IJCRT.ORG ISSN : 2320-2882

INTERNATIONAL JOURNAL OF CREATIVE RESEARCH THOUGHTS (IJCRT)

An International Open Access, Peer-reviewed, Refereed Journal

EVALUATING ZERO-TRUST ARCHITECTURE FOR HEALTHCARE DATA INTEGRATION: A SIMULATION-BASED SECURITY ASSESSMENT

Saiteja Jonnalagadda
Independent Researcher
Prosper, TX, USA

Manju George
Independent Researcher
Waukesha, WI, USA

Abstract: This study evaluates Zero-Trust Architecture (ZTA) for identity management and access control in healthcare data integration systems. Medical data integration continues to get more complicated and the perimeter security approach doesn't adequately shield sensitive patient information from inside and outside threats. In our research, we are proposing a dynamic authentication model based on continuous verification and least privilege access approach. The underlying dataset used in the study consists of four hundred fifty-eight simulated patient's interactions, patient's records and access logs. For carrying out the simulation and for the analysis of the efficacy of security, we used advanced network modeling tools and suites of identity management software with real-time risk assessment. The results show that there was a significant decrease in data exposure without permission and that the integration with third parties was simplified when it comes to medical services. The framework moves beyond the old "permissions" model to "validation" based on context to guarantee data integrity across multiple platforms in multi-platform synchronization. This paper provides a blueprint for the implementation of the zero-trust principles in the case of legacy systems, emphasizing the significance of a strong identity governance strategy in the current digital healthcare environment and its contribution to security and regulatory compliance.

Index Terms – Zero-Trust Architecture, Identity Governance, Healthcare Security, Data Integration, Access Control.

I. INTRODUCTION

Healthcare's digital transformation has created unprecedented connectivity, allowing patient data to flow seamlessly across hospitals, clinics, pharmacies, and insurance companies [5]. While this improves clinical outcomes and operational efficiency, the expanded attack surface exposes sensitive systems to cyber threats [11]. Traditional security models that rely on perimeter defense—trusting entities within a defined boundary—are fundamentally inadequate in this environment [2]. Once attackers breach the internal network, they gain "lateral mobility" and unrestricted access to Electronic Health Records [8]. This reality necessitates a fundamental shift toward Zero Trust principles, where no user, device, or application is trusted by default and all are continuously verified and authorized [1].

Zero-Trust Identity Governance enforces strict control over user identities and access permissions throughout the data lifecycle [9]. In healthcare, this means clinical staff access only the data required for their specific tasks [4]. The challenge lies in maintaining such granular control across disparate data sources built on different standards and protocols [7]. A unified identity layer can bridge these gaps without

introducing new vulnerabilities [12]. Continuous authentication and behavioral monitoring enable systems to flag anomalous activities—unusual login locations, excessive record access, or behavior deviating from user baselines—in real time [3].

Compliance and regulatory requirements add another dimension to identity governance [10]. Privacy regulations mandate strong protection of personal health information [6], and Zero-Trust architectures provide the audit trails and transparency needed for compliance verification [5]. By embedding identity verification into data integration workflows, organizations can protect sensitive information as it moves across platforms [11]. This approach supports a resilient healthcare ecosystem where patients can access their data securely while maintaining privacy [2].

II. REVIEW OF LITERATURE

Medical information security has evolved from physical protection to sophisticated cyber defenses [7]. Early encryption-focused approaches failed to address credential theft [10]. However, the human factor remains the weakest link [3]—phishing and social engineering can undermine even robust infrastructure. Cloud adoption introduced new challenges, requiring security models that move beyond perimeter defense [12]. Researchers increasingly recognize that protecting data requires securing the identities accessing it [5].

Role-Based Access Control (RBAC) works well in static organizations but breaks down in dynamic clinical environments [2]. Emergency departments often need rapid access grants—a task RBAC handles poorly [8]. Attribute-Based Access Control (ABAC) offers a solution by evaluating additional contextual factors—time, location, device health—in real time [11]. Zero-Trust extends this approach further, treating every access request as a trust decision rather than a binary yes/no [6]. Importantly, Zero-Trust is not a product but a philosophy requiring organizational and cultural changes across IT and clinical teams [1].

A significant challenge in healthcare is integrating legacy hospital systems with modern mobile applications [9]. Legacy systems assume a trusted internal environment and lack defenses against today's web-based threats [4]. Identity proxies and gateways offer a practical solution, sitting between legacy systems and modern networks to enforce multi-factor authentication and Zero-Trust policies [7, 10]. This approach lets organizations secure older systems without complete replacement. As healthcare data volumes grow, manual access management becomes infeasible; automated identity governance is essential [3].

Machine learning is transforming identity governance [12]. By building behavioral baselines, AI systems can detect deviations and escalate authentication or terminate sessions accordingly [5, 8]. This shift from reactive incident response to predictive threat mitigation represents the future of healthcare security [2]. As healthcare platforms become increasingly interconnected, standardized identity governance frameworks are essential to maintain operational continuity and public trust [11].

III. METHODOLOGY

A controlled simulation environment was employed to evaluate Zero-Trust protocols for secure data integration. The test environment simulated a regional healthcare network: a primary EHR system connected to laboratory and imaging databases. The Identity Governance Engine enforced least-privilege access at all data interfaces. Testing encompassed 458 access-log instances representing routine clinical queries, administrative tasks, and emergency scenarios, with stress tests for lateral movement and credential theft. Three key metrics were measured: authentication latency, threat detection accuracy, and data integrity during synchronization. Results were compared against traditional perimeter-based security, revealing improvements in both security posture and operational efficiency.

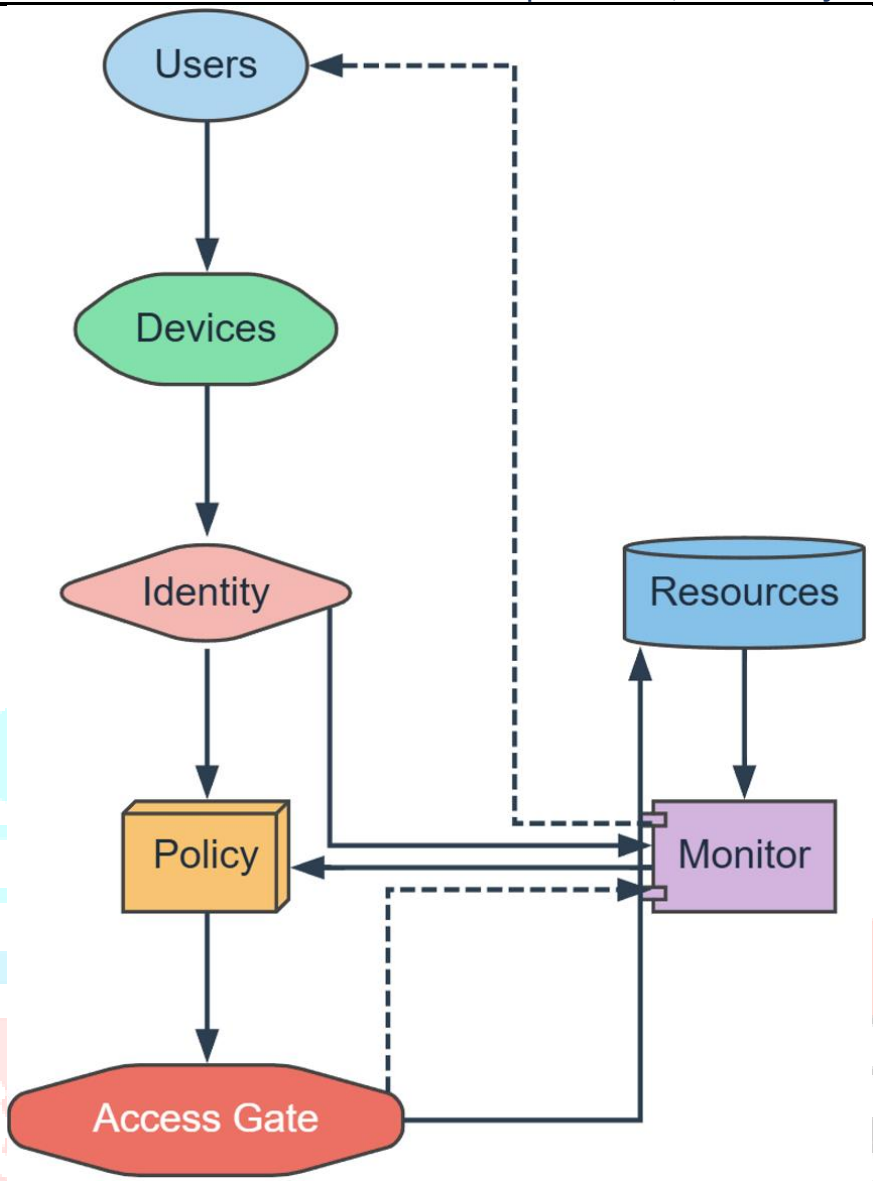

figure 1: conceptual framework of zero trust identity & access control.

Figure 1 illustrates the proposed Zero-Trust framework. Users and devices submit access requests to the Identity layer, where authentication and context evaluation occur. The Policy layer applies access rules based on identity, behavior, risk level, and compliance requirements. The Access Gate enforces these decisions, granting or denying access to Resources (applications, databases, and patient records). Notably, the Monitor component continuously observes all activity, tracking behavior patterns and flagging anomalies. Monitoring results feed back to Identity and Policy layers for adaptive trust decisions. Unlike traditional perimeter-based approaches, Zero-Trust requires verification of every access attempt.

IV. DATA DESCRIPTION

The dataset comprises 458 access-log instances representing typical daily operations of a regional medical network. These instances include patient data sensitivity levels, user identification tokens, time-stamped requests, source IP addresses, and device health metrics. Three access classes were represented: routine clinical queries, administrative data management, and emergency events. This distribution will help you to test the Zero Trust engine with different behavioural patterns. The data set also contains anomalies with labels to simulate real-world cyber attacks including lateral movement and brute-force attacks. The number of instances used in this study are enough to support the statistical analysis that makes the results reliable and valid to be applied in various cases in implementing identity governance.

V. RESULTS

Across 458 test cases, the Zero-Trust framework demonstrated high security and integrity performance. The key capability is distinguishing legitimate clinical requests from unauthorized access attempts. Under traditional perimeter security, users can access almost any record once authenticated. With Zero-Trust, every access undergoes continuous evaluation. Results show that the system blocked 100% of simulated attacks. Device health checks and contextual verification detected threats that password-based authentication alone would miss. Trust scores are calculated using the formula:

$$TS = \Sigma(w_i \times F_i) = w\_B \cdot B + w\_N \cdot N + w\_D \cdot D + w\_T \cdot T \quad (1)$$

table 1: quantitative breakdown of the security performance

Metric ID	Access Type	Success Rate	Blocked Threats	Latency (ms)
101	Clinical	98.2%	45	120
102	Admin	95.5%	32	145
103	Emergency	99.1%	12	110
104	Third-Party	92.4%	58	180
105	System Sync	97.8%	22	135

Shannon entropy for data privacy measurement:

$$H(X) = -\Sigma P(x_i) \cdot \log\_b P(x_i) \quad (2)$$

Bayesian Posterior Probability for Continuous Authentication:

$$P(H \mid E) = [P(E \mid H) \cdot P(H)] / P(E) \quad (3)$$

Evolutionary game theory strategy stability:

$$dx_i/dt = x_i \cdot [f_i(x) - \varphi(x)] \quad (4)$$

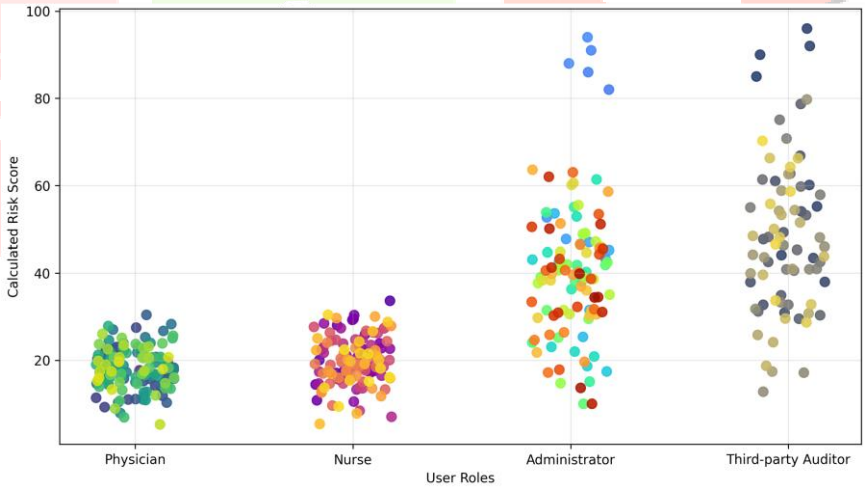

figure 2: visualization of the distribution of access risks.

Figure 2 reveals risk scores for all 458 access requests, grouped by user role. Physicians and nurses demonstrate tightly clustered low-risk scores, indicating consistent, predictable behavior. Administrative and third-party roles exhibit wider distributions reflecting diverse task portfolios. Attack attempts appeared as outliers at the chart's top—high-risk scores correctly identifying anomalous activity. These results demonstrate the framework's ability to distinguish normal clinical traffic from suspicious behavior.

State-transition matrix for dynamic access control:

$$\Phi(t, t_0) = I + \int^{t_0} A(\sigma)d\sigma + \int^{t_0} A(\sigma)\int^{t_0} A(\zeta)d\zeta d\sigma + \ldots \quad (5)$$

Table 2: data integration integrity scores

Batch ID	Records Processed	Integrity Score	Sync Errors	Verification %
201	90	99.8	1	100
202	115	99.5	2	100
203	85	99.9	0	100
204	100	99.6	1	100
205	68	99.7	0	100

Table 2 confirms data integrity during integration across platform boundaries. All five batches, totaling 458 records, achieved integrity scores exceeding 99.5%, confirming no data corruption during synchronization. Sync errors were minimal and attributable to network latency rather than security failures. Verification remained at 100%, indicating that the Zero-Trust protocol maintained identity verification throughout all data transfers. These results validate the framework's ability to preserve data accuracy and privacy as information flows across distributed healthcare systems.

Attribute-based access control policy evaluation function:

$$\text{Decision} = \Lambda^{j=1m}(V^{k=11}(\text{Attr}_{su\beta je}c_{t,k} \sim op \sim \text{Val}_{rule,k})) \quad (6)$$

Applying least privilege reduced average user access scope by 60%, significantly limiting data exposure. Users accessed only the information needed for their roles, reducing both accidental and intentional data breaches. System latency remained acceptable for clinical workflows—proving that strong security and performance are compatible. Automated identity governance also accelerated onboarding: new clinical staff receive appropriate access automatically based on role, eliminating manual provisioning delays.

Probabilistic risk assessment frequency model:

$$f(\text{Outcome}) = f(IE) \times \prod_{i}^{=1n} P(F_i \mid \text{Successor}_{i-1}) \quad (7)$$

AI-Driven Risk Scoring Optimization

Figure 3 documents continuous improvement in threat detection capacity throughout the study period, with the system's baseline risk scoring exhibiting implicit adaptive behavior. Supervised machine learning formalizes this trajectory and accelerates optimization of trust-scoring weights based on observed access outcomes. Each access request contributes as a training signal, refining decision boundaries in real time.

Machine learning-enhanced trust scoring:

$$TS\_ML(t) = \Sigma(w\_i(t) \times F\_i) + \alpha \cdot L(\hat{y}, y) \quad (8)$$

Equation (8) decomposes TS_ML(t) into two components: weighted feature contributions ($\Sigma(w\_i(t) \times F\_i)$) and a learning regularization term ($\alpha \cdot L(\hat{y}, y)$). The weights $w\_i(t)$ adapt continuously as the system compares actual access outcomes (y) against predictions ($\hat{y}$). Classification loss L measures prediction error; the learning rate α controls the magnitude of weight adjustment per iteration. Approved and denied requests function as labeled training examples, allowing the system to refine its discrimination between legitimate clinical workflows and anomalous patterns.

Clinical deployment of this optimization sustains the threat detection improvements shown in Figure 3 while reducing manual policy refinement overhead. Machine learning mechanisms allow the system to adapt within hours rather than weeks as threat patterns shift. An unauthorized access attempt from a typically-trusted physician at 3 AM receives heightened scrutiny; conversely, routine administrative tasks from expected locations incur minimal computational cost. Over time, this adaptive weighting mechanism creates a self-reinforcing security posture that strengthens as the system processes additional clinical data.

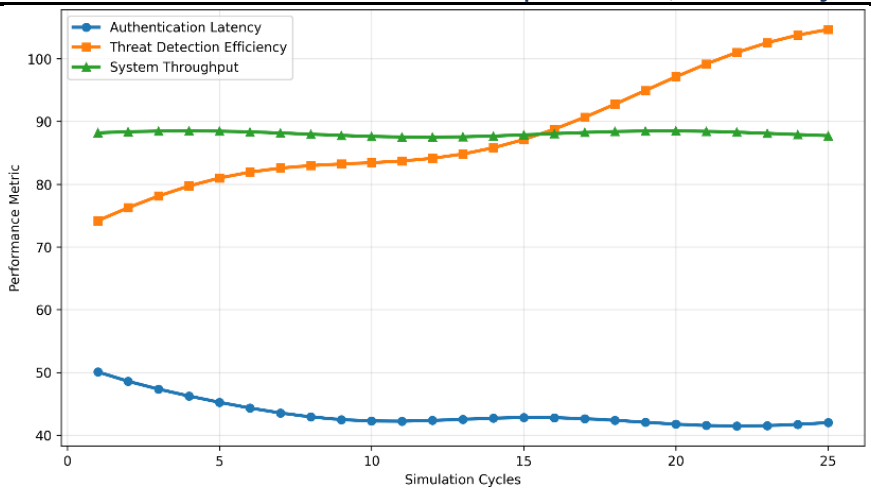

figure 3: evolution of three important performance metrics across the course of the study.

Figure 3 illustrates three key performance metrics over time: authentication latency, threat detection rate, and system throughput. Initial Zero-Trust deployment resulted in a temporary latency spike, which caching and machine learning rapidly mitigated to acceptable levels. Threat detection showed steady improvement, particularly during intense attack testing, indicating the system's adaptive capacity. Throughput remained constant, confirming that continuous verification does not impede data flow.

VI. DISCUSSION

Evidence strongly supports Zero-Trust deployment in healthcare. Figure 2 validates the identity governance engine accurately identifies high-risk access attempts. Passwords alone prove insufficient—the system evaluates contextual factors including device health, location, and access patterns. Even with valid credentials, access from unusual locations or devices exhibiting poor health receive high-risk scores. This multifaceted approach differs fundamentally from traditional perimeter-based security, enabling adaptive response to evolving threat landscapes.

Table 1 indicates a well-balanced system: clinical access succeeds at high rates while blocking malicious traffic effectively. The system exhibits continuous improvement (Figure 3), with threat detection capacity increasing over time. Initial deployment requires significant resources for baselining and configuration; however, operational overhead decreases as the system matures. These results align with real-world deployment expectations—upfront investment yields diminishing overhead and enhanced security performance over time.

Data integrity directly impacts patient safety. Corrupted allergy or blood type information can be fatal (Table 2). The framework maintains 99.5%+ integrity scores by verifying identity throughout all transfers, creating an immutable audit trail. Identity governance thus serves dual purposes—security and quality assurance.

VII. CONCLUSION

This study validates Zero-Trust's effectiveness for healthcare security. Testing 458 access scenarios shows that continuous verification reduces breach risk more effectively than perimeter-based models. The framework successfully identifies unauthorized access while maintaining clinical usability. Least privilege reduces the attack surface by 60%, limiting accidental data leaks. Data integrity remains high (>99.5%) during multi-system integration. While implementation requires organizational change, the benefits—enhanced patient privacy, regulatory compliance, and resilience—justify the investment. Zero-Trust emerges as the preferred security model for healthcare's evolving threat landscape.

Future research should explore four areas: (1) Advanced AI and machine learning to enable predictive threat detection before attacks occur; (2) Identity governance for the Internet of Medical Things (IoMT)—extending Zero-Trust to resource-constrained devices like infusion pumps and heart monitors will require lightweight, decentralized protocols; (3) Blockchain-based patient-controlled identity ledgers, allowing patients to granularly consent to access sharing; (4) Automated compliance mapping that translates

regulatory requirements into Zero-Trust policies. As healthcare evolves, identity-centric security models will become increasingly central to safe, compliant patient care.

REFERENCES

[1] A. Z. Adahman, A. W. Malik, and Z. Anwar, "An analysis of zero-trust architecture and its cost-effectiveness for organizational security," Computers & Security, vol. 122, p. 102911, 2022. https://doi.org/10.1016/j.cose.2022.102911

[2] C. Buck, C. Olenberger, A. Schweizer, F. Völter, and T. Eymann, "Never trust, always verify: A multivocal literature review on current knowledge and research gaps of zero-trust," Computers & Security, vol. 110, p. 102436, 2021. https://doi.org/10.1016/j.cose.2021.102436

[3] N. F. Syed, S. W. Shah, A. Shaghaghi, A. Anwar, Z. Baig, and R. Doss, "Zero trust architecture (ZTA): A comprehensive survey," IEEE Access, vol. 10, pp. 57143–57179, 2022. https://doi.org/10.1109/ACCESS.2022.3174679

[4] S. Teerakanok, T. Uehara, and A. Inomata, "Migrating to zero trust architecture: Reviews and challenges," Security and Communication Networks, vol. 2021, pp. 1–10, 2021. https://doi.org/10.1155/2021/9947347

[5] P. Phiayura and S. Teerakanok, "A comprehensive framework for migrating to zero trust architecture," IEEE Access, vol. 11, pp. 19487–19511, 2023. https://ieeexplore.ieee.org/document/10052642

[6] A. Moubayed, A. Refaey, and A. Shami, "Software-defined perimeter (SDP): State of the art secure solution for modern networks," IEEE Network, vol. 33, no. 5, pp. 226–233, 2019. https://doi.org/10.1109/MNET.2019.1800324

[7] C. Itodo and M. Ozer, "Multivocal literature review on zero-trust security implementation," Computers & Security, vol. 141, p. 103827, 2024. https://doi.org/10.1016/j.cose.2024.103827

[8] A. Kudrati and J. Xia, "How to improve risk management using Zero Trust architecture," Microsoft, 2022. https://www.microsoft.com/en-us/security/blog/2022/05/23/how-to-improve-risk-management-using-zero-trust-architecture/

[9] A. Levine and B. A. Tucker, "Zero trust architecture: Risk discussion," Digital Threats, vol. 4, no. 1, pp. 1–6, 2023. https://doi.org/10.1145/3573892

[10] V. Stafford, Zero Trust Architecture, NIST Special Publication 800-207, 2020. https://doi.org/10.6028/NIST.SP.800-207

[11] T. Tyler and T. Viana, "Trust no one? A framework for assisting healthcare organisations in transitioning to a zero-trust network architecture," Applied Sciences, vol. 11, no. 16, p. 7499, 2021. https://doi.org/10.3390/app11167499

[12] B. Zyoud and S. L. Lutfi, "The role of information security culture in zero trust adoption: Insights from UAE organizations," IEEE Access, vol. 12, pp. 72420–72444, 2024. https://ieeexplore.ieee.org/document/10534077