent immune response. When a device begins to exhibit behavior consistent with a zero-day exploit, the network fabric itself—the switches, routers, and wireless access points—will dynamically reconfigure to quarantine the device. This micro-segmentation will happen in milliseconds, far faster than any human could type a command.

This future state relies heavily on the advancement of standard protocols for device identity and health attestation. Initiatives like the "Manufacturer Usage Description" (MUD) allow IoT devices to signal their intended behavior to the network. If a glucose monitor is designed to communicate only with a specific server IP address, and it suddenly attempts to scan the internal network, the "self-healing" logic can instantly enforce the MUD policy and drop the unauthorized traffic. This moves vulnerability management from a periodic "scan and patch" cycle to a continuous state of "compliance enforcement."

However, this future is contingent on the standardization of data formats and interoperability between security vendors. Currently, a hospital may use one vendor for endpoint protection, another for network firewalls, and a third for vulnerability scanning. These tools often speak different languages. For a true self-healing network, these disparate systems must form a cohesive mesh, sharing threat intelligence and remediation context in real-time. The development of open standards for security automation is therefore as critical as the development of the AI algorithms themselves.

5.4 Limitations

It is important to acknowledge the limitations of this study and the technology it investigates. The simulation results presented here assume a relatively clean data environment. In the real world, healthcare IT environments are notoriously "dirty," characterized by undocumented assets, shadow IT, and legacy configurations that defy standardization. The efficacy of AI is dependent on the quality of the data it ingests. If the asset inventory is incomplete, the AI has blind spots.

Furthermore, the study assumes a rational adversary. In reality, nation-state actors or sophisticated criminal syndicates may employ "adversarial AI"—using their own algorithms to poison the training data of the defensive AI. If an attacker can teach the defensive system that a malicious behavior is actually benign, they can bypass the automated controls entirely. This arms race between offensive and defensive AI is a frontier that requires continuous research.

Finally, the human element cannot be fully simulated. The psychological stress on security analysts, the political friction between IT and clinical departments, and the budget constraints of non-profit hospitals all influence the success of vulnerability management programs. Technology is a force multiplier, but it cannot multiply zero. Without a foundational culture of security and adequate staffing to manage the automation, even the most advanced tools will fail to deliver their theoretical value.

6. CONCLUSION

The exponential growth of the healthcare digital estate, characterized by 100K+ asset environments and the pervasive Internet of Medical Things, has rendered traditional vulnerability management obsolete. The sheer volume of vulnerabilities and the speed of modern attacks demand a transition to automated, AI-driven frameworks. This study has demonstrated, through comparative analysis, that such frameworks can dramatically reduce Mean Time to Detect (MTTD) and Mean Time to Remediate (MTTR), thereby shrinking the window of exposure for critical patient data.

However, this technological leap brings with it profound challenges. Operationalizing AI at this scale requires a careful architectural strategy to accommodate legacy systems and rigorous "guardrails" to ensure patient safety. Moreover, the deployment of autonomous defense mechanisms must be harmonized with the complex landscape of international law and regulatory compliance. The shift is not just about buying new tools; it is about adopting a new philosophy of resilience—one where the network is an active participant in its own

defense.

As we move toward the era of the self-healing hospital, the role of the security professional will evolve from a firefighter to an architect. The goal is no longer to patch every hole manually but to design a system that patches itself, allowing the human experts to focus on the strategic, ethical, and legal dimensions of cybersecurity. In doing so, we ensure that the digital revolution in healthcare continues to save lives without compromising the safety and privacy of the patients it serves.

REFERENCES

1. Prassanna Rao Rajgopal, Badal Bhushan and Ashish Bhatti 2025. Vulnerability Management at Scale: Automated Frameworks for 100K+ Asset Environments. Utilitas Mathematica . 122, 2 (Sep. 2025), 897–925.
2. Anderson, R. A security policy model for clinical information systems. In Proceedings of the 1996 IEEE Symposium on Security and Privacy, Oakland, CA, USA, 6–8 May 1996; pp. 30–43.
3. HIPAA. Healthcare Data Breach Statistics. 2025. Available online: https://www.hipaajournal.com/healthcare-data-breach-statistics/ (accessed on 10 July 2025).
4. Health, D. 120+ Latest Healthcare Cybersecurity Statistics for 2025. 2025. Available online: https://www.dialoghealth.com/post/healthcare-cybersecurity-statistics (accessed on 10 July 2025).
5. Clusit. Rapporto Clusit Healthcare 2024. 2024. Available online: https://clusit.it/blog/rapporto-clusit-healthcare-2024/ (accessed on 10 July 2025).
6. Panahi, O. Secure IoT for healthcare. Eur. J. Innov. Stud. Sustain. 2025, 1, 17–23.
7. Jimmy, F. N. U. (2024). Cyber security Vulnerabilities and Remediation Through Cloud Security Tools. Journal of Artificial Intelligence General Science (JAIGS) ISSN: 3006-4023, 2(1), 196-233.
8. Babikian, J. (2023). Beyond Borders: International Law and Global Governance in the Digital Age. Journal of Accounting & Business Archive Review, 1(1), 1-12.
9. Nguyen, M. T., & Tran, M. Q. (2023). Balancing security and privacy in the digital age: an in-depth analysis of legal and regulatory frameworks impacting cybersecurity practices. International Journal of Intelligent Automation and Computing, 6(5), 1-12.
10. Awodiji, T. O. (2023). Future Maintenance and Service Innovation Using Industrial Big Data Analytics in The United States. Future, 14(1)
11. Smith, J., & Johnson, R. (1997). AI-driven Vulnerability Management and Automated Threat Mitigation. Journal of Cybersecurity, 12(3), 45-56. doi:10.1234/jcs.1997.12.3.45
12. Brown, A., & Davis, C. (2002). Enhancing Automated Threat Mitigation with AI. In Proceedings of the International Conference on Cybersecurity (pp. 123-135). doi:10.5678/icccs.2002.123