



e-ISSN: 2278-8875

p-ISSN: 2320-3765

International Journal of Advanced Research

in Electrical, Electronics and Instrumentation Engineering

Volume 15, Issue 4, April 2026



Impact Factor: 8.807

☎ 9940 572 462

📞 6381 907 438

✉ ijareeie@gmail.com

@ www.ijareeie.com



Threat Intelligence-Driven Security Architecture for Cloud-Native ERP Ecosystems: Dynamic Risk Scoring, Behavioral Anomaly Detection, and Automated Compliance Enforcement across Hybrid Landscapes

Srinivas Kakarla

SAP Architect, USA

ABSTRACT: Enterprise Resource Planning (ERP) platforms have migrated decisively toward cloud-native and hybrid deployment models over the past decade, a shift that has expanded the attack surface of the systems that underpin financial reporting, supply chain operations, human capital management, and regulatory compliance. This article presents an integrated security architecture that fuses external threat intelligence with internal telemetry to produce continuously updated, asset-aware risk scores; applies behavioral anomaly detection to identify misuse of legitimate credentials and violations of segregation-of-duties (SoD) controls; and automates the enforcement of regulatory and internal compliance policy through policy-as-code mechanisms. Drawing on an 18-month multi-organization pilot spanning on-premises SAP S/4HANA, private-cloud virtualization platforms, and public-cloud SaaS ERP suites, we quantify reductions in mean detection latency (from a median of 205 minutes to 38 minutes), false-positive alert rates (from 38% to 10.1%), and mean time to remediate compliance exceptions (from 22.4 days to 3.7 days across five control families). We further show that a hybrid scoring function combining exploit-likelihood signals from structured threat intelligence feeds with behavioral deviation scores from unsupervised sequence models outperforms either signal in isolation, achieving an area under the receiver operating characteristic curve (AUC) of 0.96 versus 0.74 for a conventional rule-based baseline. The architecture, its component models, deployment considerations, and residual limitations are described in detail, together with a forward-looking discussion of the operational and governance implications for security and audit functions responsible for cloud-native ERP ecosystems.

KEYWORDS: ERP security, threat intelligence, dynamic risk scoring, behavioral anomaly detection, UEBA, segregation of duties, policy-as-code, compliance automation, hybrid cloud, SAP S/4HANA, zero trust, SOAR orchestration, cyber risk quantification

I. INTRODUCTION

Enterprise Resource Planning platforms have historically been treated as trusted, internally-facing systems, protected primarily through network segmentation and coarse-grained authorization roles. That assumption no longer holds. Between 2021 and 2025, the majority of Fortune 2000 organizations initiated or completed a migration of core ERP workloads - finance, procurement, human capital management, and supply chain - to public or hybrid cloud infrastructure, driven by vendor roadmaps such as SAP's 2027 end-of-mainstream-maintenance deadline for ECC and Oracle's continued investment in Fusion Cloud Applications. This migration has dissolved the network perimeter that previously served as the default control boundary and has introduced a wide range of new attack surfaces: internet-facing OData and REST interfaces, third-party middleware connectors, CI/CD pipelines that can push configuration changes directly into production landscapes, and identity federation trust relationships that span multiple cloud tenants. At the same time, the threat actor ecosystem targeting ERP systems has professionalized. Ransomware affiliates increasingly pursue financial and payroll modules directly, business email compromise crews chain ERP vendor-master-record manipulation with payment redirection fraud, and nation-state-aligned groups have been observed conducting long-dwell reconnaissance inside SAP and Oracle landscapes to harvest intellectual property and supply chain data. Conventional signature-based monitoring, built for network and endpoint telemetry, is poorly suited to detect this class of activity because the actions involved - a procurement clerk approving a vendor bank-account change, a developer promoting a custom ABAP transport, an integration account querying a materials-management table - are all technically permitted operations. The security problem is not unauthorized access in the classical sense; it



is the misuse of authorized access in a manner inconsistent with an entity's established behavioral baseline, risk exposure, and compliance obligations.

This article proposes and evaluates an integrated architecture that addresses three interlocking requirements of modern ERP security programs:

- Dynamic, context-aware risk scoring that continuously re-prioritizes assets, users, and transactions using external threat intelligence rather than static annual risk assessments;
- Behavioral anomaly detection capable of surfacing SoD violations, credential misuse, and configuration drift using statistical and machine-learning baselines rather than fixed rule thresholds; and
- Automated compliance enforcement that encodes regulatory and internal control requirements as executable policy, closing the gap between control documentation and control operation.

The remainder of this article is organized as follows. Section 2 reviews related work in ERP security, cyber-risk quantification, and UEBA. Section 3 characterizes the current threat landscape. Section 4 presents the reference architecture. Sections 5 through 7 describe the three core subsystems in detail. Section 8 discusses deployment across hybrid landscapes. Sections 9 and 10 present the evaluation methodology and results from an 18-month, multi-organization pilot. Sections 11 through 14 discuss findings, limitations, governance implications, and future work, and Section 15 concludes.

Scope note: this article treats “ERP” broadly to include core financial and operational suites (SAP S/4HANA and ECC, Oracle Fusion Cloud, Workday, Microsoft Dynamics 365) together with their adjacent procurement, human-capital-management, and supply-chain modules, whether deployed on-premises, in private cloud, or as multi-tenant SaaS.

II. BACKGROUND AND RELATED WORK

2.1 ERP-Specific Security Research

Early academic and practitioner work on ERP security concentrated on role and authorization design, particularly the detection of segregation-of-duties conflicts within SAP's role-based access control model. Tools built around this problem, including SAP GRC Access Control and third-party SoD rule sets, rely on predefined conflict matrices - for example, flagging a user who holds both vendor-master maintenance and payment-run execution authorizations. These approaches remain foundational but are static: they detect the existence of a conflicting authorization combination but not whether that combination is being actively exploited, and they require manual rule maintenance as business processes evolve.

2.2 Cyber-Risk Quantification

Parallel work in cyber-risk quantification, most notably the FAIR (Factor Analysis of Information Risk) model and NIST's Cybersecurity Framework 2.0, has advanced organization-level frameworks for translating threat and vulnerability data into financial or ordinal risk estimates. These frameworks are typically applied at a slow cadence (quarterly or annual risk registers) and at a coarse granularity (business unit or system level) rather than continuously and at asset- or transaction-level granularity. The architecture in this article adapts FAIR-style factor decomposition - threat event frequency, vulnerability, and loss magnitude - into a scoring function that recomputes on a fifteen-minute cycle against individual ERP assets and user sessions.

2.3 User and Entity Behavior Analytics

UEBA has matured substantially in the general SIEM/XDR market, with peer-group baselining, sequence modeling, and graph-based correlation now standard capabilities in commercial platforms. However, general-purpose UEBA products are trained on network and identity telemetry (VPN logins, Active Directory events, cloud API calls) and have limited visibility into ERP-native semantics - for instance, the business meaning of a change to a vendor's IBAN field versus a change to a cost-center description. Effective anomaly detection for ERP requires feature engineering informed by functional business-process knowledge, not only generic access-log statistics.

2.4 Policy-as-Code and Compliance Automation

The DevSecOps movement popularized policy-as-code frameworks such as Open Policy Agent (OPA)/Rego and HashiCorp Sentinel for infrastructure guardrails. Their application to ERP compliance - encoding SOX ITGC, GDPR Article 32, and PCI DSS 4.0 control requirements as machine-evaluable policy against ERP configuration and transaction data - remains comparatively immature, in part because ERP configuration objects (authorization roles, workflow steps, table maintenance settings) do not expose the same declarative interfaces as cloud infrastructure



resources. Section 7 describes an adapter pattern that closes this gap without requiring ERP vendors to natively support policy-as-code toolchains.

III. THREAT LANDSCAPE FOR CLOUD-NATIVE ERP ECOSYSTEMS

The pilot's threat-intelligence ingestion layer classified confirmed incidents and near-miss findings across six categories over an 18-month observation window. Table 1 summarizes the resulting taxonomy, which subsequently informed both the risk-scoring factor design (Section 5) and the anomaly-detection feature set (Section 6).

Table.1. ERP Threat Taxonomy and Representative Techniques

Category	Representative Technique	Primary ERP Target	Typical Entry Vector
Excessive / SoD Access	Conflicting authorization exploitation	Financial & procurement roles	Insider misuse, stale access
Credential Abuse	Password spraying, session token replay	Fiori/portal front-ends	Exposed endpoints, phishing
Insecure Interfaces	Unauthenticated RFC/OData calls	Integration middleware	Misconfigured gateway
Configuration Drift	Unauthorized table/parameter change	Basis & security config	Change-control bypass
Custom Code Flaws	ABAP/Apex injection, missing auth checks	Custom transactions/apps	Insecure custom development
Third-Party Connector Risk	Compromised iPaaS/EDI credentials	B2B & EDI gateways	Supply-chain compromise




Figure.8. Distribution of detected threat categories across the 18-month pilot (n = 4,812 confirmed findings).

Excessive and SoD-conflicting access accounted for the largest single share of confirmed findings (27%), consistent with prior audit literature indicating that authorization sprawl - the gradual accumulation of unused or overly broad entitlements - remains the dominant residual risk in mature ERP landscapes even after initial role design is complete. Credential abuse (19%) and insecure interfaces (16%) reflect the expanded external attack surface introduced by cloud



and hybrid deployment; both categories were materially more prevalent in the public-cloud SaaS segment of the pilot population than in the on-premises segment.

3.1 Regulatory Drivers

Six overlapping regulatory and control frameworks were in scope across the pilot's participating organizations: NIST CSF 2.0, ISO/IEC 27001:2022, Sarbanes-Oxley IT General Controls (SOX ITGC), GDPR Article 32, PCI DSS 4.0, and the Cloud Security Alliance Cloud Controls Matrix (CCM) v4. Each framework imposes distinct evidentiary and cadence requirements - SOX ITGC, for example, requires point-in-time attestation tied to a financial reporting period, whereas PCI DSS 4.0's 2024 customized-approach provisions permit continuous, telemetry-based control validation. Section 7 describes how the compliance-enforcement engine maps a single underlying control observation to multiple framework citations to avoid duplicative evidence-collection effort.

IV. REFERENCE ARCHITECTURE

Figure 1 presents the end-to-end reference architecture. Four external and internal data sources - structured threat intelligence feeds, MITRE ATT&CK and CISA Known Exploited Vulnerabilities (KEV) mappings, vulnerability and patch management signals, and native ERP telemetry - feed two parallel analytical engines: the Dynamic Risk Scoring Engine (Section 5) and the Behavioral Anomaly Detection subsystem (Section 6). Outputs from both engines are consumed by a Policy-as-Code Compliance Enforcement Engine and an Adaptive Response Orchestrator, which converge into a Unified Risk and Compliance Register that drives executive and auditor-facing dashboards.




Figure.1. Reference architecture for threat intelligence-driven ERP security, showing data flow from intelligence sources through scoring and detection engines to enforcement and reporting layers.

The architecture is deliberately layered so that each engine can be deployed incrementally: organizations with mature vulnerability-management programs but no behavioral analytics capability can adopt the risk scoring engine independently, while organizations with existing UEBA tooling can integrate the compliance-enforcement layer without re-platforming their detection stack. This modularity was a deliberate design response to pilot partner feedback that a single monolithic platform would face prohibitive replacement costs against incumbent SIEM and GRC investments.



4.1 Design Principles

- Signal fusion over signal replacement - the architecture augments existing SIEM, GRC, and vulnerability management tooling rather than requiring their removal.
- Asset- and identity-centric scoring - risk and anomaly scores attach to ERP business objects (vendor master records, payment runs, custom transactions) and identities, not only to network hosts.
- Explainability by construction - every automated action is traceable to the specific intelligence, behavioral, or policy signal that triggered it, a requirement driven by audit and regulatory review.
- Graduated automation - enforcement actions escalate from advisory notification to soft-block to hard-block based on confidence and business-criticality, avoiding blanket automated denial of legitimate transactions.

4.2 Component Inventory

Component	Primary Function	Representative Technology Pattern
Intelligence Normalizer	Ingests STIX/TAXII feeds, KEV, vendor advisories	Message queue + schema mapper
Risk Scoring Engine	Computes composite asset/identity risk	Streaming feature store + scoring service
UEBA Engine	Detects behavioral deviation and SoD violation	Sequence model + graph analytics
Policy Engine	Evaluates control policy against config/transactions	OPA/Rego-style policy evaluation
Response Orchestrator	Executes graduated remediation actions	SOAR playbook engine
Unified Register	Aggregates risk, findings, and control status	Risk/GRC data warehouse + BI layer

V. DYNAMIC RISK SCORING ENGINE

The risk scoring engine computes a composite score for each monitored asset (an ERP system, module, custom transaction, or integration endpoint) and identity (a named user or service account) on a fifteen-minute recomputation cycle. The scoring function is a weighted combination of four factors, adapted from FAIR-style risk decomposition:

- Asset criticality (Ac) - business-impact weighting derived from data classification, financial materiality, and process dependency mapping;
- Exploit likelihood (El) - derived from CVSS scores, CISA KEV listing status, and observed exploitation activity in threat-intelligence feeds relevant to the asset's technology stack;
- Exposure (Ex) - network and identity reachability, including internet-facing status, authentication strength, and number of entities with access; and
- Control mitigation (Cm) - a discount factor reflecting compensating controls already in force (patch status, network segmentation, MFA enforcement).

The composite score is computed as:

Risk Score = $[(Ac \times 0.35) + (El \times 0.30) + (Ex \times 0.25)] \times (1 - Cm \times 0.5)$, scaled to a 0–100 range

Weights were calibrated using historical incident data from pilot partners via logistic regression against confirmed-incident outcomes, then reviewed and adjusted by each partner's risk committee to reflect organization-specific priorities; the weights shown above represent the pooled cross-partner calibration and are provided as a reference starting point rather than a fixed constant.




Figure.2. Dynamic risk scoring data pipeline, from raw signal ingestion through temporal decay to action routing.

5.1 Temporal Decay

A critical departure from static annual risk assessments is the introduction of temporal decay: a score elevated by a newly disclosed vulnerability or a threat-intelligence indicator decays exponentially toward its baseline value over a configurable half-life (default 72 hours) absent reinforcing signal, such as confirmed patch deployment or continued attacker interest in the associated CVE. This prevents "alert fatigue accumulation," in which risk registers grow monotonically because stale findings are never algorithmically retired.

5.2 Illustrative Scoring Matrix

Table 2 and Figure 6 present the resulting composite risk matrix across four asset-criticality tiers and five threat-likelihood bands, holding exposure and control mitigation at population-median values. Mission-critical assets facing an "almost certain" likelihood band (e.g., an internet-facing payment run interface with an actively exploited CVE) reach the top decile of the scoring range, triggering automatic escalation to the Adaptive Response Orchestrator described in Section 7.




Figure.6. Dynamic risk score matrix across asset criticality and threat likelihood, at median exposure and control-mitigation values.

**Table.2. Risk Tier Thresholds and Recommended Response Posture**

Score Range	Risk Tier	Recommended Response Posture	Default SLA
85–100	Critical	Automated soft-block + incident bridge activation	15 minutes
65–84	High	Analyst triage with pre-staged remediation playbook	1 hour
40–64	Moderate	Queued review within standard SOC shift	8 hours
20–39	Low	Logged for trend analysis; no immediate action	72 hours
0–19	Minimal	Baseline monitoring only	N/A

VI. BEHAVIORAL ANOMALY DETECTION

Where the risk scoring engine answers "how exposed is this asset," the behavioral anomaly detection subsystem answers "is this specific action, by this specific identity, consistent with expected behavior." The subsystem operates on ERP-native event logs (SAP Security Audit Log, change documents, table access logs; Oracle Fusion audit trails; Workday activity logs) rather than generic network telemetry, allowing feature engineering that reflects business-process semantics.




Figure.3. Behavioral anomaly detection pipeline (UEBA), from log collection through SoD graph correlation to alert prioritization.

6.1 Feature Engineering

Features fall into four families: temporal (time-of-day and day-of-week deviation from an entity's 90-day rolling baseline), volumetric (transaction count and monetary-value deviation), relational (new or unusual counterparty entities, e.g., a first-time vendor payment), and structural (authorization-graph position relative to defined SoD conflict rules). Peer-group baselining compares each entity not only against its own history but against a cohort of entities with equivalent job function, reducing false positives arising from legitimate role changes.

6.2 Sequence Modeling

Two complementary unsupervised model families were evaluated during the pilot: an LSTM-based sequence autoencoder trained on chronologically ordered transaction sequences per entity, and a graph-based SoD correlation engine that continuously evaluates the authorization graph against a maintained conflict matrix. The autoencoder's reconstruction error serves as an anomaly score for sequence-level deviation (e.g., an unusual order of operations culminating in a payment), while the SoD engine independently flags structural conflicts regardless of observed behavior. A combined score - the maximum of normalized autoencoder error and SoD conflict severity - was found during calibration to outperform either signal alone, consistent with the classifier comparison in Section 10.




Figure.9. ROC curve comparing the combined behavioral anomaly classifier against a conventional rule-based baseline

6.3 Alert Prioritization

Raw anomaly scores are combined with the corresponding asset/identity risk score from Section 5 before prioritization, ensuring that a moderate behavioral deviation on a mission-critical payment process is ranked above a severe deviation on a low-criticality reporting query. Table 3 summarizes the resulting detection-latency improvement across five representative threat categories relative to the pilot partners' pre-existing SIEM correlation rules.

Table.3. Median Detection Latency, Legacy SIEM vs. Proposed Architecture

Threat Category	Legacy Baseline (min)	Proposed Architecture (min)	Improvement
Credential Misuse	182	34	81.3%
SoD Violation	240	41	82.9%
Data Exfiltration	205	38	81.5%
Privilege Escalation	168	29	82.7%
Configuration Drift	310	52	83.2%




Figure.4. Median detection latency by threat category, legacy SIEM baseline versus proposed architecture.

VII. AUTOMATED COMPLIANCE ENFORCEMENT

The compliance enforcement engine encodes control requirements from the six frameworks introduced in Section 3.1 as declarative policy evaluated continuously against ERP configuration state, authorization assignments, and transaction data. Because ERP platforms do not natively expose OPA/Rego-compatible configuration interfaces, the engine relies on an adapter layer that extracts configuration and transaction data into a normalized object model, against which policy is then evaluated using standard policy-as-code tooling.

7.1 Control-to-Policy Mapping

A single control observation frequently satisfies evidentiary requirements across multiple frameworks. For example, a policy verifying that all vendor bank-account changes require dual approval within a four-eyes workflow simultaneously supports SOX ITGC change-management controls, PCI DSS 4.0 requirement 7 (access restriction), and CCM v4 domain IAM-04. Table 4 provides a representative excerpt of this mapping; the full matrix used in the pilot contained 214 control-to-policy mappings and is summarized further in Appendix B.

Table.4. Representative Control-to-Policy Mapping (Excerpt)

Policy ID	Control Objective	NIST CSF 2.0	SOX ITGC	PCI DSS 4.0	CCM v4
POL-014	Dual approval for vendor bank changes	PR.AA-05	Yes	Req. 7	IAM-04
POL-027	SoD: no combined PO create + approve	PR.AA-05	Yes	-	IAM-09
POL-041	MFA enforced on privileged Fiori roles	PR.AA-03	Yes	Req. 8.4	IAM-02
POL-058	Transport approval before production import	PR.PS-01	Yes	Req. 6.5	CCC-01
POL-073	Encryption at rest for PII tables	PR.DS-01	-	Req. 3.5	DSP-06



7.2 Graduated Enforcement Actions

Consistent with the design principle of graduated automation (Section 4.1), the engine supports four escalating response tiers: advisory notification to the control owner; workflow-embedded soft-block requiring supervisory override with logged justification; hard-block preventing transaction posting until remediated; and automatic access revocation for the highest-severity, highest-confidence violations (e.g., a terminated employee's credentials still active in a production role).




Figure.10. Mean time to remediate (MTTR) across five control families, before and after automated compliance enforcement

Table.5. Mean Time to Remediate, Before vs. After Automation

Control Family	Before (days)	After (days)	Reduction
Access Recertification	21.0	3.2	84.8%
Policy Exception Closure	14.0	2.1	85.0%
SoD Conflict Remediation	26.0	4.8	81.5%
Critical Patch Deployment	18.0	5.5	69.4%
Audit Evidence Collection	33.0	2.8	91.5%

VIII. HYBRID LANDSCAPE DEPLOYMENT MODEL

Pilot partner landscapes spanned three or more of the following zones simultaneously: on-premises data centers running SAP ECC or early S/4HANA releases, private-cloud virtualization platforms hosting migrated or satellite workloads, public-cloud infrastructure-as-a-service environments hosting re-platformed S/4HANA instances, and multi-tenant SaaS ERP modules such as Workday or SAP Ariba. Each zone requires a distinct control-point pattern, illustrated in Figure 11.




Figure.11. Security control points across the hybrid ERP landscape, unified beneath a central control plane and a zero-trust identity fabric.

8.1 Zone-Specific Control Patterns

Zone	Primary Control Point	Telemetry Source	Latency Consideration
On-Prem Data Center	Secure gateway + log forwarder	Security Audit Log, Syslog	Batch upload, 1–5 min lag
Private Cloud	Service mesh sidecar policy	Mesh access logs, hypervisor events	Near real-time
Public Cloud IaaS	Cloud-native CSPM/CWPP	Cloud provider control-plane logs	Real-time streaming
SaaS ERP	API gateway + CASB	Vendor audit API, SCIM logs	Polling interval dependent

8.2 Identity Fabric

All four zones are unified beneath a federated, zero-trust identity fabric providing single sign-on, just-in-time privileged access provisioning, and consistent MFA policy enforcement independent of the underlying ERP deployment model. This layer proved to be the single highest-leverage investment reported by pilot partners: organizations that had already consolidated identity federation prior to pilot onboarding achieved full architecture deployment 40% faster, on average, than those that federated identity concurrently with the pilot.

Practitioner note: landscapes with unresolved identity federation debt (multiple unlinked directories, orphaned service accounts across zones) should treat identity consolidation as a prerequisite work stream, not a parallel one - every downstream engine in this architecture depends on a reliable, de-duplicated identity graph.



IX. EVALUATION METHODOLOGY

9.1 Pilot Population

The evaluation draws on an 18-month pilot conducted across eleven organizations spanning manufacturing, retail, financial services, and healthcare verticals, collectively operating 34 distinct ERP landscape instances. Table 6 summarizes the pilot population by deployment model and organizational size band.

Table.6. Pilot Population by Deployment Model and Organization Size

Deployment Model	Instances	Small (<2K users)	Mid users (2K–10K)	Large (>10K users)
On-Premises	9	2	4	3
Private Cloud	11	3	5	3
Public Cloud IaaS	10	1	5	4
SaaS ERP	4	2	2	0

9.2 Metrics

- Detection latency - elapsed time between the earliest reconstructable evidence of a threat event and its confirmed detection by the monitoring architecture;
- False-positive rate - proportion of generated alerts subsequently dispositioned as benign by analyst review, tracked monthly;
- Mean time to remediate (MTTR) - elapsed time between confirmed finding and verified control closure;
- Classifier performance - AUC, precision, and recall for the behavioral anomaly classifier against a labeled holdout set of 1,240 confirmed incidents and 8,600 confirmed-benign events;
- Compliance coverage - percentage of applicable control objectives per framework with continuous, automated evidence collection versus manual, point-in-time attestation.

9.3 Baseline Comparison

Each metric is compared against a documented pre-pilot baseline representing each partner's incumbent SIEM correlation rules, quarterly manual SoD reviews, and point-in-time compliance attestation process. Baseline data was reconstructed from twelve months of historical incident and audit records preceding pilot onboarding, cross-validated against each partner's internal audit findings for the same period.

X. RESULTS

10.1 Detection Latency

Across all five tracked threat categories, the proposed architecture reduced median detection latency by 81–83% relative to baseline (Table 3, Figure 4). The largest absolute improvement occurred in configuration-drift detection (310 minutes to 52 minutes), attributable to continuous policy evaluation replacing the partners' prior reliance on periodic configuration snapshots.

10.2 False-Positive Rate Trend

Figure 5 tracks the monthly false-positive rate over the full 18-month pilot window (first 12 months shown for trend clarity). The rate declined from 38% in month one - reflecting cold-start baselining before sufficient behavioral history had accumulated - to 10.1% by month twelve, as peer-group baselines matured and analyst feedback was incorporated into supervised threshold tuning for the highest-volume alert categories.




Figure.5. False-positive rate reduction over the 12-month pilot ramp period

10.3 Alert Volume Composition

Total monthly alert volume also declined over the observation window even as detection sensitivity improved, driven by the temporal-decay mechanism (Section 5.1) retiring stale low-confidence findings and by supervised tuning suppressing recurring benign patterns identified during month-over-month analyst review.




Figure.12. Monthly alert volume by severity tier across the 12-month pilot ramp period.

10.4 Classifier Performance

The combined behavioral anomaly classifier (sequence autoencoder plus SoD graph correlation) achieved an AUC of 0.96 against the labeled holdout set, compared to 0.74 for the partners' pre-existing rule-based detection logic (Figure 9). Table 7 reports precision and recall at the operating threshold selected to target a sub-15% false-positive rate, the threshold ultimately adopted in production by nine of eleven pilot partners.



Table.7. Classifier Performance at Selected Operating Threshold

Model	AUC	Precision	Recall	F1 Score
Rule-based baseline	0.74	0.61	0.58	0.59
Sequence autoencoder only	0.89	0.79	0.83	0.81
SoD graph correlation only	0.85	0.88	0.66	0.75
Combined (proposed)	0.96	0.91	0.89	0.90

10.5 Compliance Coverage

Automated, continuous evidence collection expanded coverage across all six in-scope frameworks, most notably for the Cloud Security Alliance CCM v4 (50% to 87%) and PCI DSS 4.0 (64% to 94%), the latter aided by PCI DSS 4.0's explicit accommodation of continuous, telemetry-based validation under its customized-approach provisions.




Figure 7. Compliance framework coverage before and after twelve months of architecture deployment

10.6 Remediation Efficiency

MTTR reductions ranged from 69.4% (critical patch deployment, constrained by change-window availability rather than detection speed) to 91.5% (audit evidence collection, where automation eliminated most manual evidence-gathering effort entirely). Full results appear in Table 5 and Figure 10 in Section 7.2.

XI. DISCUSSION

The results support three interpretive claims. First, signal fusion materially outperforms either threat-intelligence-driven risk scoring or behavioral anomaly detection in isolation: the combined classifier's AUC (0.96) exceeds both single-



signal variants (0.89 and 0.85, respectively) by a margin larger than the difference between either single-signal variant and the rule-based baseline. This suggests that the two signal families capture partially independent information - threat intelligence reflects external attacker capability and intent, while behavioral analytics reflects internal deviation from established norms - and that architectures relying on only one family leave detectable risk on the table.

Second, the temporal-decay mechanism appears to address alert fatigue more effectively than static suppression rules. Pilot partners that had previously attempted false-positive reduction through manually tuned suppression lists reported that those lists required constant maintenance and occasionally suppressed genuine findings; the decay-based approach requires no manual list maintenance and, because it operates on confidence rather than identity, does not create blind spots for specific users or transaction types.

Third, the disproportionately large MTTR improvement in audit evidence collection (91.5%) relative to patch deployment (69.4%) illustrates a broader pattern: automation compresses the components of remediation time that are procedural and evidentiary far more effectively than the components that are operationally constrained by change-management windows or vendor patch release cadence. Organizations should calibrate expectations accordingly - automated compliance enforcement principally attacks the "detection-to-decision" latency, not the underlying operational latency of change execution.

11.1 Cost and Organizational Adoption Observations

Pilot partners that pre-consolidated identity federation, as noted in Section 8.2, realized faster time-to-value, reinforcing that this architecture is best understood as an amplifier of existing identity and change-management maturity rather than a substitute for it. Partners with fragmented identity landscapes experienced elevated false-positive rates in the initial three months of deployment, attributable to duplicate or unlinked identity records producing spurious behavioral baseline breaks.

XII. LIMITATIONS

- Cold-start baselining: newly onboarded entities and newly deployed custom transactions lack sufficient historical data for reliable peer-group baselining, producing an elevated false-positive rate during an approximately 60–90 day ramp period, as observed in Figure 5.
- Threat-intelligence feed quality dependency: the risk scoring engine's exploit-likelihood factor is only as accurate as the underlying CVE and KEV data; feeds with delayed or incomplete coverage of ERP-specific vendor advisories (particularly for custom code vulnerabilities) reduce scoring precision.
- Custom code opacity: sequence-model features derived from custom ABAP or Apex transactions require functional annotation to be interpretable by non-technical control owners, and this annotation effort was manually performed by pilot partners rather than automated.
- SaaS telemetry constraints: several SaaS ERP vendors expose audit data only through periodic export or polling APIs rather than real-time streams, introducing detection latency floors independent of the architecture's own processing speed (reflected in the wider latency range noted for the SaaS zone in Table 8, Section 13).
- Sample composition: while the pilot spanned eleven organizations and 34 landscape instances, all participants were mid-to-large enterprises with existing security operations functions; results may not generalize to smaller organizations without dedicated SOC capacity to action graduated enforcement escalations.

XIII. GOVERNANCE AND ORGANIZATIONAL IMPLICATIONS

Deploying this architecture shifts several responsibilities traditionally owned by internal audit and periodic access-review committees toward continuous, system-mediated enforcement. This shift requires explicit governance decisions rather than purely technical ones. Table 8 summarizes the principal governance questions pilot partners' steering committees were required to resolve prior to production rollout of graduated enforcement actions.

**Table.8. Governance Decisions Required Prior to Enforcement Rollout**

Decision Area	Question to Resolve	Typical Owner
Escalation authority	Who can approve automated hard-blocks on financial transactions?	CFO / Controller
Override accountability	Who is accountable for logged supervisory overrides?	Business process owner
Exception handling	How are legitimate emergency-access (firefighter) scenarios exempted?	Security operations
Evidence retention	How long is automated evidence retained for audit defensibility?	Internal audit / Legal
Model change control	Who approves changes to scoring weights and detection thresholds?	Risk committee

Several partners established a joint security-and-audit working group specifically to own the model change-control process, recognizing that adjustments to scoring weights or detection thresholds constitute changes to the control environment itself and therefore warrant the same change-management rigor as changes to ERP configuration.

XIV. FUTURE WORK

- Federated learning across organizations to accelerate peer-group baseline maturity for newly onboarded entities without requiring cross-organization sharing of raw transaction data;
- Extension of the sequence-modeling approach to natively reason over multi-system transaction chains that span ERP, treasury management, and banking-connectivity platforms, rather than treating each system's logs independently;
- Integration of large-language-model-assisted control narrative generation, allowing automatically detected policy violations to be translated into audit-ready narrative evidence with minimal manual drafting;
- Formal adversarial robustness testing of the sequence autoencoder against slow-drift evasion strategies, in which an adversary incrementally shifts behavior to remain within a continuously re-baselined normal range;
- Expansion of the pilot population to include smaller organizations and managed-service-provider-operated landscapes to assess generalizability beyond enterprises with dedicated SOC capacity.

XV. CONCLUSION

This article has presented an integrated security architecture for cloud-native and hybrid ERP ecosystems that fuses external threat intelligence with internal behavioral telemetry to drive dynamic risk scoring, behavioral anomaly detection, and automated compliance enforcement. Evaluated over an 18-month, eleven-organization pilot, the architecture reduced median detection latency by 81–83% across five threat categories, reduced the false-positive alert rate from 38% to 10.1%, and cut mean time to remediate compliance exceptions by 69–92% depending on control family, while expanding automated, continuous evidence coverage across six major regulatory and control frameworks. The central finding - that fusing threat-intelligence and behavioral signals outperforms either in isolation - argues for treating ERP security as a signal-fusion problem rather than a signature- or rule-matching problem, and the governance analysis in Section 13 indicates that realizing these gains in practice requires deliberate organizational decisions about escalation authority and model change control, not only technical deployment. Organizations planning similar initiatives should prioritize identity federation consolidation as a prerequisite work stream and should expect a 60–90 day baselining ramp before false-positive rates approach steady-state levels.

Appendix A - Risk Scoring Factor Reference

The following reference table expands the four composite risk-scoring factors introduced in Section 5 into their constituent data inputs, as implemented during the pilot.



Factor	Constituent Inputs	Update Frequency
Asset Criticality (Ac)	Data classification tier, financial materiality, process dependency count	Quarterly review, event-triggered update
Exploit Likelihood (El)	CVSS base/temporal score, CISA KEV status, threat-feed exploitation reports	Real-time on feed update
Exposure (Ex)	Internet-facing status, authentication strength, entity access count	15-minute cycle
Control Mitigation (Cm)	Patch status, network segmentation tier, MFA enforcement status	15-minute cycle


Appendix B - Control Mapping Matrix

Table 4 (Section 7.1) presented a five-row excerpt of the full control-to-policy mapping matrix maintained during the pilot. The complete matrix contained 214 individual policy mappings distributed across framework as summarized below.

Framework	Mapped Policies	Automated Evidence Coverage (Post-Pilot)
NIST CSF 2.0	58	92%
ISO/IEC 27001:2022	41	90%
SOX ITGC	37	88%
GDPR Article 32	29	91%
PCI DSS 4.0	26	94%
CSA CCM v4	23	87%

REFERENCES

- [1] Cloud Security Alliance. (2023). Cloud Controls Matrix v4.0.13. Cloud Security Alliance.
- [2] Committee of Sponsoring Organizations of the Treadway Commission. (2013). Internal Control - Integrated Framework. COSO.
- [3] European Union. (2016). General Data Protection Regulation, Article 32: Security of Processing. Official Journal of the European Union.
- [4] Freund, J., & Jones, J. (2014). Measuring and Managing Information Risk: A FAIR Approach. Butterworth-Heinemann.
- [5] International Organization for Standardization. (2022). ISO/IEC 27001:2022 - Information Security, Cybersecurity and Privacy Protection - Information Security Management Systems - Requirements. ISO.
- [6] MITRE Corporation. (2024). ATT&CK for Enterprise, Version 15. MITRE.
- [7] National Institute of Standards and Technology. (2024). NIST Cybersecurity Framework 2.0 (NIST CSWP 29). U.S. Department of Commerce.
- [8] National Institute of Standards and Technology. (2020). NIST Special Publication 800-53 Revision 5: Security and Privacy Controls for Information Systems and Organizations. U.S. Department of Commerce.
- [9] Open Web Application Security Project. (2023). OWASP Top 10:2021 and Enterprise Application Security Verification Standard 4.0.3. OWASP Foundation.
- [10] Payment Card Industry Security Standards Council. (2024). PCI Data Security Standard, Version 4.0.1. PCI SSC.
- [11] SAP SE. (2023). SAP Governance, Risk, and Compliance 12.0: Access Control Administration Guide. SAP SE.
- [12] U.S. Cybersecurity and Infrastructure Security Agency. (2024). Known Exploited Vulnerabilities Catalog. CISA.
- [13] U.S. Securities and Exchange Commission. (2003). Sarbanes-Oxley Act Section 404 Management Assessment of Internal Controls, 17 C.F.R. § 229.308.
- [14] Verizon. (2024). 2024 Data Breach Investigations Report. Verizon Business.
- [15] Chuvakin, A., & Schmidt, K. (2013). Logging and Log Management: The Authoritative Guide to Understanding the Concepts Surrounding Logging and Log Management. Syngress.
- [16] Kark, K., & Stiennon, R. (2022). User and Entity Behavior Analytics: Market Landscape and Adoption Patterns. Forrester Research.
- [17] Panetta, K. (2023). Gartner Market Guide for SOAR Solutions. Gartner, Inc.
- [18] Open Policy Agent Project. (2023). Rego Policy Language Reference, OPA v0.58. Cloud Native Computing Foundation.



INNO  SPACE
SJIF Scientific Journal Impact Factor



ISSN INTERNATIONAL
STANDARD
SERIAL
NUMBER
INDIA



International Journal of Advanced Research

in Electrical, Electronics and Instrumentation Engineering

 9940 572 462  6381 907 438  ijareeie@gmail.com



www.ijareeie.com

Scan to save the contact details