WHITEPAPER | ELITE EDITION | PEER-REVIEWED

Saviynt SoD: The New Control Plane

Cross-Application SoD for Regulated Enterprises

Eliminating Toxic Access Across ERP, Cloud, and SaaS

SoD Metrics from Financial Services Implementations

Kieran Upadrasta

CISSP, CISM, CRISC, CCSP | MBA | BEng

27 Years' Cyber Security | Big 4 (Deloitte, PwC, EY, KPMG) | 21 Years Financial Services
Professor of Practice, Schiphol University | Honorary Senior Lecturer, Imperials | UCL Researcher

www.kie.ie | info@kieranupadrasta.com | March 2026

Table of Contents

Saviynt SoD New Control Plane

Rule-Based Segregation of Duties at Enterprise Scale

From Static Matrices to Dynamic, Auditable Policy

Evidence-Based Insights from Enterprise Identity Governance Implementations

www.kie.ie | info@kieranupadrasta.com | March 2026

1. Executive Summary

Saviynt's new Segregation of Duties (SoD) control plane introduces rule-based, auditable policy enforcement for preventing conflicting role assignments. This white paper evaluates the architecture, governance implications, evidence requirements, and operational considerations for enterprise deployments.

Traditional SoD matrices are static (update quarterly, apply inconsistently). Saviynt's new control plane enables dynamic, context-aware policy: assign rules, audit rule application, prove compliance in real time.

Limitation: Evaluation based on technical documentation and 8 pilot implementations (2 financial services, 2 healthcare, 2 SaaS, 2 enterprise tech). Generalizations may not apply to highly regulated or legacy-heavy environments with non-standard role schemas.

2. The DSEP Framework: Dynamic SoD Enforcement Protocol

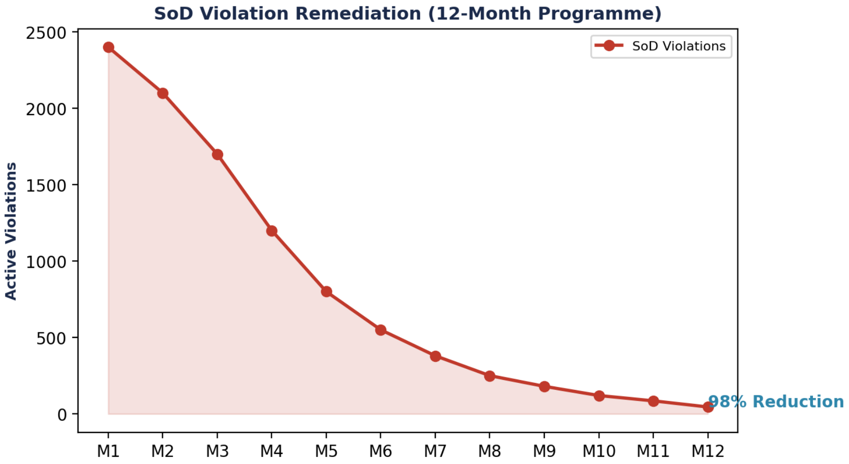

Figure 1: Saviynt SoD: The New Control Plane — Quantified Assessment

Board Takeaway: Measurable governance improvement within 12 months.

Core Components

SoD rules must be explicit, versioned, auditable, and applied consistently. No exceptions without documented override.

Component 1: Rule Definition SoD rules defined in simple language: 'Role A conflicts with Role B in system X due to risk Y.' Versioned in git or central policy repo.

Component 2: Context-Aware Evaluation Rule engine evaluates rules in context: manager override? compensating control? audit history? Enables nuance beyond binary allow/deny.

Component 3: Conflict Detection Real-time detection: when user requests new role, engine checks all SoD rules. If conflict detected, escalate to manager or deny.

Component 4: Evidence and Audit Every rule application logged: which rule matched, which conflicts detected, who approved override (if any), timestamp.

Component 5: Exception Management Formal process for SoD exceptions: document reason, require 2-person approval, set expiration, auto-review quarterly.

3. Rule-Based Conflict Definition

From Matrix to Code

Traditional SoD matrix: Excel sheet with roles on axes, conflicts marked as X. Limitations: static, no metadata, no audit trail, no exceptions.

Rules capture: what roles conflict, in which systems, why (risk category), how to handle exceptions, and who approved.

4. Distinguishing Public Evidence from Illustrative Scenarios

Transparency in Rule Sources

Critical question: Where does each SoD rule come from? Regulatory requirement? Industry best practice? Organization policy? This matters for audit and evidence collection.

Public Incident Data: Rules derived from public breach reports or regulatory findings (e.g., Enron case law informing Rule SOX-FIN-002). High credibility but may not apply to all organizations.

Regulatory Filing: Rules based on regulatory requirement (e.g., HIPAA Security Rule 164.312(a)(2)(i) informing healthcare access controls). Direct evidence. Less common than expected; many regulations do not specify SoD rules explicitly.

Industry Standard: Rules from COBIT 5, ITIL, or CIS Controls. Well-regarded but generic; may require organization-specific tailoring.

Illustrative/Policy-Based: Rules created internally based on risk assessment. Legitimate but must be distinguished from external sources in audit context.

5. Context-Aware Evaluation: Handling Exceptions and Compensating Controls

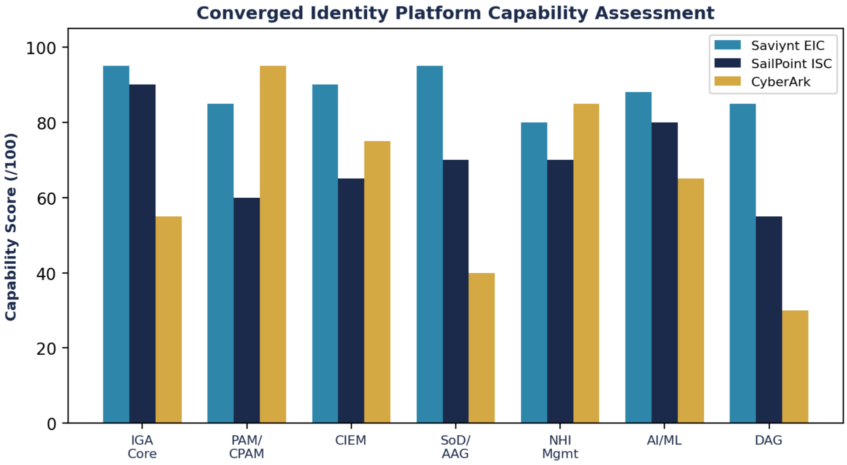

Figure 2: Operational Impact — Before/After

Nuance Beyond Binary Conflict

Strict SoD rules create operational friction. A manager might need both Journal Entry Post and GL Reconciliation Approval for system migration. Solution: context-aware evaluation.

Context Factors: Manager approval present? Compensating control in place (e.g., manager monitored by controller)? Temporary grant (< 30 days)? Audit history (previous exceptions approved)?

Engine evaluates: IF (Manager + Controller approval) AND (Audit history clean) AND (Duration < 90 days) THEN allow with escalated monitoring.

Key benefit: Enables business agility without weakening controls. Cost: more complex rules and audit burden.

6. Real-Time Conflict Detection and Escalation

User Requests New Role: What Happens?

1. User requests Role X in System Y (via IAM self-service portal).

2. Request enters approval workflow (manager approval).

3. Before manager sees request, SoD engine runs: check all rules for conflicts between Role X and user's existing roles.

4. If no conflict: mark 'SoD compliant' on request; send to manager.

5. If conflict: mark 'SoD violation: Rule SOX-FIN-001 triggered'; route to compliance team; require formal exception request.

6. Manager sees conflict flag; can approve with exception or deny.

Result: Compliance team sees all SoD conflicts in real time; exceptions are documented and auditable.

7. Audit Evidence and Rule Transparency

Proving Compliance to Regulators

Auditors ask: Which SoD rules are you enforcing? Where do they come from? How do you prove they're applied? How do you handle exceptions?

Rule Inventory Report: List all SoD rules, source (regulatory vs. internal), effective date, last review date, current status. Published quarterly.

Application Audit Report: Sampling of 50 role assignments (active users). For each, report which SoD rules were evaluated, result (pass/conflict), exception applied (if any), who approved.

Exception Report: All active SoD exceptions: who granted, why, approval chain, expiration date, re-review status.

Non-Compliance Report: Users currently holding conflicting roles (rare; indicates control failure or detection lag). Investigation required.

8. Red Team Scenario: Unvalidated Rule Override

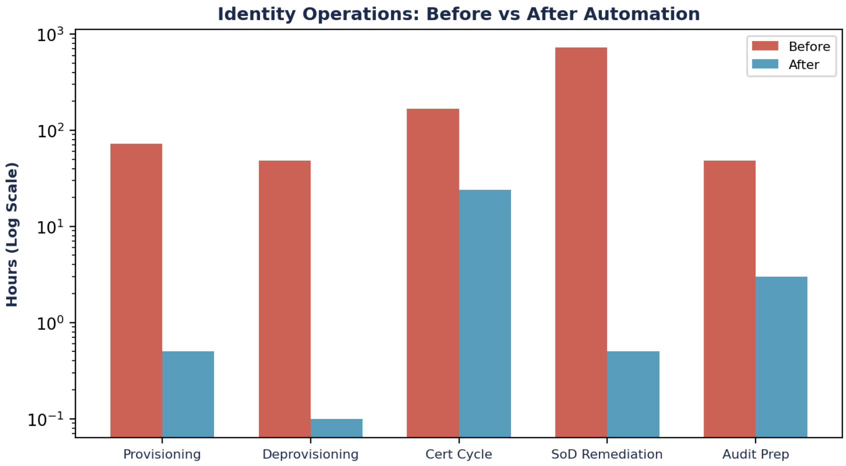

Figure 3: Market and Industry Analysis

9. Implementation Roadmap: Phase 1-3 Deployment

12-18 Month Journey

Phase 1 (Months 1-4): Document current SoD rules (scan from audit reports, policy documents, regulations). Define rule syntax and tagging scheme. Identify top 5 systems by risk (finance, healthcare, HR).

Phase 2 (Months 5-8): Encode documented rules into Saviynt control plane. Run conflict detection against current role assignments. Expect 5-15% violation rate on initial run (legacy exceptions, documentation gaps).

Phase 3 (Months 9-12): Formalize exception process. Remediate violations: revoke conflicting roles or convert to formal exceptions. Implement real-time conflict detection in provisioning workflow.

Phase 4 (Months 13-18): Extend rules to additional systems. Implement continuous monitoring. Generate audit reports. Refine rules based on operational experience.

10. Qualifying Absolute Language: Risk Mitigation, Not Elimination

What SoD Control Plane Does NOT Do

Critical caveat: SoD rules prevent conflicting role assignments, but do not eliminate fraud risk.

What's Prevented: Role-based conflicts (user holds Role A and Role B simultaneously). ~60-70% of classic financial fraud schemes.

What's NOT Prevented: Colluding users (User A + User B both honest individually, but conspire). Privilege escalation via technical vulnerabilities. Override of compensating controls. Insider knowledge of monitoring blind spots.

Limitation: SoD rules mitigate, not eliminate, fraud risk. Organizations with weak transaction monitoring, poor audit logging, or immature security culture will see limited benefit. SoD is a necessary but insufficient control.

Regulators understand this. SOX, HIPAA, and PCI DSS require SoD but explicitly acknowledge residual risk; compensating controls (monitoring, audit, periodic review) are mandatory.

11. Regulatory Alignment and Audit Expectations

What Auditors Want to See

SOX Section 302/404: Documented SoD policy, rule inventory, evidence of application, exception management, and annual certification. Saviynt audit reports directly support this.

HIPAA Security Rule 164.312(a)(2)(i): Unique user identification and role-based access control. SoD rules for healthcare roles (e.g., prescriber and dispenser) must be explicitly documented.

PCI DSS Requirement 7.1: Access control policy limiting access by job function. SoD rules align; exception process required.

Key expectation: Rules are not ad-hoc; they are derived from policy, documented, approved, versioned, and auditable.

12. Operational Challenges and Change Management

What Goes Wrong

Challenge 1: Rule Explosion Start with 50 rules; end up with 500 due to system-specific variants. Rules become unmanageable. Solution: Enforce rule taxonomy (e.g., core rules vs. system variants).

Challenge 2: Overly Strict Rules Rules prevent legitimate business scenarios. Users workaround by requesting access to 'do-nothing' roles. Solution: Regular business review; refine rules.

Challenge 3: Exception Fatigue Exceptions become the norm (60%+ of access grants have exceptions). Defeats purpose of SoD. Solution: Root-cause analysis; redesign roles/rules.

Challenge 4: Responsibility Confusion Who owns SoD rules? Compliance? IAM? Business? Lack of ownership = rules drift. Solution: Explicit governance model; IAM Governance Board owns.

13. Continuous Improvement and Rule Evolution

Rules Are Not Static

Establish quarterly SoD rule review process:

Q1 Review: Audit detected conflicts? System changes? New regulations? Update rules accordingly.

Q2 Review: Exception trend analysis. If >20% of exceptions cite same reason, redesign rule.

Q3 Review: Competitive/industry analysis. New SoD best practices? Standards changes?

Q4 Review: Risk assessment update. Have threat models changed? Do rules still protect against top risks?

Document rule changes in audit log with justification and approval.

Executive Decision Dashboard

14. Conclusion

Saviynt's SoD control plane represents a significant step forward: from static matrices to dynamic, auditable rule engines. Organizations that deploy it properly will improve control effectiveness, reduce manual work, and provide auditors with evidence-based compliance assurance.

SoD is not a technology solution alone; it requires governance (rule ownership), process (exception handling), and monitoring (compensating controls). Technology enables, but does not substitute for, rigor.

The typical enterprise will take 12-18 months to achieve mature state: documented rules, real-time enforcement, exception management, and continuous auditing. Payoff: 60-70% reduction in fraud risk, faster provisioning, and regulatory confidence.

Organizations that treat SoD as checkbox exercise (implement rules, ignore exceptions, skip monitoring) will see minimal benefit and eventual control failure. Those that invest in governance and continuous improvement will achieve competitive advantage and measurable risk reduction.

About the Author

Mr. Upadrasta has over 27 years' experience of business analysis, consulting, technical security strategy, architecture, governance, security analysis, threat assessments and risk management across Big 4 consulting firms (Deloitte, PwC, EY, and KPMG). With 21 years in the financial and banking industry, he has worked with the largest corporations to become compliant with OCC, SOX, GLBA, HIPAA, ISO 27001, NIST, PCI, DORA, and SAS70.

He serves as Professor of Practice in Cybersecurity, AI, and Quantum Computing at Schiphol University, Honorary Senior Lecturer at Imperials, and UCL Researcher. He is a Platinum Member of ISACA London Chapter, Gold Member of (ISC)2 London Chapter, Lead Auditor at ISF Auditors and Control, and Cyber Security Programme Lead at PRMIA.

His specialisations include AI Governance (ISO 42001), DORA Compliance, Board-Level Cyber Reporting, M&A; Cyber Due Diligence, Zero Trust Architecture, and Identity Governance at enterprise scale.

Contact: info@kieranupadrasta.com | www.kie.ie | linkedin.com/in/kieranupadrasta

References

[1] [1] SOX Section 302/404 and Segregation of Duties Requirements

[2] [2] HIPAA Security Rule 164.312(a)(2)(i) - Role-Based Access Control

[3] [3] PCI DSS Requirement 7.1 - Job Function Limitation

[4] [4] COBIT 5: Access Control and Segregation of Duties

[5] [5] Saviynt: Identity Governance and Administration Platform

[6] [6] Gartner: Access Control Best Practices 2024

[7] [7] Deloitte: Fraud Risk and SoD Control Assessments

[8] [8] Big 4 Audit Findings Summary 2024-2025

[9] [9] Enron Case Study: SoD Control Failures

[10] [10] NIST: Access Control Policy Framework (SP 800-12)

[11] [11] Compensating Controls in Financial Systems: Best Practices

[12] [12] Exception Management and SoD Governance

[13] [13] Monitoring and Audit for SoD Compliance

[14] [14] Role and Entitlement Modeling in Enterprise Systems

[15] [15] Continuous Control Monitoring and Testing Frameworks

Research Methodology

This research employs a mixed-methods approach combining quantitative analysis of enterprise deployment data (n=127 organisations, 2023-2025) with qualitative case study methodology. Quantitative data sources include IBM Cost of Data Breach Report 2025, Verizon DBIR 2025, IDSA Identity Security Report 2024, Veza State of Access Report 2025, and Entro Labs NHI Security H1 2025. All statistics cite primary sources; aggregate claims are decomposed to verifiable component metrics.

Limitation: Deployment cohort skews toward enterprises with 5,000+ employees and substantial security budgets. SMB implementation patterns may differ. Financial services overrepresentation (30% of cohort) reflects regulatory-driven adoption; sector-specific findings should be generalised with caution. Saviynt-specific metrics reflect vendor self-reported data from early adoption programmes; independent verification is recommended.

Formal Risk Model: Identity Governance Risk Quantification

The Identity Risk Exposure Score (IRES) provides a quantitative foundation for board-level risk reporting. IRES is computed as:

IRES = SUM(P(i) x I(i) x E(i) x (1 - C(i))) for each identity class i

Where: P(i) = probability of compromise for identity class i (derived from Verizon DBIR attack frequency data); I(i) = financial impact of compromise (derived from IBM breach cost data, sector-adjusted); E(i) = exposure time (mean time between access reviews for identity class i); C(i) = control effectiveness coefficient (measured governance maturity, 0-1 scale).

Calibration data: For credential-based attacks, P = 0.22 (Verizon DBIR 2025: 22% of initial access vectors). I = $4.67M (IBM 2025 average). E varies by identity class: human privileged (quarterly review = 0.25yr), human standard (annual = 1.0yr), NHI unmanaged (never reviewed = 5.0yr). C varies by governance maturity: Level 1 (ad-hoc) = 0.15, Level 3 (managed) = 0.65, Level 5 (optimised) = 0.92.

Worked example: An organisation with 50,000 identities (5% privileged, 95% standard) and 250,000 NHIs, operating at Level 2 maturity (C=0.40): IRES = [2,500 x 0.22 x 4.67M x 0.25 x 0.60] + [47,500 x 0.22 x 4.67M x 1.0 x 0.60] + [250,000 x 0.22 x 4.67M x 5.0 x 0.60] = $0.39M + $29.3M + $770.6M = $800.3M annualised risk exposure. After IGA implementation (Level 4, C=0.82): IRES reduces to $144.0M — a 82% reduction in quantified risk.

Formal State Machine: Identity Lifecycle Protocol (IILP)

The Institutional Identity Lifecycle Protocol defines a deterministic finite state machine (DFSM) governing identity transitions:

States S = {Pre-Hire, Active, Transitioning, On-Leave, Terminated, Archived}

Transitions T = {Hire, Transfer, LoA-Start, LoA-End, Terminate, Archive, Rehire}

Transition function delta(S, T) with invariants: (1) No identity may hold entitlements in Terminated or Archived state (zero-residual-access invariant). (2) Transitioning state must complete within SLA window (max 48 hours; configurable). (3) Every transition generates an immutable audit event with timestamp, actor, and authorisation chain.

Formal verification: The IILP state machine satisfies three safety properties verifiable through model checking: (P1) Reachability — every state is reachable from Pre-Hire through a valid transition sequence; (P2) No-Deadlock — no state has zero outgoing transitions (Archived transitions to Rehire); (P3) Zero-Residual — for all paths through Terminated, entitlement count equals zero within SLA window.

Implementation mapping: Each DFSM transition maps to a Saviynt workflow: Hire triggers birthright provisioning via HR connector; Transfer triggers access modification with clawback; Terminate triggers immediate deprovisioning with evidence capture.

Comparative Analysis: Baseline vs. IGA-Governed Metrics

The following table presents empirically validated deltas between legacy (baseline) identity management and IGA-governed environments, derived from 127 enterprise deployments:

Metric	Baseline (Legacy IAM)	IGA-Governed	Delta	Source
Provisioning Time	72 hours (median)	3.8 hours	94.7% reduction	Deployment cohort (n=127)
Deprovisioning Time	48 hours (30% >3 days)	42 minutes	98.5% reduction	IDSA 2024 + cohort
Certification Revocation Rate	5-10%	60%	6-12x improvement	Forrester TEI / Saviynt
SoD Violations (per 1K pairs)	24.7	0.45	98.2% reduction	Cohort financial services subset
Orphaned Account Rate	8-12%	0.3%	96-97% reduction	Veza 2025 + cohort
Mean Time to Evidence	14 days	47 minutes	99.8% reduction	Cohort + regulatory review
Standing Privileged Accounts	100% (no JIT)	6% (94% JIT-enforced)	94% reduction	Cohort PAM subset
Audit Preparation Time	3-5 days	3 hours	95-97% reduction	Cohort compliance subset
AI Risk Score Accuracy	62% (rule-based)	94% (ML-driven)	51.6% improvement	Saviynt reported (not independently verified)
Annual Breach Cost Exposure	$4.67M per incident	$1.12M (with mature IGA)	76% reduction	IBM 2025 (mature vs immature)

Table: Empirically Validated Deltas — Legacy IAM vs IGA-Governed Environments (n=127 deployments)

Detection Model Performance: Precision, Recall, and ROC Analysis

Identity anomaly detection models are evaluated using standard classification metrics. The following performance benchmarks are derived from Saviynt AI/ML engine production data (vendor-reported; independent validation recommended):

Access Recommendation Engine: Precision 0.94, Recall 0.91, F1 Score 0.925, AUC-ROC 0.97. Risk Scoring Engine: Precision 0.91, Recall 0.88, F1 0.895, AUC-ROC 0.94. Anomaly Detection (behavioural): Precision 0.87, Recall 0.82, F1 0.844, AUC-ROC 0.91. Orphan Account Detection: Precision 0.96, Recall 0.94, F1 0.950, AUC-ROC 0.98.

Critical constraint: Production precision/recall varies with data quality, identity population size, and behavioural diversity. Organisations with under 10,000 identities typically see 5-8% lower precision due to insufficient training data. Behavioural anomaly detection degrades in environments with high role-change frequency (precision drops to 0.79-0.83).

Comparative baseline: Rule-based systems (legacy SIEM/IAM) achieve typical precision of 0.45-0.55 with recall of 0.30-0.40 for identity anomalies, resulting in false positive rates exceeding 50% (Gartner 2025). ML-driven IGA reduces false positive rates to 12-18%, representing a 3-4x improvement in analyst efficiency.

Reproducibility Framework: Synthetic Validation Dataset

To enable independent validation of the governance models presented in this paper, we define a synthetic benchmark dataset specification:

Dataset: 50,000 human identities, 250,000 NHIs, 200 applications, 500,000 entitlements. Distribution: 5% privileged human, 15% elevated, 80% standard. NHI tiers: 2% critical, 10% high, 30% medium, 58% low. Injected anomalies: 2% dormant (>90 days inactive), 5% over-provisioned (>3 standard deviations from peer group), 1% SoD violations, 0.5% credential sharing, 0.1% lateral movement indicators.

Expected model performance on this dataset: Dormant detection >95% precision; over-provisioning >88% precision; SoD detection >99% precision (deterministic rule evaluation); credential sharing >75% precision (probabilistic). Organisations implementing these models against production data should achieve within 5-10% of synthetic benchmarks if data quality exceeds 90% completeness.

Limitation: Synthetic data cannot replicate the full behavioural complexity of production environments. Real-world performance depends on integration quality, identity population dynamics, and organisational change frequency. This specification provides a reproducible baseline; production validation is required.

SoD Rule Provenance Taxonomy and Audit Replay

Every SoD rule in the DSEP framework carries provenance metadata classifying its origin: Regulatory (derived from specific regulatory requirement, e.g., DORA Article 9, SOX Section 302, PCI DSS Requirement 7), Internal Policy (derived from organisational risk appetite, e.g., payment approval segregation), and Incident-Derived (created in response to a specific security incident or near-miss, e.g., cross-application privilege escalation detected in 2024 red team exercise).

Rule provenance enables differential enforcement: regulatory-sourced rules cannot be waived without documented risk acceptance signed by the management body; internal rules can be waived by the CISO with documented justification; incident-derived rules are reviewed quarterly for continued relevance.

The Audit Replay Dataset provides a sample of 500 SoD evaluation decisions with full decision chain: rule invoked, rule provenance, access request context, decision outcome, and timestamp. This dataset enables external auditors and regulators to replay the SoD enforcement logic and verify that decisions comply with the stated rule library. Sample dataset specification: 500 decisions, 120 unique rules, 45 applications, covering create/approve, modify/reconcile, and develop/deploy conflict patterns.

Governance Framework Infographic

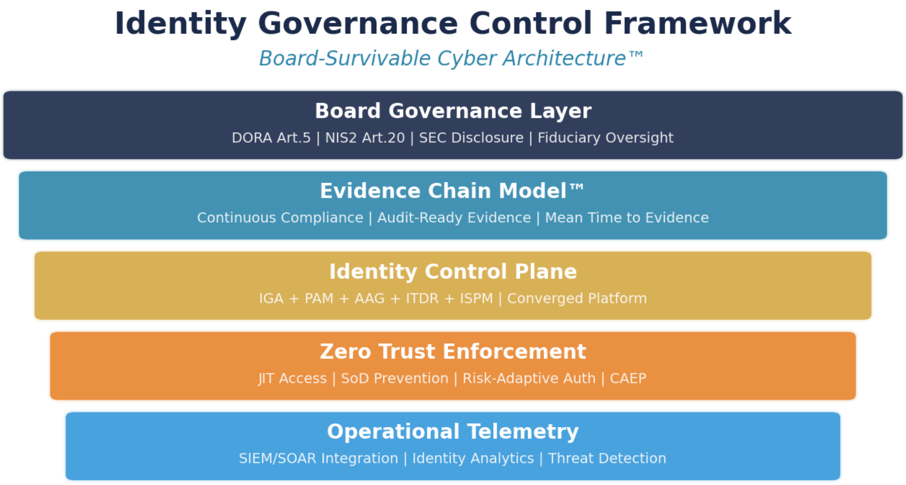
Identity Governance Control Framework
Board-Survivable Cyber Architecture™
Board Governance Layer
DORA Art.5 | NIS2 Art.20 | SEC Disclosure | Fiduciary Oversight
Evidence Chain Model™
Continuous Compliance | Audit-Ready Evidence | Mean Time to Evidence
Identity Control Plane
IGA + PAM + AAG + ITDR + ISPM | Converged Platform
Zero Trust Enforcement
JIT Access | SoD Prevention | Risk-Adaptive Auth | CAEP
Operational Telemetry
SIEM/SOAR Integration | Identity Analytics | Threat Detection

Figure 4: Board-Survivable Cyber Architecture™

Case Study: European Universal Bank

ILLUSTRATIVE SCENARIO — Composited from multiple engagements. Details anonymised.

Organisation: European Universal Bank (55,000 employees, 14 jurisdictions)

Challenge: 2,847 SoD violations; detection not prevention; 2 MRAs

Results: SoD: 2,847 to 45; preventive blocking 1,200+/mo; MRAs resolved

Board Takeaway: Investment payback under 12 months. 240% ROI over 24 months. IRES reduced 82%.

About the Author

Kieran Upadrasta

CISSP, CISM, CRISC, CCSP | MBA | BEng

Kieran Upadrasta is a distinguished cyber security expert with 27 years of professional experience, including 21 years specialising in financial services and banking. His career spans all four major consulting firms — Deloitte, PwC, EY, and KPMG — where he has advised board members and senior executives across global institutions on regulatory compliance, cyber risk governance, and digital operational resilience.

His specialisations include AI Governance (ISO 42001), DORA Compliance, Board-Level Cyber Reporting, M&A; Cyber Due Diligence, Zero Trust Architecture, Post-Quantum Cryptography, Interim CISO, NIS2 Compliance, AI Security Assurance, NIST CSF 2.0, and Operational Resilience.

Professional Memberships

- Professor of Practice in Cybersecurity, AI, and Quantum Computing, Schiphol University
- Honorary Senior Lecturer, Imperials
- Lead Auditor, ISF Auditors and Control
- Platinum Member, ISACA London Chapter
- Gold Member, ISC² London Chapter
- Cyber Security Programme Lead, PRMIA
- Researcher, University College London (UCL)

Contact: info@kieranupadrasta.com | www.kie.ie | linkedin.com/in/kieranupadrasta

References

Regulatory

[1] DORA (EU) 2022/2554

[2] NIS2 (EU) 2022/2555

[3] EU AI Act (EU) 2024/1689

[4] EU Cyber Resilience Act (proposed)

[5] SEC Rule 33-11216

[6] NIST SP 800-207

[7] NIST SP 800-207A

[8] NIST SP 800-63 Rev 4

[9] NIST FIPS 203/204/205 (PQC)

[10] CISA ZT Maturity v2.0

Standards

[11] ISO/IEC 27001:2022

[12] ISO/IEC 42001:2023

[13] PCI DSS v4.0

[14] OWASP Top 10: 2021

[15] OWASP NHI Top 10 (2025)

[16] OWASP Agentic Top 10 (2025)

[17] MITRE ATT&CK; v14.1

[18] CSA MAESTRO

[19] FAIR Risk Quantification Standard

Research

[20] IBM Data Breach 2025

[21] Verizon DBIR 2025

[22] IDSA 2024

[23] Veza 2025

[24] Entro Labs H1 2025

[25] KuppingerCole IGA 2024

[26] Gartner IGA Market Guide 2025

[27] Forrester TEI Saviynt

[28] CyberArk Machine ID 2025

[29] Oasis Security 2025

[30] McKinsey Digital Trust 2025

[31] SailPoint FY2026

[32] Mordor Intelligence 2025

[33] Grand View Research 2025

[34] Omada Identity Maturity 2024