Radhika Kande[1], Chaithanya Kotla[2], Krishna kanth Thottempudi[3]
[1]Sagarsoft Inc, USA
[2]Devops and Cloud lead, State of Maryland, USA
[3]Infosys Limited, USA

Received: 11-07-2022; Revised: 29-08-2022; Accepted: 05-10-2022

Zero Trust Architecture with ML-Driven Behavioral Analytics for Healthcare ERP on Microsoft Azure

Abstract
Healthcare ERP systems on Microsoft Azure require access controls that can evaluate not only user identity, but also session context, privilege behavior, transaction activity, and network movement. This article presents a Zero Trust architecture enhanced by ML-driven behavioral analytics for adaptive access enforcement in healthcare ERP workloads. The framework integrates Azure AD sign-ins, MFA status, conditional access decisions, RBAC assignments, privileged role activation, device compliance, ERP transaction logs, NSG flow records, private endpoint activity, Sentinel alerts, and runtime telemetry. ML models generate user behavior risk scores, privilege deviation scores, session trust probabilities, ERP transaction anomaly scores, and network behavior threat scores to support continuous verification. Results show that identity trust accuracy, privilege deviation detection, and ERP transaction anomaly recognition improve across Zero Trust verification cycles, while healthcare ERP access scenarios reveal distinct variations in user behavior risk, session trust, and network threat exposure. The study concludes that Zero Trust becomes more effective for healthcare ERP security when Azure-native controls are combined with behavioral intelligence and ERP-specific transaction monitoring.
Keywords: Zero Trust architecture, healthcare ERP security, Microsoft Azure, behavioral analytics, adaptive access control, ERP transaction anomaly detection, identity risk scoring.

1. Introduction

Healthcare ERP systems on Microsoft Azure operate at the intersection of clinical administration, finance, procurement, insurance processing, workforce scheduling, inventory control, and patient-facing operational workflows. These systems contain sensitive business and healthcare data, and their access patterns are often more complex than ordinary cloud applications because users may include physicians, nurses, billing staff, pharmacy teams, procurement officers, administrators, external auditors, and emergency-support personnel. Zero Trust architecture is increasingly relevant in this context because it removes implicit trust from internal users, applications, devices, and network locations, requiring continuous verification before access is granted or maintained [1]. In healthcare ERP environments, this principle becomes critical because an authenticated user may still behave abnormally after login.

Zero Trust research emphasizes identity verification, least-privilege access, segmentation, device context, continuous monitoring, and adaptive access enforcement. Broader surveys on Zero Trust architecture show that static perimeter models are insufficient when users, workloads, and applications are distributed across cloud and hybrid environments [2]. Intelligent Zero Trust work further highlights the role of machine learning in strengthening adaptive access decisions when user behavior and contextual risk change over time [3]. These ideas are especially suitable for healthcare ERP systems, where a user's role, device, session context, transaction behavior, and access timing can all influence whether a request should be trusted. A healthcare ERP access decision must therefore consider both identity validity and behavioral legitimacy.

The main limitation in many ERP security models is that access control remains role-centered rather than behavior-centered. A billing officer may have permission to access claims records, but unusual after-hours bulk access may still indicate misuse. A privileged administrator may have authority to modify configuration settings, but repeated role changes from an unfamiliar location may indicate account compromise. Behavioral analytics in Zero Trust environments has shown that identity-based threat segmentation can become more effective when access decisions use real-time behavioral risk scores instead of static rules alone [4]. This reveals a gap in healthcare ERP protection: legitimate permissions must be continuously checked against expected behavior.

This problem matters because healthcare ERP compromise can affect clinical continuity, financial integrity, regulatory compliance, and patient-data confidentiality. A weak access decision may allow abnormal finance postings, unauthorized procurement changes, sensitive record browsing, payroll manipulation, or misuse of clinical supply workflows. ML-enhanced Zero Trust and user/entity behavior analytics models have been proposed to strengthen SIEM, SOAR, and adaptive security response by combining identity context with behavioral risk interpretation [5]. Healthcare ERP systems require the same type of adaptive trust logic, but the analytics must be aligned with ERP-specific

transaction behavior and Azure-native identity controls. The goal is not only to block attackers, but to prevent trusted accounts from becoming uncontrolled risk channels.

This article proposes a Zero Trust architecture enhanced by ML-driven behavioral analytics for healthcare ERP systems deployed on Microsoft Azure. The framework combines Azure AD identity signals, conditional access, MFA status, RBAC assignments, privileged role activation, device compliance, ERP module activity, transaction logs, NSG flow records, Sentinel alerts, and policy state. Machine learning models generate user behavior risk scores, privilege deviation scores, session trust probabilities, ERP transaction anomaly scores, and network behavior threat scores. The study evaluates how behavioral analytics can strengthen continuous verification, least-privilege enforcement, ERP transaction monitoring, and adaptive access decisions in Azure-hosted healthcare ERP environments.

2. Methodology

The architecture is designed around a behavioral trust loop rather than a one-time access decision. When a healthcare ERP user attempts access, the system evaluates identity, device, session, role, transaction intent, network path, and historical behavior before deciding whether to allow, restrict, challenge, or block the request. Machine learning research in identity and access management shows that authentication, authorization, and auditing can be strengthened when user behavior and access patterns are analyzed continuously [6]. This principle is applied to Azure healthcare ERP access by treating every session as a changing trust condition. The system therefore verifies both who the user is and whether the current action matches expected ERP behavior.

Identity verification begins with Azure AD sign-ins, MFA completion, conditional access result, sign-in risk, service principal use, and managed identity behavior. These signals are transformed into identity-risk features such as failed login rate, unusual login location, impossible travel, unfamiliar device, atypical time window, and abnormal token activity. AI-powered Zero Trust access evaluation using behavioral fingerprinting has shown that user access decisions can become more adaptive when they are linked to normal and abnormal behavioral patterns [7]. In the proposed model, identity verification is not finished after authentication. The identity state continues to influence session trust throughout the ERP interaction.

Least-privilege monitoring focuses on RBAC assignments, privileged role activation, emergency access roles, role elevation frequency, ERP module permissions, and administrative actions. A healthcare ERP user may require access to specific financial, inventory, billing, or clinical-administration functions, but role expansion beyond normal behavior can indicate misuse. Behavioral Zero Trust access-control work on insider-threat scenarios shows that access decisions should account for deviations in user behavior, role use, and contextual activity [8]. The framework therefore generates a privilege deviation score that increases when a user activates unusual roles, performs unexpected administrative operations, or

accesses modules outside expected responsibility. This score is used to trigger step-up authentication, temporary restriction, or security review.

ERP transaction behavior is analyzed separately because ordinary identity monitoring does not fully capture healthcare business risk. The model evaluates transaction frequency, module type, posting value, record-access volume, approval sequence, failed transaction attempts, bulk exports, and unusual workflow combinations. Identity and access management adoption in ERP systems has been linked to governance quality, role-based control, and security accountability [9]. In this architecture, ERP transaction analytics extends those controls by checking whether permitted users are performing expected actions. A finance posting, inventory adjustment, patient-account update, or procurement approval is therefore evaluated not only for authorization, but also for behavioral consistency.

Network and session trust are assessed through device compliance, session location, Azure private endpoint usage, NSG flow behavior, firewall events, VPN status, and service path consistency. Cloud database access-control research using Zero Trust and machine learning shows that adaptive access decisions benefit from combining identity, behavioral, and infrastructure-level access signals [10]. The proposed framework follows this direction by linking session trust with network behavior. For example, a normal ERP transaction from a compliant managed device may be allowed, while the same transaction from an unmanaged device over an unusual network route may require step-up verification or limited access.

The Zero Trust behavioral control map used by the framework is summarized in Table 1. The table is positioned here because it links Azure control signals with healthcare ERP risk indicators and the machine learning outputs used for adaptive trust decisions. It also shows that the architecture is not limited to identity verification. It covers privilege behavior, device trust, ERP transaction behavior, network segmentation, and continuous verification as connected Zero Trust layers.

Table 1. Zero Trust Behavioral Control Map for Healthcare ERP Workloads on Microsoft Azure

Zero Trust Layer	Azure Control Signal	Healthcare ERP Risk Indicator	ML Behavioral Analytics Output
Identity verification	Azure AD sign-ins, MFA status, conditional access	Abnormal user login, failed authentication, risky session	User behavior risk score
Least-privilege access	RBAC assignments, privileged role activation	Excessive ERP access, privilege escalation, role misuse	Privilege deviation score
Device and session trust	Device compliance, session context, location signal	Untrusted device, impossible travel, unmanaged access	Session trust probability
ERP transaction behavior	ERP access logs, finance/clinical module activity	Abnormal posting, unusual record access, bulk transaction activity	ERP transaction anomaly score
Network segmentation	NSG flow logs, private endpoint activity, firewall events	Lateral movement, external exposure, unusual service path	Network behavior threat score

Continuous verification	Sentinel alerts, policy state, runtime telemetry	Repeated anomaly, unresolved exposure, compliance drift	Adaptive access decision trigger

The behavioral analytics model uses supervised and unsupervised learning components. Supervised classification is applied where labeled access events, policy violations, or known abnormal transaction cases are available. Unsupervised anomaly detection is used for rare ERP transaction patterns, unknown insider-risk behavior, and unusual cross-module activity. Each model output is normalized into a trust-control score and passed to the decision engine. The decision engine applies four responses: allow normal access, require step-up verification, restrict the ERP function, or block and escalate. This allows the architecture to enforce Zero Trust without treating every anomaly as a full security incident.

Evaluation is performed using simulated healthcare ERP access scenarios on Azure. The scenarios include normal clinical access, routine finance posting, remote login, privileged maintenance, emergency access, anomalous bulk record access, unusual role activation, and suspicious network behavior. The model is compared with static RBAC, conditional-access-only enforcement, and rule-based ERP monitoring. Evaluation metrics include identity trust accuracy, privilege deviation detection, ERP transaction anomaly recognition, user behavior risk scoring, session trust probability, network behavior threat scoring, false challenge rate, and adaptive access decision accuracy. These metrics are selected because healthcare ERP Zero Trust must be judged by both security strength and operational usability.

3. Results and Discussion

Behavioral verification improves the Zero Trust model by converting Azure identity and ERP transaction signals into measurable trust evidence. Figure 1 shows that identity trust accuracy increases across verification cycles as the model learns normal login timing, device context, MFA behavior, and session patterns. Privilege deviation detection improves when role activation and administrative actions are evaluated against historical behavior rather than fixed permission lists. ERP transaction anomaly recognition also rises because the model learns which transaction sequences, module combinations, posting volumes, and record-access patterns are unusual for each user group. This result confirms that healthcare ERP access control becomes stronger when identity and transaction behavior are analyzed together.

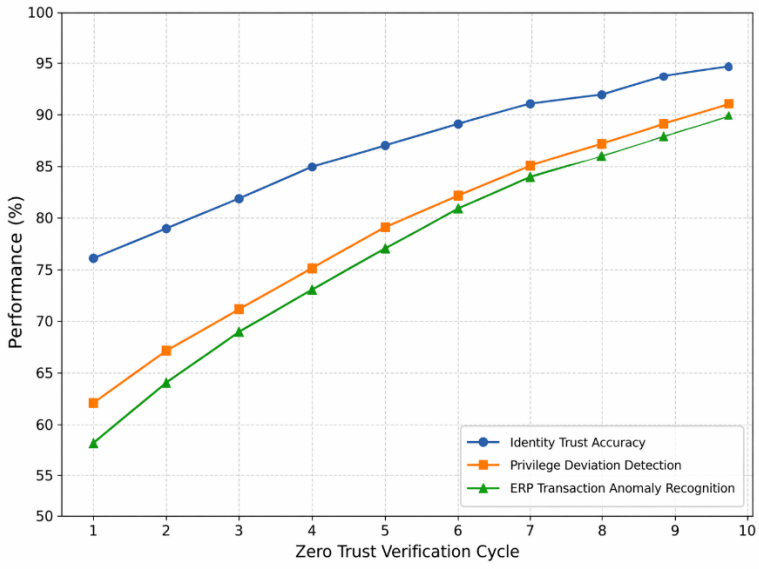

Figure 1. Identity Trust Accuracy, Privilege Deviation Detection, and ERP Transaction Anomaly Recognition Across Zero Trust Verification Cycles

The strongest improvement appears when the system connects role behavior with ERP activity. A user may authenticate successfully and hold a valid role, but abnormal transaction behavior can still indicate risk. For example, a finance user performing a high volume of claims adjustments outside normal hours may require different treatment from a clinician accessing patient scheduling records during a shift. The ML model supports this distinction by generating separate scores for identity trust, privilege deviation, and ERP transaction anomaly. This prevents the Zero Trust decision engine from relying only on login status or role assignment.

Access-scenario analysis provides a more operational view of healthcare ERP risk. Figure 2 shows that user behavior risk is lowest during normal clinical access but increases during remote login, privileged maintenance, emergency access, and anomalous bulk access. Session trust probability decreases when access occurs from unmanaged devices, unfamiliar locations, or unusual time windows. Network behavior threat score becomes more visible in scenarios involving remote access or abnormal bulk operations because unusual service paths and traffic patterns are more likely to appear. This result shows that Zero Trust enforcement must adapt to scenario context rather than apply the same control intensity to every ERP session.

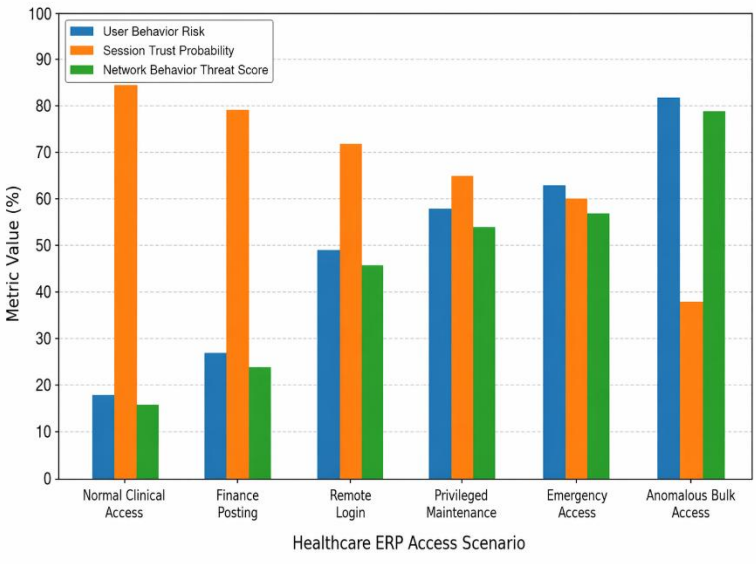

Figure 2. User Behavior Risk, Session Trust Probability, and Network Behavior Threat Score Across Healthcare ERP Access Scenarios

The comparison with static RBAC highlights why ML-driven behavioral analytics is needed. Static RBAC can define who is allowed to access each ERP function, but it cannot detect when a permitted user behaves unusually. Conditional access can challenge risky logins, but it may not understand transaction-level healthcare ERP behavior. Rule-based ERP monitoring can detect known violations, but it may miss abnormal sequences that do not match predefined rules. The proposed architecture improves on these approaches by combining Azure identity signals, access context, ERP transaction behavior, and network telemetry. This creates a more complete trust decision that reflects both technical access state and business-process behavior.

The practical implication is that healthcare ERP security on Azure should be implemented as continuous verification rather than periodic access review. Security teams can use ML-driven behavioral scores to trigger adaptive MFA, restrict sensitive ERP functions, investigate privileged role misuse, or escalate abnormal transaction activity. ERP administrators can retain role-based workflows while still detecting deviations that static access models miss. The architecture also supports auditability because each access decision is linked to identity evidence, privilege behavior, session trust, transaction pattern, and network context. This makes Zero Trust more explainable and useful for healthcare governance.

4. Conclusion

Healthcare ERP systems on Microsoft Azure require Zero Trust controls that understand both technical identity signals and ERP-specific business behavior. This article proposed a Zero Trust architecture enhanced by ML-driven behavioral analytics for identity verification, least-privilege control, session trust evaluation, transaction anomaly detection, network behavior scoring, and continuous verification.

The framework moves beyond static RBAC by evaluating whether a valid user behaves consistently with expected role, device, session, and transaction patterns. Its main contribution is the integration of Azure-native identity and security telemetry with healthcare ERP behavioral analytics.

The results show that identity trust accuracy, privilege deviation detection, and ERP transaction anomaly recognition improve across repeated verification cycles. Healthcare ERP access scenarios also show that user behavior risk, session trust probability, and network behavior threat scores vary according to context. This means that normal clinical access, finance posting, remote login, privileged maintenance, emergency access, and anomalous bulk access require different trust decisions. The framework therefore supports adaptive enforcement rather than uniform blocking or unconditional access.

Future work should validate the framework using production ERP logs, Azure Sentinel incidents, conditional access records, and real healthcare role structures. Additional research should examine explainable behavioral scoring, privacy-preserving access analytics, false challenge reduction, and integration with ERP-specific audit controls. The model can also be extended to multi-cloud healthcare ERP environments where identity, network, and application logs are distributed across several platforms. A practical Zero Trust architecture for healthcare ERP must remain measurable, explainable, and sensitive to both security risk and clinical-operational continuity.

References

1. Weinberg, A. I., & Cohen, K. (2024). Zero trust implementation in the emerging technologies era: Survey. *arXiv preprint arXiv:2401.09575*.
2. He, Y., Huang, D., Chen, L., Ni, Y., & Ma, X. (2022). A survey on zero trust architecture: Challenges and future trends. *Wireless Communications and Mobile Computing*, *2022*(1), 6476274.
3. Ramezanpour, K., & Jagannath, J. (2022). Intelligent zero trust architecture for 5G/6G networks: Principles, challenges, and the role of machine learning in the context of O-RAN. *Computer Networks*, *217*, 109358.
4. Ahmadi, S. (2025). Autonomous identity-based threat segmentation for zero trust architecture. *Cyber Security and Applications*, *3*, 100106.
5. Hassan, A., Rauf, A., Shafqat, N., Latif, R., & Khan, H. (2025). ZenGuard a machine learning based zero trust framework for context aware threat mitigation using SIEM SOAR and UEBA. *Scientific Reports*, *15*(1), 35871.
6. Aboukadri, S., Ouaddah, A., & Mezrioui, A. (2024). Machine learning in identity and access management systems: Survey and deep dive. *Computers & Security*, *139*, 103729.

7. Yadav, H., Subramanian, G., & Sharma, K. (2025). AI-Powered Zero Trust Access Evaluation Using Behavioral Fingerprinting. *International Journal of Computer Applications*, *187*(16), 19-22.
8. Fojude, M. (2025). Insider Threat Agent: A Behavioral Based Zero Trust Access Control Using Machine Learning Agent.
9. Yoheswari, S. (2025). Adoption of Identity and Access Management in Educational ERP Systems for Role-Based Security and Data Access Governance through Centralized Authentication Frameworks.
10. Mangayarkarasi, V. A., Vinayakan, K., & Kumar, A. D. (2024). Secure Cloud Data Storage with a Zero Trust Security Foundational Deep Learning Algorithm. *International Journal of Advanced Trends in Engineering and Technology*, *9*(2), 87-93.