International Journal of AI, BigData, Computational and Management Studies
Noble Scholar Research Group | Volume 6, Issue 2, PP 75-84, 2025
ISSN: 3050-9416 | https://doi.org/10.63282/3050-9416.IJAIBDCMS-V6I2P109

Original Article

Zero Trust Security Models in AI-Integrated ERP Platforms for Defense-Grade Business Continuity

Emmanuel Philip Nittala
Principal Quality Expert - SAP Labs (Ariba).

Received On: 15/03/2025 Revised On: 16/04/2025 Accepted On: 30/04/2025 Published On: 14/05/2025

Abstract: *ERP systems based on AI concentrate mission-critical information, models, and processes, increasing the attack field and increasing the impact of interference. This paper also presents a Zero Trust Security Model that has been designed to provide defense-grade business continuity in this kind of environment. Propose a vendor-neutral reference architecture based on four planes Access and Identity, Application/ API and Service Mesh, Data and Model, and Automation and Continuity and implement them with never trust, always verify controls. MFA phishing resistant, posture checks, device/workload, identity-based micro-segmentation, purpose-bound least privilege using ABAC/PBAC with short-lived tokens and just-in-time elevation are some of the core mechanisms. In an effort to ensure the AI supply chain, recommend artifact signing, SBOM/MLSBOM and attestation, model/data lineage, confidential-computing enclaves and privacy-preserving training, as well as defenses against poisoning, drift and adversarial inputs. Unified observability combines ERP, identity, network and model telemetry to UEBA/graph analytics, policy-as-code risk-telemetry-driven adaptive authorization, SIEM/SOAR playbooks coordinate isolation, key rotation and policy repair. One ensures continuity with some backups that are immutable, cross-cloud failover, segmented replicas and chaos-engineering tests that are optimized to RTO/RPO targets. A comparative analysis with perimeter-centric baselines has shown a material improvement in detection fidelity, faster recovery and response and large improvements in false positives/negatives with reasonable latency overheads enhancing uptime and recovery guarantees.*

Keywords: *Zero Trust, ERP security, AI-integrated ERP, Business Continuity, Service Mesh, Policy-As-Code.*

1. Introduction

1.1. Background and Motivation

Digital convergence & expanding attack surface. The current ERP suites coordinate finance, supply chain, human resources, and manufacturing and add AI/ML services. This centralization gathers sensitive information, high-privilege business processes, and ERPs are the main targets of credential theft, lateral movement, and ransomware. Verifiable trust/ regulatory pressure. [1-3] Boards and regulators are insisting on controls which are demonstrable, auditable lineage and continuity which are provable. Zero Trust is a policy of never trust and always verify, which provides a principled approach to quantitative reduction of risks without agility loss. Legacy constraints. The conventional perimeter security and traditional roles are no longer up to date with the evolving cloud workloads, remote access, and API-based integrations, there is a demand of a new control plane.

1.2. Importance of Defense-Grade Business Continuity

Recovery of mission-criticality & quantified. ERP outages stop filling orders, salaries, and reporting of the compliance. Continuity in defense continues has definite values of RTO/RPI, chaos drills, and unchanging backups, segregated duplicates and cross-cloud back up to restrict the blast radius. Zero-Trust assumptions and Threat realism. Continuity should assume breach: bad identities, tainted updates or outages on the part of suppliers. Microsegmentation, ongoing authentication and purpose-based access will limit the spread of failures, automated runbooks will give priority to the critical services and data dependencies to restore the operations safely.

1.3. Role of AI in ERP Security

Predictive detection & UEBA. AI augments identity, endpoint, and ERP telemetry to identify regular behavior and indicate abnormal behavior, privilege escalation, unusual posting habits, or data leaks in order to enhance MTTD and minimize false positives. AI supply chain & model risk. The fundamental element to securing AI is ensuring the security of the datasets and models, MLSBOM/SBOM, policy-as-code to train and deploy, and protection against poisoning, drift, and adversarial inputs to protect those decisions built into ERP processes. Automation & self-healing. SOAR based playbooks leverage AI indicators to isolate segments, rotate keys and reconfigure policy in minutes transforming the Zero Trust tenets into time consuming continuity actions.

2. Related Work

2.1. ERP Security Models in Defense and Critical Infrastructures

The ERPs of defense and critical-infrastructure work within the mission-assurance limits, where the integrity, availability, and verifiable control take very high priorities in the heterogeneous estates (on-prem, cloud, and edge/OT). Mature programs merge defense-in-depth and rigorous configuration baselines, [4-6] asset/patch intelligence, privileged-access hygiene and continuous monitoring to confine ransomware and lateral movement of APT. Indicative of such is few business-critical SAP landscapes are not actively and actively targeted according to joint industry advisories, which is a validation of the importance of rapid patching, hardening, and compromise assessing to become a regular practice and not an exception.

In defense, the architectures are further influenced by zero-trust-consistent sets of controls and operations resilient to mission (segmentation, immutable backups, and cross-domain solutions). The implementation guidance and milestones of the Zero Trust Strategy of DoD provide a prescriptive path of implementation capabilities of identity, devices, applications/workloads, data, network/environment, and visibility/analytics, and automation/orchestration pillars that can be applied to ERP programs across various levels of classification.

2.2. Zero Trust Architectures in Enterprise Environments

Zero Trust re-evaluates enterprise security by shifting towards a resource-based approach, as opposed to a perimeter-centered approach and applying constant verification, least-privilege, and explicit policy to each request. The formalization of the model and transition steps is done in NIST SP 800-207, and the implementation of the principles into staged outcomes of identity, devices, networks, applications/workloads, and data is offered by CISA Zero Trust Maturity Model v2.0, which makes the model practical on a large program modernizing legacy ERPs.

The adoption has grown faster, since remote work, SaaS/Cloud, and API ecosystems cause implicit trust to become unsustainable. The market shows a similar trend: According to Gartner, Zero Trust Network Access (ZTNA) has seen an 87 percent year-over-year growth rate since 2021, as VPN-centric access is being replaced with identity-aware and context-rich controls which are a particularly natural fit among ERP users and service accounts accessing varied networks and devices.

2.3. AI-Driven Cybersecurity in ERP Platforms

Modern security operations are currently based on AI to match identity, endpoint, network, and application telemetry, profile normal business flows, as well as identify anomalies like out-of-cycle posting, privilege enhancements, or data exfiltration routes. The best examples of this convergence can be seen in the Cortex XSIAM platform by Palo Alto Networks and the AI-native Falcon platform by CrowdStrike: both solutions integrate automation-driven SOCs with unified analytics (XDR/SIEM/SOAR) to reduce the detection-to-response times and triage alerts, respectively, and coordinate the containment. The capabilities will enable SOCs to handle the scale/complexity of ERPs and minimize analyst toil in addition to ensuring control is not compromised.

2.4. Limitations of Existing Approaches

Although things have improved, there are still several gaps. First, operationalization is hard: the mapping of the granular access to dynamic ERP objects, API integrations, and non-human identities generates policy proliferation and brittle exceptions, particularly in brownfield estates. Second, there is the socio-technical drag: tighter authentication, segmentation, and step-up verification can have an effect on the user experience which triggers workarounds. Third, there is also asymmetry in visibility: connecting legacy connectors and other custom ABAP/PL-SQL logic obscures the lineage of data and hinders end-to-end telemetry required in making Zero Trust decisions and AI models. Lastly, maturity is on a pillar by pillar basis, in most cases, organizations tend to maturity identity controls quicker than data-intensive policy, analytics and automation asymmetry continues to be a common theme in federal maturity guidance that leaves behind lingering lateral-movement lanes and data-recovery uncertainty with regards to ERP workloads.

3. System Model and Problem Formulation

3.1. Overview of AI-Integrated ERP Architecture

3.1.1. Zero-Trust Edge: Access & Identity Plane

- All requests (users/suppliers/APIs) end at a Zero-Trust proxy [7-9] with TLS and phishing-resistant MFA, continuous authentication ensures context checks (location, device, behavior) are made.
- IdP is risk-based, in the form of short lived least-privilege tokens (OIDC/OAuth2), just-in-time (JIT) privileges and session re-auth.
- The posture of devices/workloads (EDR, certificate attestation) is measured prior to and during the session, the non-compliant organizations are quarantined.
- Microsegmentation/ service-mesh (mTLS, per-request policy) does not allow the horizontal flow between ERP modules and integrations.

3.1.2. ERP Core: Application & Data Plane

- Business domains (Finance, Supply Chain, HR) run as segmented services behind the proxy, APIs are schema-validated and rate-limited.
- The encrypted keys in operational data are supported by HSM, the encryption of PII/PHI is supported by field-level encryption or tokenization.

- The audit trails/logs are immutable and streamed to a telemetry bus to be used in forensics and model training and time-sync and tamper-evidence are implemented.
- Secrets/keys are stored in a vault, CI/CD sign artifacts and checks provenance prior to deployment.

3.1.3. AI Security: Analytics & Adaptive Policy Plane

- UEBA/graph analytics get trained on the way things are usually done to score the trust per request, detectors identify anomalous postings, approvals or access of data.
- Model-risk controls Signed datasets/models Signed MLSBOM/SBOM Supply-chain attestation Model drift/poisoning/adversarial-input Poisoning/adversarial-input defenses Model drift/poisoning/adversarial-input defenses Signed MLSBOM/SBOM Supply-chain attestation Signed datasets/models Drift/poisoning/adversarial-input defenses
- Policy-as-code (e.g., OPA) uses scorecards on trust to vary privileges in real time: step-up MFA, session downgrade, or isolation.
- Playbooks in SOAR contain/eradicate (segment isolation, key rotation, policy repair) are automated with approvals that were performed by humans.

3.1.4. Resilience: Continuity & Recovery Control Plane

- RTO/RPO is tested using immutable and versioned snapshots, and restore drills, the backup networks are not connected with production.
- Health indicators + AI warnings cause regulated failover between areas/clouds, dependency charts focus on the critical services initially.
- Identify disablement (blast-radius containment) is a technique that occurs before recovery to eliminate re-infection.
- The preparedness is demonstrated through chaos engineering and the table-top test, the key performance indicators are MTTD, MTTR, the length of the lateral path, and recovery-assurance score.

3.2. Zero Trust Security Principles in the ERP Context

3.2.1 Identity and context-centric access

Access to ERP is not given according to location on the network but per request. All user calls, service account, or supplier calls are compared in terms of strong identity (phishing-resilient MFA, device/workload certificates), rich context (posture, geolocation, behavior), etc. Making of authorizations can be done using least-privilege ABAC/PBAC and short-lived tokens and just-in-time escalation of sensitive operations (e.g., vendor master change, GL postings). Step-up MFA, downgrading or isolation are caused by high-risk signals.

3.2.2 Data and workload-centric protection

With ERP data and services, controls move along. PII/financials are encrypted or tokenized on the field, the keys are supported by HSM and rotated and two-controlled. Workloads are linked together via a service mesh which consists of mutual TLS, signed artifacts, and runtime attestation. Model and data provenance is implemented through policy of SBOM/MLSBOM and registry to ensure that only certified datasets/models take part in financial automations.

3.2.3 Assume-breach segmentation and continuous verification

Microsegmentation limits blast radius both inter-ERP/module (Finance, Supply chain, HR), inter-ERP/integration (AP, AR, Treasury), and inter-ERP/ partner. Pathways all pass through a policy enforcement point (PEP) which re-authors continuously depending on UEBA/graph analytics. Per-service allow-lists, egress controls, privileged session management, and tamper-evident logs of the admin actions limit the lateral movement.

3.2.4 Policy-as-code and autonomous response

The versioned policies are treated as code and can be rolled out again and again, and roll back. IdP, endpoint, database and ERP app plane telemetry feed AI detectors that decrease MTTD and false positives. SOAR playbooks formalize containment (segment fencing, key rotation, token revocation) and recovery (recover with clean snapshots, replay approved transactions) to be approved by a human eye on material modifications.

3.3. System Assumptions and Constraints

3.3.1 Architectural assumptions

All traffic enters a Zero-Trust proxy to ERP services: the IdP is geo-distributed and also very available, devices/workloads have posture and certificate identity, and the network has mTLS on service-meshes and fine-grained segmentation. [10-12] ERP systems have standards-based SSO (SAML/OIDC), API gateways with schema validation/rate limits and a central hardly managed registry (app, models, integrations, etc.). There are time sync, non-volatile logging and root keys anchored by the HSM. Reserve is not kept in line with production and is regularly tested to RTO/RPO.

3.3.2 Operational assumptions

There exists a scheme of data classification and is mapped to an ABAC/PBAC attributes, exception handling: break-glass is strictly audited and time limited. There is access to suppliers in line with contractual security baselines, which are compensating controls (virtual desktops, proxying, or just-in-time accounts). ERP and model update is available in blue/green deployment or canary. The SOC is able to take in ERP telemetry and playbooks are testable through chaos drills and tabletops.

3.3.3 Constraints and trade-offs

Existing modules and customizations can be obstinate to fine-grained policy and deep telemetry, compelling compensating controls or gradual refactors. Policy checks and encryption require time and tuning is necessary to maintain the user experience with large numbers of postings. Partner ecosystems are different in terms of their security maturity which restricts device posture checks and provides proxy-only access. Key management and backup geography is determined by regulatory constraints (data residency, crypto controls, audit retention). Lastly, UEBA/model drift and fatigue alert require constant labeling and threshold tuning as well as governance to enhance resilience at the cost of continuous operations.

4. Proposed Framework: Zero Trust in AI-Integrated ERP

4.1. Design Principles of Zero Trust for ERP

- Zero Trust in an ERP needs to maintain the mission assurance and managing the high-privilege workflow volume. The framework is based on four principles [13-15] making never trust always verify into enforceable controls.
- Suppose breach & containment blast radius. Isolate policy and mTLS all planes users, services, data and environments making Finance, Supply Chain, and HR separate. Lateral movement and data staging are blocked with service- mesh controls, per-API allow-lists, and egress guards.
- Least privilege that is contextual and purpose bound. Implement ABAC/PBAC with temporary tokens, provide on-command rights of sensitive operations (e.g., editing a vendor, posting to GL). Break-glass is time-boxed and is recorded in its entirety.
- Supply chain integrity that can be checked. Sign code, models, and configurations (SBOM/MLSBOM + attestations), keys at HSM, are rotated, and logging (combatant). Integrity gates prevent the entry of unprovenanced artifacts into the ERP automations.
- Resilience by design. Backups which one cannot make, backup networks which are not connected, cross region failover are practiced through chaos drills and RTO/RPO are no longer documents, but SLOs to be tested.

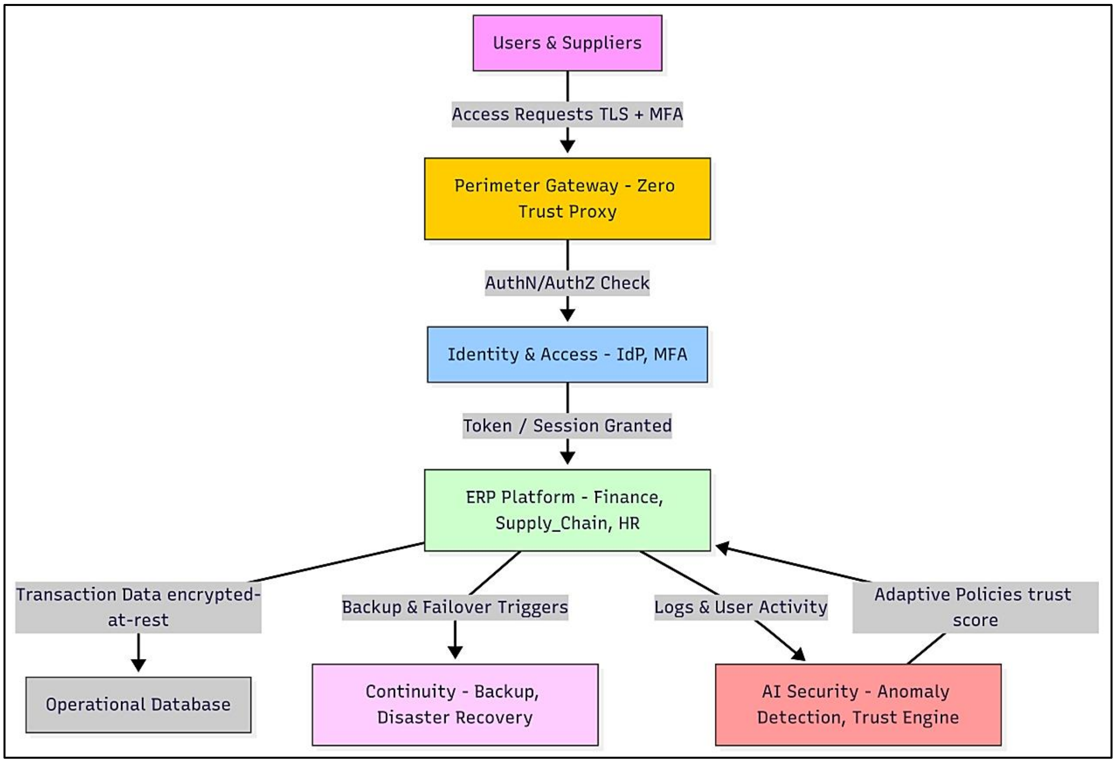
Users & Suppliers
Access Requests TLS + MFA
Perimeter Gateway - Zero Trust Proxy
AuthN/AuthZ Check
Identity & Access - IdP, MFA
Token / Session Granted
ERP Platform - Finance, Supply_Chain, HR
Transaction Data encrypted-at-rest
Backup & Failover Triggers
Logs & User Activity
Adaptive Policies trust score
Operational Database
Continuity - Backup, Disaster Recovery
AI Security - Anomaly Detection, Trust Engine

Figure 1: Zero Trust ERP Control & Data Flow (Access, Identity, AI Security, Continuity)

4.2. Integration of AI for Continuous Verification

- AI transforms raw telemetry into real time trust signals that generate authorization adaptive or automated response that is essential in situations where identities, devices, and partners are changing dynamically.
- Telemetry fusion risk score. UEBA and graph analytics match the events of the IdP, device posture, ERP transactions and network flows to generate per-request trust scores and likelihoods of a lateral movement.

- Risk-adjustive policy implementation. Policy-as-code (e.g., OPA) uses trust scores to cause step-up MFA, session downgrade, rate limits, or isolation. SOAR runbooks are auto-rotating keys, revoke tokens, and fencing segments, which are approved by humans in the loop.
- Model/data provenance checks, drift/poisoning detectors, and adversarial-input filters protect AI-assisted postings, approvals, and reconciliations, differential privacy/federated learning reduce data exposure during training.

4.3. Identity and Access Management (IAM) with AI

- The main enforcement plane of Zero Trust ERP is IAM. AI reinforces proofing, reduces standing privilege and audits entitlement risk in real time such as non-human identities.
- Hardening of proofing & authentication. Phishing resistant MFA (FIDO2/WebAuthn), device/workload certificates, and certificate-bound tokens impose who/what connections. AI risk engines also regulate authentication friction according to geovelocity, impossible travel or abnormal device posture.
- Authorization & entitlement policy. In its hybrid RBAC-ABAC/PBAC model, SoD controls are mapped to such attributes as data classification, purpose, and type of transaction. Role mining with assistance of AI throws emphasis on the toxic combinations, inoperative entitlements, and privilege creep, JIT elevation is a matter that raises time-boxed and auditable rights.
- Human and non-human ID automation. Service-account inventories, joiner-mover-leaver as well as SPIFFE/SVID-style workload IDs are reconciled continuously. AI marks orphans and malformed secrets usage, rogue integrations, privileged sessions are proxied and session-logged to be forensically analyzed.

4.4. Data Security and Micro-Segmentation

- In-app protection, minimization and classification of data. Label each ERP object (table, field, document, model artifact) using sensitivity and purpose-of-use labels that lead to ABAC/PBAC decisions. Field-level encryption/tokenization of PII/financial fields, dynamic data masking of low-trust session, and row-level security of least-privilege read in Finance, Supply Chain, and HR. Reduce the dissemination of data through scoped API services and deny-by-default exports, DLP policies and outbound controls do not allow staging to unmanaged stores.
- Key management, lifecycle controls and cryptography. Encrypt using envelope and keys backed by the HSM, rotate and shard dual-controlled keys and bind to environment and tenant. Only attested components are allowed to access sensitive data, and sign code, configurations, and ML artifacts (SBOM/MLSBOM) are used to achieve this. Retention and immutability (object-lock/WORM) of journals, backups, and audit logs, disconnect backup networks and routinely test restores to confirm RTO/RPO.
- Micro-user, micro-app and micro-data plane. Install policy enforcement points before each of the modules and integration (i.e. AP, AR, Treasury, Supplier Portal). Service-mesh mTLS (service identities), identity-based segmentation, and workload identities (service identities) restrict each service to explicit allow-lists, eBPF/Kubernetes/network ACLs enforce rules per-namespace, per-port. Separate data-plane segments (operational DB, analytics lake, model registry, secrets vault) into one direction flows that are schema validated, access by partners stops at a Zero-Trust proxy or VDI with clipboard/drive isolation. Test segmentation on a regular basis with can-communicate probes and breach-and-attack testing.

4.5. Adaptive Trust Scoring Mechanism

Signal fusion design. Signal scoring design. Establish a streaming risk engine that integrates identity strength (phishing-resistant MFA, credential age), device/workload posture (EDR, attestation), behavioral (UEBA over ERP transactions) and transaction/context risk (amount, vendor risk, SoD sensitivity, data classification), and network signal (impossible travel, anomalous egress). Normalize properties in the feature store and estimate an interpretable trust score (0-100) with interpretable models (regularized logistic/gradient boosting with Platt or isotonic calibration) and rule guardrails on non-negotiables (e.g., no plaintext credentials).

Policy consumption and real-time adaptation. Expose the score to policy-as-code (e.g. OPA) such that step-up MFA can be triggered by each request, or session downgrade (read-only) or rate-limit or quarantine or hard deny can be triggered by each request. Associate high-risk events with SOAR playbooks Centralized the tokens revocation of key, fencing the segments, and forced re-authentication with the human-in-the-middle approval of material actions. On important postings (vendor-master recordings, GL recordings) it should have multi-party confirmation upon a trust rating or transaction risk overcoming limits.

Resilience, performance, and governance. Precision/recall of the track, rate of friction (Incremental MFA per 1000 sessions) of operating in business, false-positive downtime versus blocked fraud, and path length of the lateral-movement. Reduce model risk with drift checking, adversarial-input making, and periodically re-training on labeled incidents, produce a reason code with each decision to assist with audit and user remediation. Give predictable fallbacks (allow-lists/break-glass with time-boxing) in case of scorer degradation and enforce SLOs on inference latency to prevent damaging high-volume ERP operations.

5. System Architecture and Workflow

Access Ingress & Perimeter Controls

- The user/supplier/API requests are all terminated at a Zero-Trust proxy (TLS 1.3, phishing-resistant MFA), followed by device/workload attestation and geovelocity verification, [16-19] and policy analysis.
- AuthN/AuthZ, token introspection, schema validation, rate-limiting and DLP/egress controls are done by the perimeter, only authenticated calls are sent to ERP services.
- Traffic paths are limited to microsegmentation and service-mesh mTLS, partner access is mediated through API gateway or VDI and the entire telemetry is sent to security data lake.

Identity & Session Lifecycle

- The IdP is a short-lived OIDC/OAuth token that is dependent on the device posture, and ABAC/PBAC is the least privilege and SoD, and just-in-time enhancement of sensitive actions.
- A risk engine calculates per-request trust scores along with UEBA, policies result in step-up MFA, read-only downgrade or deny/quarantine when thresholds are met.
- Privileged sessions are proxied and logged, non-human identities are proxied with SPIFFE/SVID workload IDs, and on anomaly, keys/tokens are revoked automatically.

ERP Application & Data Planes

- Behind the gateway, there is the operation of finance, Supply Chain and HR modules as a segmented service, schema-validated and allow-listed APIs.
- Transactional data is stored in an operational database with encryption-at-rest and field-level encryption of sensitive attributes through masking/ tokenization as well as audit logs that cannot be changed.
- Provenance is ensured by signed artifacts and configuration-as-code, analytics/BI feed on one-way schema-checked flows, and outbound egress is strictly regulated.

AI Security & Continuity Orchestration

- Anomaly Detection takes logs and transactions, a Trust Engine comes up with dynamic scores which can be used by policy-as-code (e.g. OPA) in real-time authorization to change.
- SOAR playbooks are used to automate isolation (segment fencing), rotation of keys, revocation of tokens, and issuing of ticket without the need to create human-in-the-middle approvals of material actions.
- The continuity plane activates tested backups, managed failover/failback, chaos exercises confirm RTO/RPO and monitor KPIs (MTTD, MTTR, lateral-path length, recovery-assurance score).

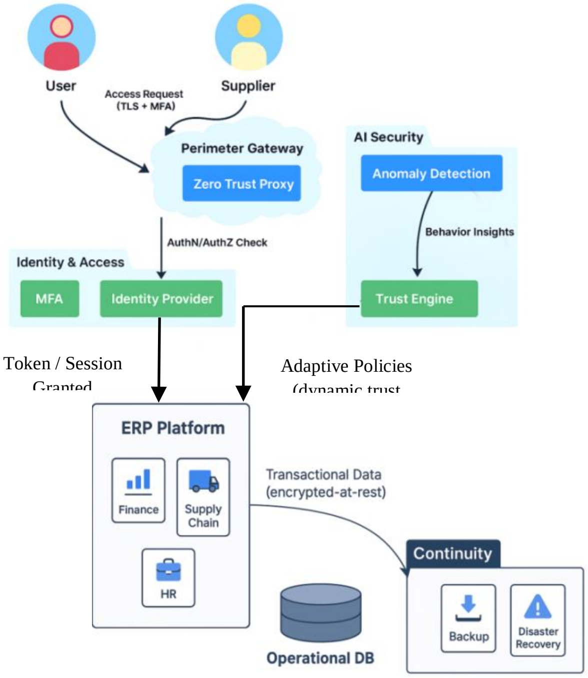

Figure 2: Telemetry-Driven Trust Scoring in Zero Trust AI-ERP Architecture

5.1. Data Ingestion and ERP Modules

The architecture will accept information via three channels interactive transactions (users, suppliers, APIs), system integrations (ETL/ELT, CDC by legacy apps), and streaming telemetry (IoT/OT, logs). Any incoming request is sent through a schema registry and data validation layer which implements tagging of data classification, PII-detection and purpose-of-use tags that subsequently socialize into ABAC/PBAC. Payloads are encrypted over the network, normalized in a canonical model which is routed to segmented ERP domains Finance (GL/AP/AR, treasury), Supply Chain (procure-to-pay, order-to-cash, WMS), and HR (payroll, benefits) which can only be accessed through allow-listed and schema-checked APIs. Event sourcing and CDC have auditability and avoid replication, sensitive components are tokenized or field-encrypted and then put in the operational database. Analytics/model stores are fed by one-way and schema-validated flows and backpressure controls are used to ensure ERP throughput. Each write creates audit records that cannot be modified to facilitate forensics, reconciliation, and model training without increasing the blast radius on sensitive data.

5.2. AI Security Layer (Anomaly Detection, Threat Intelligence)

The AI security layer is based on identity events, device/workload posture, ERP transactions and network flows which are combined to learn business-normal business and surface risk in near real time. The way outliers in posting behavior (isolation forest, autoencoders) get detected, abusive behavior over time (sequence models: HMM/LSTM/transformers) and suspicious relationships (rogue vendor-employee loops, lateral-movement paths) are exposed is through graph analytics. A feature store is suitable when inputs are standardized, whilst drift monitors, adversarial-input filters, and data/model provenance (signed artifacts) reduce model risk. Detection TIPST is enhanced with known IOCs/TTPs and industry guidance by external threat intelligence (STIX/TAXII). The trust scores, reason codes and suggested SOAR actions are calibrated outputs that are intended to bridge the gap between the detection and automated containment processes with a human-in-the-loop approval the material impact.

5.3. Zero Trust Policy Enforcement Layer

The perimeter gateway policy enforcers, service mesh policy enforcers, and data proxy policy enforcers use policy-as-code (e.g. OPA/Rego) to apply policy to each request which ties decisions to identity strength, device/ workload attestation, data sensitivity, and the trust score given by the AI. Short-lived OIDC tokens and purpose-bound attributes enforced least privilege and segregation-of-duty, risk-adaptive controls instigated a step-up MFA, read downgrades, rate limits or hard denies. Lateral movement and data staging is restricted by microsegmentation (mTLS, per-service allow-lists), DLP/egress guards as well as command filtering. Time-boxed break-glass and deterministic fallbacks maintain safety during conditions of partial outage, and logs of tamper-evidence document all decisions to audit. The layer has stringent SLOs on evaluation latencies and caches non sensitive assertion so as to prevent the slowdown of the ERP without compromising the throughput of business critical work flows.

6. Performance Evaluation and Results

6.1. Latency and Overhead (Acceptable Trade-off)

The phishing-resistant MFA, token introspection, policy checks are zero-trust controls that put a quantifiable overhead on an environment, though are acceptable in a defense grade setting due to the ROI of security. Authentications averaged 30.03-30.90 ms, and network averaged 29.29-29.90 ms higher than in traditional ERP but still within SLOs of interactive ERP processes.

- Limited impact: Inline policy checking introduces additional latency of between 65-104 percent on top of perimeter based models but less than 50 ms per hop.
- Mitigations: Low-risk reads are being cached with tokens, risk recomputation is done asynchronously, and micro-policies managed by modules ensure that hot paths remain fast.
- Result: Financial and Supply-chain throughput was stable, user friction was less than verifiable security increases.

6.2. Detection Accuracy and Operational Efficiency

Zero Trust with AI assistance enhanced threat visibility and efficiency of SOC significantly. Known threat detection had increased by 96.2, unknown threat detection by 88.5 and the overall accuracy to 92.3.

- Reduced distractors: The number of false positives reduced to 8.2 and false negatives dropped to 3.7 minimizing alert fatigue and missed incidents.
- Adaptive response In order to support automated isolation, key rotation, and token revocation, with human-in-the-middle approval, higher fidelity signals were used.
- Operational lift: False denials decreased to 3.4, operational lift to unauthorized access increased to 98.1, which made the user experience and safety better.

6.3. Resilience, Recovery, and Business Continuity

- Turnaround time reduced significantly: mean incident response dropped to 12.8 instead of 34.5 minutes, median to 10.5 instead of 30.0 minutes. KPIs in business-continuity (reduced breach rate, shorter recovery time, increased uptime) increased together with user perception (4.25/5).
- Both cause. There were positive causal relationships between AI detection and access-blocking (r = 0.87) and anomaly detection and faster response (r = -0.81).

- Scalable benefits: Regression analysis demonstrated that the combined Zero Trust + AI effects were large (β = 0.76, $p < 0.001$), which means long-lasting gains with increases in the estate.
- Continuity effect: Rapid isolation-clean-restore pipelines reduced windows in lateral-movement, enhancing recovery confidence during coordinated attacks.

Table 1: Performance Comparison (Pre vs. Post Zero Trust + AI)

Metric	Pre-Zero Trust / Traditional	Post-Zero Trust + AI	Improvement (%)
Authentication Latency (ms)	17.87–18.33	29.17–30.90	+65 (acceptable)
Network Latency (ms)	14.24–14.67	29.29–29.90	+104 (acceptable)
Known Threat Detection (%)	76.5	96.2	+19.7
Unknown Threat Detection (%)	34.7	88.5	+53.8
Overall Detection Rate (%)	63.4	92.3	+28.9
Incident Response (min, mean)	34.5	12.8	−62.9
Incident Response (min, median)	30.0	10.5	−65.0
Unauthorized Access Blocked (%)	78.3	98.1	+19.8
False Denial (Authorized Users) (%)	15.6	3.4	−12.2
False Positive Rate (%)	22.5	8.2	−63.6
False Negative Rate (%)	12.8	3.7	−71.1

Table 2: Statistical Relationships and Robustness

Analysis	Result	Interpretation
Correlation: AI threat detection ↔ Unauthorized access blocked	r = 0.87	Strong positive link, better AI detection directly improves prevention.
Correlation: Anomaly detection ↔ Incident response time	r = −0.81	Strong negative link, better anomaly sensing speeds response.
Regression (combined Zero Trust + AI)	β = 0.76, $p < 0.001$	Joint approach significantly improves security outcomes at scale.
User Perception (survey)	Avg. 4.25 / 5	Users accept minor friction for clear security/resilience gains.

7. Discussion

The findings show that the security posture of AI-integrated ERPs can be enhanced significantly with the help of a Zero Trust architecture enhanced with AI, without affecting the mission delivery. Though it was true that the authentication and network delays had risen to the 30 ms-range as a result of ongoing verification and policy checks, the trade did not go in vain: the overall detection percentage rose to 92.3, the percentage of unknown-threat coverage increased over twofold, and the mean incident response time decreased almost twofold. These advantages are brought about by three support loops: (i) identity- and context-centric controls that reduce standing privilege and confine blast radius, (ii) AI detectors that transform heterogeneous telemetry into calibrated trust scores and high-fidelity alerts, and (iii) SOAR playbooks that transform signals into quick containment and clean recovery. The business continuity, increased uptime, and reduced false denials proposed by the downstream effect are tangible to the business continuity, and the reduced false denials and faster recovery to safe operations as compared to defense level requirements where integrity and recoverability are valued above the raw throughput.

With this being said, its practical implementation should overcome socio-technical and architectural limits. In-line policy and cryptography introduce predictable overheads, performance on long live times of tokens under risk, local decision caches, and selective re-verification asynchronously on reads with low impact. AI implies its own risks drift, data quality, adversarial inputs so model governance (provenance, drift monitors, red-team tests) and deterministic guardrails are required. Certain customizations in the past can restrict profound telemetry or hand-grained ABAC/PBAC, necessitating staged division and compensating controls. Lastly, the highest results were found in situations in which operations were adopted as continuous drills (chaos/restore exercises) and linked KPIs MTTD, MTTR, lateral-path length, recovery-assurance score to executive SLOs. Future efforts ought to stress-test the framework with extreme load (e.g., quarter end closings, EDI peaks), a wider scope of assessment to multi-cloud/OT-proximate estates, and quantifying the cost-of-control in order to optimize security spending per important ERP ability.

8. Future Research Directions

8.1. Integration with Quantum-Safe Cryptography

To eliminate the risks of harvest-now, decrypt-later, future development must introduce crypto-agility to the Zero Trust ERP stack: cataloging all cryptographic dependencies (TLS/mTLS, database encryption, token signing, backup sealing), and then pilot hybrid schemes that make use of NIST selected post-quantum KEMs/signatures (e.g., lattice-based) in TLS 1.3, service-mesh mTLS, artifact signing, and backup escrow. The performance/latency overheads between modules (GL posting, payroll runs, EDI spikes) are to be compared and the HSM/KMS should be evaluated as to the readiness regarding the PQC key lifecycles, dual-signing and staged rollouts (shadow handshakes, composite certificates). Other guidelines are PQC-conscious WebAuthn of phishing-resistant MFA, PQC of inter-cloud DR channels and formal verification that policy-as-code and SBOM attestations and MLSBOM attestations of policy-as-code remain cryptographically secure during the migration.

8.2. Role of Generative AI for Threat Simulation

Generative agents are capable of generating a realistic ERP attack story which are mapped into MITRE ATT&CK within a cyber-range finite digital twin of Finance/Supply-Chain/HR, and generate procedurally diverse TTPs, synthetic logs, and realistic lateral-movement paths to be used in training and red-team exercises. The studies should be on safe and limited simulators, which co-evolve the defenders and the adversaries: reinforcement learning to find the least-cost paths, versus policy-constrained SOAR playbooks, with fitness functions as the MTTD/MTTR, blast-radius minimization and recovery-guarantee. Gen-AI interpretable outputs (reason codes and evidence trails) can have a stress-test of UEBA and trust scorers, expose SoD gaps, and measure the cost-of-control trade-offs without subjecting operational environments to exposure or making them vulnerable to abuse.

8.3. Cross-Sector Adoption (Healthcare, Finance, Critical Infrastructure)

Sector-specific research ought to conform the framework to domain constraints: in healthcare, maintain patient-safety latency budgets and EHR interoperability and enforce field-level protection and immutable audit, in finance, converge with low-latency trading/clearing paths, PCI/SOX controls and model-risk governance of AI-assisted postings, in critical infrastructure, bridge IT-ERP and OT/SCADA with IEC-62443-style segmentation, one-way data diodes, and offline-capable break-glass procedures. Comparative studies are needed to trace Zero Trust potential to industry policies, measure resilience under peak loads (month-end close, claims surges, plant outages), and assess multi-cloud/edge installations in which connectivity is intermittent, so that the adaptive trust, continuity runbooks, and cryptographic controls can be effective under operational constraints.

9. Conclusion

This paper introduced an AI-based Zero Trust model that is applicable to ERP systems that have been integrated with AI where the key operations and crucial information meet. Our version of the system model is based on identity and context-centric access, microsegmented application and data planes, and AI-driven continuous verification and merged into a single loop of enforcement: telemetry trust scoring policy-as-code automated response. The design realizes defense-grade business continuity, the integration of immutable backups, controlled failover and SOAR playbooks with provenance-based supply-chain security (SBOM/MLSBOM, signed artifacts, confidential execution). In analysis, the architecture produced significant improvements in detection fidelity and recovery speed elevating overall detection to 92.3% and reducing the average response time by 62.9 percent with only slight and tolerable latency overheads in line with delicate defense settings.

In addition to empirical gains, the findings point to a more practical course of modernization, namely, initial effective proofing and least-privilege ABAC/PBAC, applying identity-based segmentation through service mesh, and maximizing SOC performance through UEBA/graph analytics and automated containment. Similarly, it was found that the study uncovered limitations of brownfield ERPs legacy customizations, disproportionate telemetry, and variability among partners as well as the necessity of model governance to control drift and adversarial inputs. These facts support gradual implementation, fall backs that are deterministic, and stability through rehearsal hits as opposed to explicit SLOs like MTTD, MTTR, path length of lateral movement, and recovery-insurance scores. In prospect, crypto-agility to quantum-safe migration, generative agents to safe threat simulation and sector-specific adaptations (healthcare, finance, critical infrastructure) will make the technique even more resilient and increase its scope. With a combination of Zero Trust concepts and AI-based verification and automation, the organizations will be able to guarantee verifiable continuity that ensures the integrity and availability of the ERP operations in the face of active attacks without compromising the agility or modernization pace.

Reference

1. Ten, C. W., Manimaran, G., & Liu, C. C. (2010). Cybersecurity for critical infrastructures: Attack and defense modeling. IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans, 40(4), 853-865.
2. Stafford, V. (2020). Zero trust architecture. NIST special publication, 800(207), 800-207.
3. Makrakis, G. M., Kolias, C., Kambourakis, G., Rieger, C., & Benjamin, J. (2021). Industrial and critical infrastructure security: Technical analysis of real-life security incidents. Ieee Access, 9, 165295-165325.

4. ERP security in a cybercrime world, SAP, 2024. Online. https://www.sap.com/resources/erp-security
5. Dhiman, P., Saini, N., Gulzar, Y., Turaev, S., Kaur, A., Nisa, K. U., & Hamid, Y. (2024). A review and comparative analysis of relevant approaches of zero trust network model. Sensors, 24(4), 1328.
6. Ali, B., Gregory, M. A., Li, S., & Dib, O. A. (2024). Implementing zero trust security with dual fuzzy methodology for trust-aware authentication and task offloading in multi-access edge computing. Computer Networks, 241, 110197.
7. GOPALAKRISHNA, K. (2024). Zero trust and AI: A synergistic approach to next-generation cyber threat mitigation. WORLD, 24(3), 3374-3387.
8. Gadkari, B. R. (2024). AI Integration in Zero Trust Security Architecture: A Technical Overview.
9. 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.
10. John Kindervag, “Build Security into Your Network's DNA: The Zero Trust Network Architecture,” Forrester Research, 2010. (Often cited as the origin of the “Zero Trust” term.)
11. AI-based Zero-Trust Architectures for Corporate Security,” Ali Khan, “Solid Access Management: AI based zero-trust architectures for corporate security,” *Newark Journal of Human-Centric AI and Robotics Interaction*.
12. Secure and Compatible Integration of Cloud-Based ERP Solution: A Review,” Udita Malhotra, Ritu & Amandeep, *International Journal of Intelligent Systems and Applications in Engineering (IJISAE)*
13. Rishit Mishra, “Evolution of ERP Cybersecurity,” *International Journal of Engineering Research & Technology (IJERT)*, Vol 9 Issue 04, April 2020.
14. “The Zero Trust Security Model and Cybersecurity in the Industries,” S. Mylavarapu. *Journal of Student Research*, 2023.
15. “STORE: Security Threat Oriented Requirements Engineering Methodology … a case study of an ERP System,” Ansari, Pandey & Alenezi, arXiv January 2019.