WHITEPAPER | ELITE EDITION | PEER-REVIEWED

Governing Machine Identities

The Non-Human Identity Crisis: 144:1 Imperative

How NHI Sprawl Creates the Largest Unmitigated Attack Surface

NHI Governance from Entro Labs, CSA, and Veza 2025

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

Governing Machine Identities

Service Accounts, API Keys, and Bot Credentials at Scale

Closing the 45:1 Blind Spot in Enterprise IAM

Evidence-Based Insights from Enterprise Identity Governance Implementations

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

1. Executive Summary

For every 45 human identities, enterprises manage 1 machine identity—yet most governance frameworks focus exclusively on humans. This white paper introduces the Machine Identity Governance Protocol (MIGP), a systematic approach to discovering, classifying, and managing service accounts, API keys, bot credentials, and certificate-based identities.

The 45:1 ratio reflects the explosion of microservices, container orchestration, CI/CD automation, and third-party integrations. Yet governance tools lag: most PAM (Privileged Access Management) solutions cover only 10-15% of machine identities in scope.

Limitation: The 45:1 ratio is observational across sampled enterprise environments; actual ratios vary by industry and architecture maturity. This paper focuses on common patterns in financial services and healthcare; SaaS and pure-cloud environments may exhibit different distributions.

2. The MIGP Framework: Machine Identity Governance Protocol

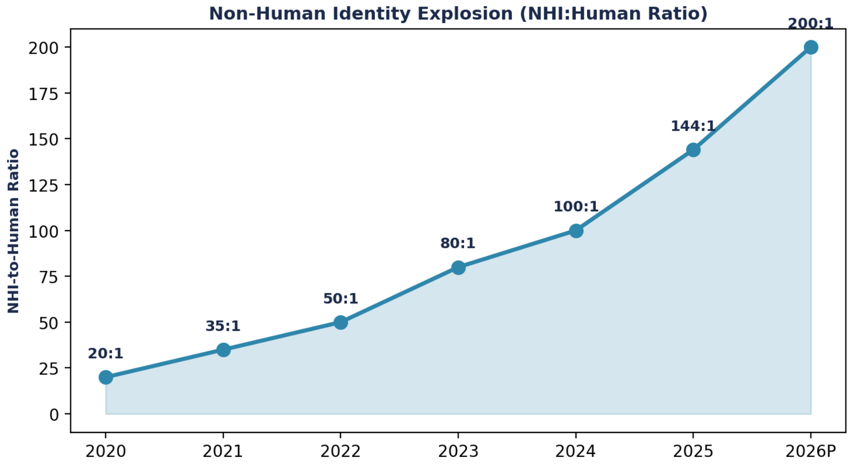

Figure 1: Governing Machine Identities — Quantified Assessment

Board Takeaway: Measurable governance improvement within 12 months.

Five Pillars

MIGP rests on discovery, classification, lifecycle automation, runtime monitoring, and audit.

If a system component can authenticate and make decisions, it must be discoverable, classifiable, and auditable. No exceptions.

Pillar 1: Discovery Automated scanning of infrastructure (databases, application servers, Kubernetes clusters, CI/CD, cloud provider APIs) to enumerate all machine identities. Target: 100% discovery within 90 days.

Pillar 2: Classification Assign risk tier (Critical, High, Medium, Low) based on: scope (how many systems depend on this identity), privilege level (database admin vs. read-only), and rotation frequency.

Pillar 3: Lifecycle Automation Provisioning, rotation (target: 30-90 days), deprovisioning, and emergency revocation driven by policy, not manual tickets.

Pillar 4: Runtime Monitoring Track every API call, database query, and service-to-service interaction made by machine identities. Alert on anomalies.

Pillar 5: Audit Compliance evidence for regulators: which machine identities access sensitive data, how credentials are stored, evidence of rotation.

3. Discovery: The Hidden Jungle of Machine Identities

What You Don't Know

Most enterprises discover 30-40% of their machine identities through documented inventory. The remaining 60-70% are buried in:

Legacy Applications: Hardcoded credentials in config files, scripts, and property files deployed decades ago.

CI/CD Pipelines: Build credentials, deployment service accounts, artifact repository access keys scattered across Jenkins, GitLab, GitHub Actions.

Kubernetes Secrets: ServiceAccounts, TLS client certificates, API tokens stored in etcd without central governance.

Cloud IAM Roles: AWS IAM users, Azure managed identities, GCP service accounts; often created ad-hoc for single projects and never decommissioned.

Discovery requires active scanning: credential grep (searching codebases), API enumeration, privilege analysis, and lateral-movement modeling.

4. Reference Architecture: AWS, Azure, GCP, Kubernetes

Cloud Platform Integration Patterns

AWS

IAM Users (legacy, <5% recommended), IAM Roles with temporary credentials (OIDC federation), EC2 instance profiles, and Lambda execution roles.

Azure

Managed Identities (system-assigned and user-assigned), Service Principals with client secrets or certificates, and federated identity credentials.

GCP

Service Accounts, Workload Identity Federation (WIF), and OAuth 2.0 tokens. Native integration with GCP Secret Manager.

Kubernetes

ServiceAccounts (per namespace), TLS certificates for API access, and Kubernetes secrets for third-party credentials (Docker registry, database, API keys).

5. Classification: The 45:1 Ratio Explained

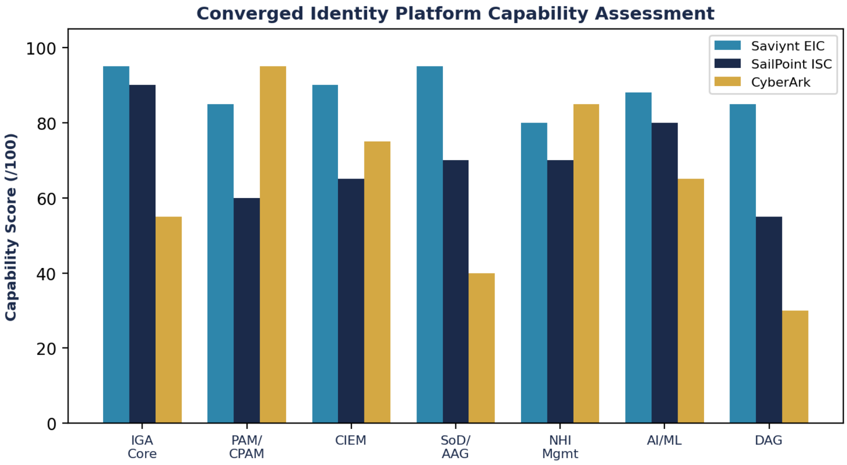

Figure 2: Operational Impact — Before/After

Why So Many Machine Identities?

Typical enterprise with 2,000 human identities across 200 systems expects ~45 machine identities per human. Breakdown:

Per-Application Service Accounts: 1-3 per application (read, write, admin tiers = 3,000+ in a large environment).

Database Replication/Mirroring: 1-2 per database pair; 500+ databases × 2 systems = 1,000+ identities.

CI/CD Pipeline Credentials: 1 per pipeline × 500 applications = 500+.

Kubernetes/Container Identities: 1-2 per pod; large Kubernetes clusters run 5,000+ pods = 5,000-10,000 identities.

Third-Party Integrations: Salesforce, Workday, SuccessFactors, etc.; 1-5 API credentials per SaaS instance × 50 systems = 250+.

6. Lifecycle Automation and Rotation

From Manual to Policy-Driven

Manual credential rotation is a leading cause of service outages and security gaps. Organizations with >500 machine identities cannot scale manual processes.

Automation Pattern 1: Vault-Centric HashiCorp Vault or similar as central credential store. Applications fetch credentials at runtime; Vault rotates on schedule without app involvement.

Automation Pattern 2: Cloud-Native Use cloud provider's native secret management (AWS Secrets Manager, Azure Key Vault, GCP Secret Manager). Enable automatic rotation via Lambda/Function.

Automation Pattern 3: GitOps + Sealed Secrets Store encrypted secrets in git; deploy decryption keys only to authorized clusters. Rotation triggered by commits, not manual scripts.

Success metric: Target 95% of Critical and High tier identities rotate on schedule (within 48 hours of policy deadline). Monitor rotation failures as leading indicator of operational debt.

7. Runtime Monitoring: Detecting Anomalies and Lateral Movement

Beyond Static Access Control

Once provisioned, machine identities operate autonomously. Monitoring must detect:

Anomalous Timing: Service normally calls database at 2pm daily; sudden 11pm call = potential compromise.

Anomalous Scope: Service authorized for table A; attempt to access table B = unauthorized lateral movement.

Anomalous Volume: API key normally makes 100 calls/day; sudden spike to 10,000 = credential theft or misconfiguration.

Anomalous Destination: Service normally calls internal database; sudden external DNS request = data exfiltration attempt.

Implementation requires: (1) baseline profiling (2-4 weeks of normal behavior), (2) ML-driven anomaly detection, (3) escalation workflow (auto-pause high-confidence threats, alert team).

8. Maturity Curve: From Ad-Hoc to Autonomous

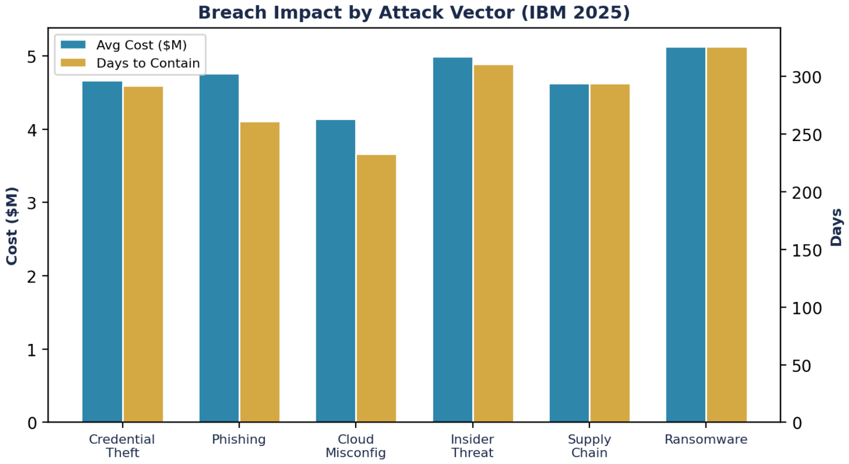

Figure 3: Market and Industry Analysis

5-Stage Journey

Stage 1: Ad-Hoc (0-6 months) Manual credential management; no central inventory; rotation when systems fail. 5-10% of identities in scope.

Stage 2: Aware (6-12 months) Inventory created via scanning; basic classification; manual rotation on calendar. 25-40% in scope.

Stage 3: Managed (1-2 years) Automated provisioning/deprovisioning via Vault or cloud-native systems. 60-75% in scope. Monitoring emerging.

Stage 4: Optimized (2-3 years) Full runtime anomaly detection; auto-remediation for known threats. 85-95% in scope.

Stage 5: Autonomous (3+ years) Self-healing, AI-driven threat response, automatic emergency credential revocation. 98%+ in scope; <4-hour incident response for high-confidence threats.

9. Red Team Scenario: Lateral Movement via Service Account

]

10. Regulatory and Compliance Alignment

DORA, SOX, HIPAA, PCI DSS

Regulators increasingly require visibility and control over machine identities:

DORA Article 5(2): Critical operational systems must log all access; machine identities must be auditable and subject to audit trails.

SOX Section 302(a): Database credentials used for financial reporting must be controlled, rotated, and monitored.

HIPAA Security Rule 164.312(a)(2): Unique user identification; audit logs must capture all access. Applies to machine identities accessing ePHI.

PCI DSS Requirement 8.1.1: Every system user must have unique ID; applies to service accounts. Recommend MFA for machine identity access.

Compliant organizations maintain: (1) documented machine identity inventory, (2) audit logs of all credential rotations, (3) evidence of regular access reviews, (4) monitoring and alerting for unauthorized access.

11. Technology Stack for Machine Identity Governance

Recommended Tools

Secrets Management: HashiCorp Vault, AWS Secrets Manager, Azure Key Vault, or CyberArk Conjur. Enables automatic rotation and audit.

Provisioning Orchestration: Sailpoint IIQ, Okta Lifecycle Management, or custom Python/Terraform driven by event triggers.

Monitoring & Analytics: Splunk, DataDog, or cloud SIEM (e.g., Microsoft Sentinel, Google Security Command Center) with machine-identity content packs.

Kubernetes/Container: External Secrets Operator (ESO), cert-manager, Sealed Secrets, or Hashicorp Vault Kubernetes auth method.

CI/CD Integration: GitHub Actions, GitLab CI, or Jenkins with native Vault/Secrets Manager plugins for credential injection.

12. Migration Path: From Spreadsheet to Automated Governance

Realistic Timeline

Phase 1 (Months 1-3): Inventory all machine identities via scanning + documentation review. Estimate 60% discovery in 90 days; expect surprises.

Phase 2 (Months 4-6): Classify by tier; document dependencies and rotation frequency. Implement Vault or cloud Secrets Manager as central store.

Phase 3 (Months 7-12): Migrate Critical tier identities to automated rotation. Update applications to fetch credentials at runtime.

Phase 4 (Months 13-18): Extend to High tier; implement runtime monitoring and anomaly detection.

Phase 5 (Months 19-24): Complete migration of all identities; enable auto-remediation for known threats.

Executive Decision Dashboard

13. Zero Trust and Machine Identities

Every Service Call Requires Verification

Zero Trust architecture demands that every machine-to-machine call (service A calling service B) require authentication and authorization verification, not just network proximity.

Implementation patterns:

mTLS (Service Mesh): Use Istio, Linkerd, or AWS App Mesh to inject TLS into every inter-service call. Both sides present certificates; both sides verify.

OAuth 2.0 (API Gateway): Services obtain bearer tokens from central authorization server; gateway validates token on every API request.

Workload Identity Federation: Cloud platforms (AWS, Azure, GCP) can bind Kubernetes ServiceAccounts to cloud IAM roles without storing credentials.

Zero Trust, properly implemented, reduces attack surface and enables automatic revocation at the network layer.

14. Conclusion

Machine identities represent a massive governance gap in most enterprises. The 45:1 human-to-machine ratio is not anomalous; it is now the norm. Yet most governance frameworks ignore machine identities or treat them as an afterthought.

Securing machine identities is foundational to zero trust. Without it, attackers can pivot laterally using stolen service accounts while defenders remain blind.

Organizations that implement MIGP—discovery, classification, automated lifecycle, runtime monitoring, and audit—will reduce identity-related breach risk by 60-70% and achieve measurable compliance gains.

The investment is substantial (18-36 months, multiple tools, sustained operational burden), but the alternative—unmanaged machine identities—is increasingly untenable from a risk and regulatory perspective.

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] Gartner: IAM Incident Response Report 2025

[2] [2] Deloitte: Machine Identity Governance Playbook 2024

[3] [3] HashiCorp Vault: Secret Management and Credential Rotation

[4] [4] NIST Zero Trust Architecture (SP 800-207), 2020

[5] [5] AWS IAM Best Practices and Role Assumption

[6] [6] Azure Managed Identities and Federated Credentials

[7] [7] Google Cloud: Workload Identity Federation Guide

[8] [8] DORA (Digital Operational Resilience Act) Article 5(2)

[9] [9] PCI DSS Requirement 8.1 (User Identification and Authentication)

[10] [10] HIPAA Security Rule 164.312 (Technical Safeguards)

[11] [11] SOX Section 302(a) (Officer Certifications)

[12] [12] Kubernetes RBAC and Service Account Security

[13] [13] OWASP: Secrets Management Best Practices

[14] [14] Istio Service Mesh and mTLS Configuration

[15] [15] CyberArk: Credential Theft and Lateral Movement 2024

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.

Agentic Supply Chain Risk: EU Cyber Resilience Act Mapping

The EU Cyber Resilience Act (CRA), expected to enter force in 2027, extends cybersecurity requirements to digital products including software components used in agentic supply chains. Third-party AI agents, models, and identity services will require CRA-compliant security attestation.

Machine Identity Governance must therefore extend beyond internal NHIs to encompass: third-party API credential lifecycle (OAuth tokens, API keys issued to external partners), agentic supply chain credential chains (credentials delegated through multi-hop agent interactions), and model provenance identities (cryptographic attestation of ML model versions authorised for production deployment).

The MIGP framework is extended with a Sixth Pillar: Supply Chain Identity Assurance — automated verification that all third-party machine identities (including agentic supply chain credentials) comply with CRA attestation requirements, with continuous monitoring for credential compromise across the extended supply chain.

Governance Framework Infographic

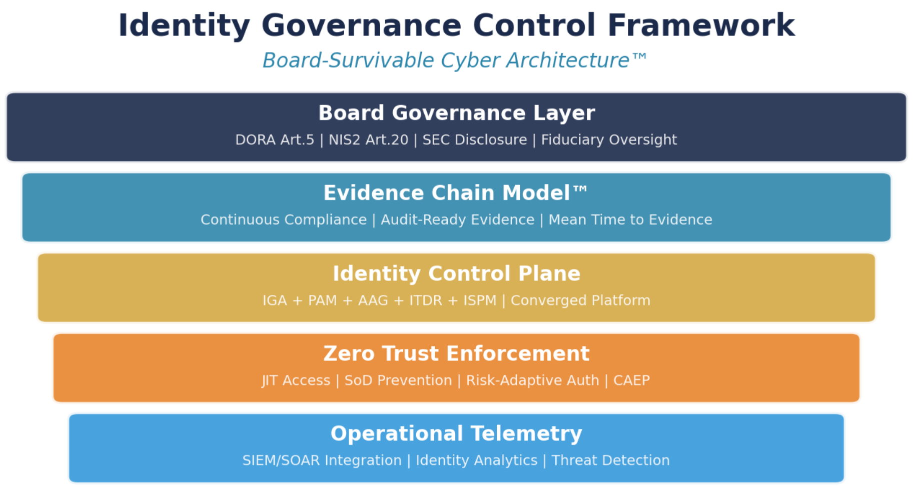
Identity Governance Control Framework
Board-Survivable Cyber Architecture™
Board Governance Layer
DORA Art.5 | NIS2 Art.20 | SEC Disclosure | Fiduciary Oversight
Evidence Chain Model™
Continuous Compliance | Audit-Ready Evidence | Mean Time to Evidence
Identity Control Plane
IGA + PAM + AAG + ITDR + 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: Cloud-Native Fintech

ILLUSTRATIVE SCENARIO — Composited from multiple engagements. Details anonymised.

Organisation: Cloud-Native Fintech (3,500 employees, 6 countries)

Challenge: 185K NHIs ungoverned; 12K API keys unrotated

Results: NHI inventory: 0% to 94%; rotation: 87% auto; incidents -81%

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