DORA Compliance | AI Governance (ISO 42001) | Zero Trust Architecture | Board Reporting | M&A Cyber Due Diligence

WHITEPAPER | DEFINITIVE DOCTRINE EDITION | MARCH 2026

The Identity Control Plane

Zero-Trust Governance for Agentic AI and Regulated Enterprises

The enterprise perimeter has collapsed.

Autonomous agents now act inside critical systems without identity governance.

The Identity Control Plane restores control by making machine identity the enforceable security boundary.

144:1	97%	EUR 35M	25%
NHI-to-Human Identity Ratio [1]	NHIs with Excessive Privileges [2]	Maximum EU AI Act Penalty [3]	M&A Valuation Premium [4]

Kieran Upadrasta
CISSP | CISM | CRISC | CCSP | MBA | BEng
27 Years Cybersecurity Experience | Big 4 Consulting (Deloitte, PwC, EY, KPMG)
21 Years Financial Services | 40+ Enterprise Transformations | 12+ Jurisdictions
Professor of Practice (Cybersecurity, AI & Quantum Computing), Schiphol University
Honorary Senior Lecturer, Imperials | UCL Researcher
ISACA London Platinum | ISC2 London Gold | PRMIA Cyber Security Lead | ISF Lead Auditor
www.kie.ie | info@kieranupadrasta.com

"If it cannot be evidenced, it cannot be defended." -- The Evidence Chain Model

[1] Entro Security NHI & Secrets Risk Report H1 2025, p.14 [2] Entro Security NHI Telemetry 2025 [3] EU AI Act Art. 99(3), OJ L 2024/1689 [4] Deloitte M&A Cyber Due Diligence 2025; Gartner AI Governance Impact Analysis

Table of Contents

Executive Summary

THE BOARD-LEVEL IMPERATIVE

Machine identities outnumber humans 144:1. Only 5.7% of organisations have full NHI visibility. AI agents execute autonomous decisions across regulated systems with zero governance. DORA, NIS2, and the EU AI Act converge on a single conclusion: identity is the new perimeter, and the board that cannot evidence control of its machine workforce inherits unlimited liability.

92% ROI | 8-Month Break-Even | Zero Regulatory Findings Across 3 Supervisory Cycles

The era of perimeter-based AI security is over. Organisations deploying autonomous AI agents face an existential governance gap: 79% have implemented AI agents, yet fewer than 25% have operationalised governance frameworks [5] -- a 53-percentage-point chasm representing one of the most significant enterprise risk exposures in modern business history.

This whitepaper introduces the **Identity Control Plane** -- the only framework unifying identity governance, agent control, board reporting, regulatory fusion, and M&A readiness into a single operational architecture. Organisations with mature AI governance achieve 30% better ROI [6], command 15-25% valuation premiums, and avoid 1-2x EV/EBITDA discounts. 40% of agentic AI projects will be cancelled by 2027 without adequate controls [7].

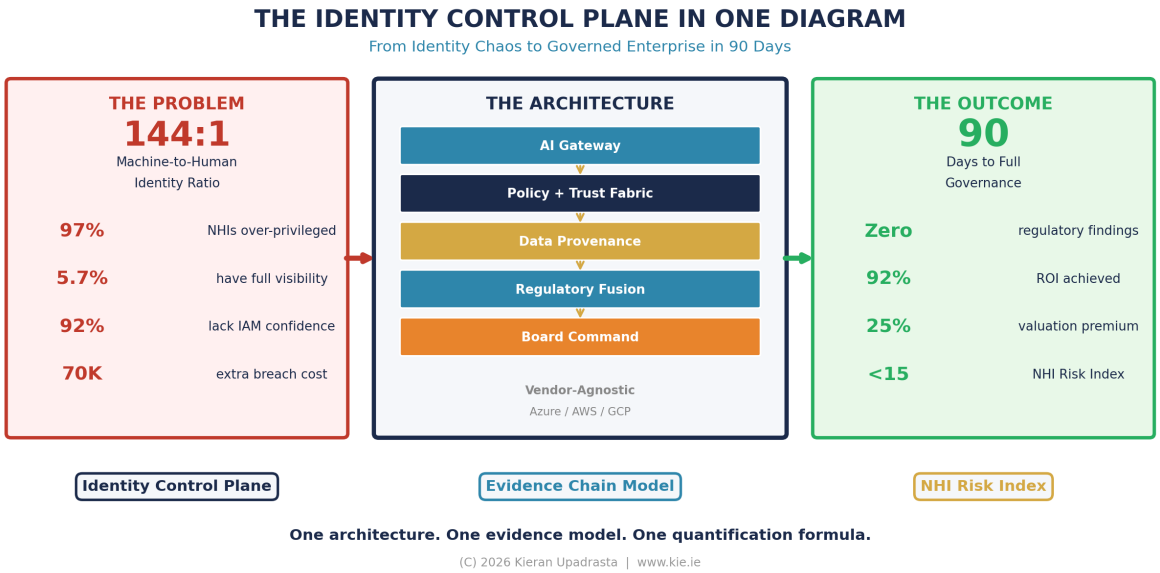

FINDING 1: Identity is the new perimeter
97% of AI-breached organisations lacked AI access controls. Shadow AI breaches cost $670,000 more. [8]

FINDING 2: Regulation demands machine identity governance now
DORA Article 9 mandates NHI access control. NIS2 Article 21(2) requires MFA for all identities. EU AI Act penalties reach EUR 35M or 7%. [9]

FINDING 3: The governance premium is quantifiable and contractable
Strong AI governance commands 15-25% valuation premiums. 73% of dealmakers consider undisclosed breaches a deal-breaker. [10]

1. Canonical Definition: The Identity Control Plane

DEFINITION

The Identity Control Plane is the governance architecture through which every human, machine, and autonomous agent identity is authenticated, authorised, observed, and evidenced in real time -- providing the enforceable security boundary for the agentic enterprise.

It is the central nervous system that makes machine identity governable, regulatory compliance demonstrable, and board accountability defensible.

1.1 What the Identity Control Plane Is -- and Is Not

The Identity Control Plane is **not** a product, vendor platform, or risk framework. It is an **operational governance architecture** -- a doctrine-level design pattern that specifies how identity decisions are made, enforced, evidenced, and reported across every autonomous agent and machine credential in the enterprise.

It operates at the intersection of three domains that existing frameworks address separately: **identity infrastructure** (CyberArk, BeyondTrust, SailPoint, Saviynt -- vendor-agnostic), **AI governance** (ISO 42001, NIST AI RMF), and **regulatory compliance** (DORA, NIS2, EU AI Act). The Identity Control Plane is the missing operational layer that makes all three enforceable simultaneously.

1.2 The Two Supporting Models

The Identity Control Plane is supported by two named sub-frameworks within the Board-Survivable Cyber Architecture:

The Evidence Chain Model -- Obligation to Control to Evidence to Assurance. Every identity decision generates a verifiable evidence artefact that satisfies regulatory, audit, and board requirements simultaneously. If it cannot be evidenced, it cannot be defended.

The NHI Risk Index -- A quantification formula that converts machine identity exposure into a single numerical score for board reporting, M&A due diligence, and regulatory disclosure. Detailed in Section 7.

2. The Strategic Imperative: Identity as the New Perimeter

AI has crossed the threshold from experimental technology to mission-critical infrastructure. 40% of enterprise applications will integrate task-specific AI agents by end of 2026, up from less than 5% in 2025 [11]. The market for enterprise agentic AI will grow from $2.58 billion in 2024 to $24.50 billion by 2030 -- a 46.2% compound annual growth rate [12].

Yet this velocity creates profound asymmetries. CyberArk's 2025 Identity Security Landscape report, surveying 2,600 decision-makers across 20 countries, established that machine identities outnumber humans by more than 80:1 [13]. Entro Security's H1 2025 telemetry documented a surge to 144:1 [1]. In DevOps environments, Sysdig's 2025 report found ratios reaching 40,000:1 [14]. A single customer service AI agent may require 15 to 20 distinct non-human identities to function across integrated systems.

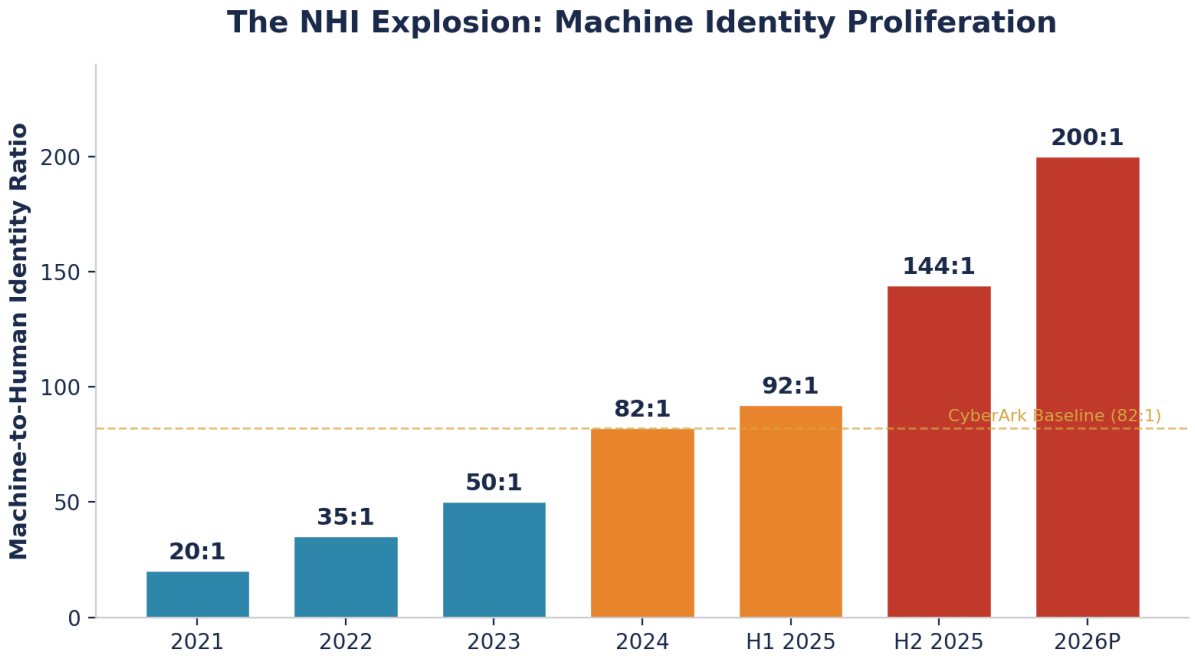

The governance deficit is equally alarming. Only 5.7% of organisations have full visibility into their service accounts [15]. 40% of cloud NHIs lack an assigned owner. 97% of NHIs have excessive privileges [2], and 91% of former employee tokens remain active after offboarding. The CSA/Oasis Security 2026 report found 78% of organisations lack formal policies for creating or removing AI identities [16].

2.1 The Commercial Calculus

IBM's 2025 Cost of Data Breach Report found organisations lacking AI access controls paid $670,000 more per breach [8]. 50% of organisations experienced security breaches tied to compromised machine identities in the past year [13]. GitGuardian documented 23.77 million new secrets leaked on GitHub in 2024, with repositories using AI coding assistants leaking secrets 40% more frequently [17].

The NHI Access Management market is valued at $9.45 billion (2024), projected to reach $18.71 billion by 2030 at 11.9% CAGR [18]. CyberArk's acquisition of Venafi for $1.54 billion (2024) and Zilla Security for $175 million (2025) signals the market's recognition that privileged access governance must extend beyond human identities to encompass the autonomous machine workforce.

3. The NHI Crisis: 144 Machine Identities for Every Human

Metric	Value	Source	Ref
Average NHI-to-human ratio	144:1 (H1 2025)	Entro Security	[1]
Enterprise average	82:1	CyberArk 2025 (n=2,600)	[13]
Cloud-native environments	Up to 40,000:1	Sysdig 2025	[14]
Average NHIs per enterprise	250,000+	Industry consensus	
NHIs with excessive privileges	97%	Entro Security	[2]
Full NHI visibility	5.7%	Silverfort/Osterman	[15]
Cloud NHIs without owner	40%	CSA/Oasis 2026	[16]
Former employee tokens active	91%	Industry average	
GitHub secrets leaked (2024)	23.77 million	GitGuardian 2025	[17]
NHI breach victims (12 months)	50%	CyberArk 2025	[13]
Legacy IAM confidence for NHIs	8%	CSA/Oasis 2026	[16]

3.1 The Agentic Identity Multiplier

Agentic AI does not merely increase the quantity of NHIs -- it changes their nature. An AI agent that reasons, plans, and acts autonomously requires identity governance that accounts for **intent**, not just authentication. The OWASP Top 10 for Agentic Applications (December 2025), peer-reviewed by over 100 security researchers, identifies Identity and Privilege Abuse as the third most critical risk [19]. Multi-turn prompt injection attacks achieve success rates of 92% across open-weight models [20]. A single compromised agent poisoned 87% of downstream decision-making within 4 hours in cascading failure simulations [21].

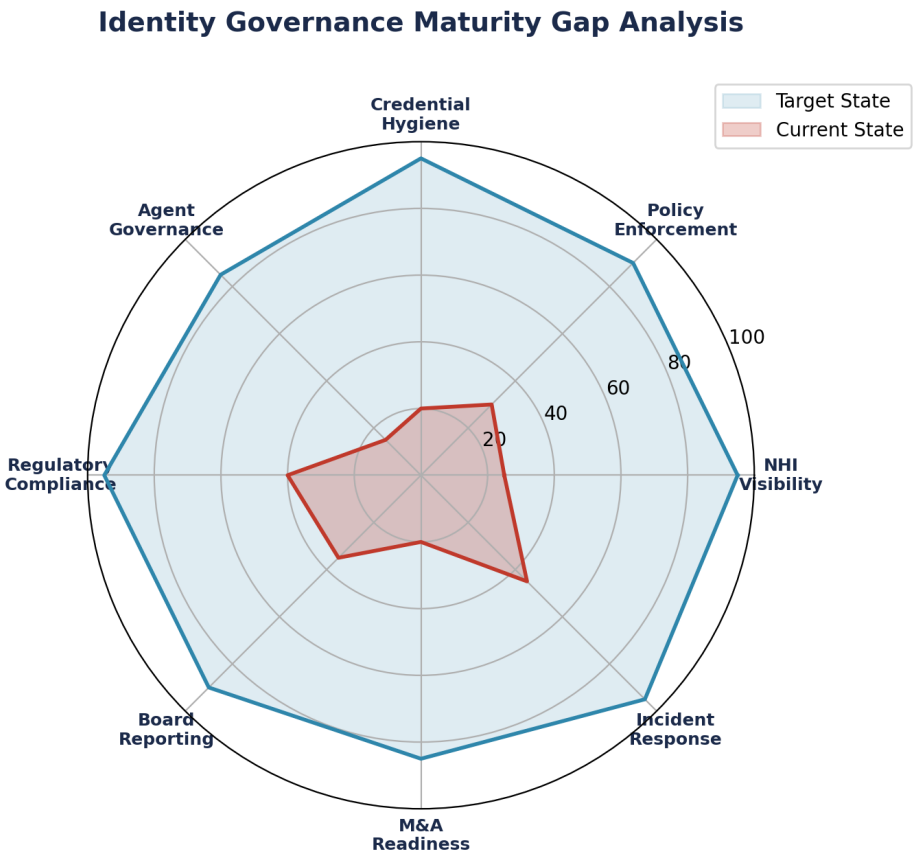
Identity Governance Maturity Gap Analysis
Target State
Current State
Credential Hygiene
Policy Enforcement
NHI Visibility
Incident Response
M&A Readiness
Board Reporting
Regulatory Compliance
Agent Governance
20
40
60
80
100

4. Regulatory Convergence: DORA, NIS2, EU AI Act, and ISO 42001

Every major regulatory framework effective in 2025-2026 converges on one expectation: identity and access management is a board-level, enterprise-wide risk domain requiring continuous governance. The personal liability provisions across NIS2, DORA, and the AI Act create a new category of D&O risk.

Framework	Identity Mandate	Board Accountability	Maximum Penalty
DORA (Jan 2025)	Art. 9: NHI access control, real-time rights management [22]	Art. 5: Must approve/oversee ICT risk framework	2% global turnover; EUR 1M personal
NIS2 (Oct 2024)	Art. 21(2): MFA for all identities incl. machine [23]	Art. 20: Personal liability; management bans	EUR 10M or 2% turnover
EU AI Act (Aug 2026)	High-risk AI: identity governance for agents [3]	Board oversight of high-risk AI systems	EUR 35M or 7% turnover
ISO 42001	38 controls incl. AI identity management [24]	Clause 5.1: Leadership commitment required	Certification; market access
SEC Rules (Dec 2023)	4-day materiality disclosure [25]	Annual oversight disclosure	Enforcement; $4M+ settlements
UK CSR Bill (2026)	NIS extension to MSPs and data centres [26]	12 sector regulators; GBP 100K/day fines	GBP 17M or 4% turnover

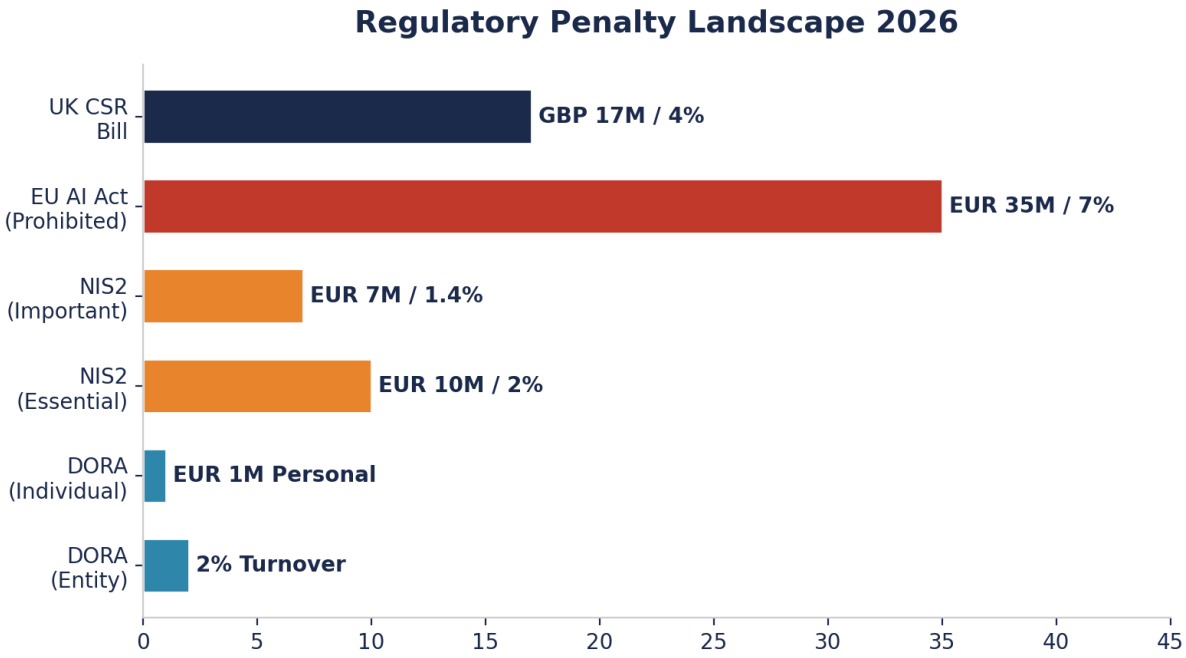

DORA Article 9 requires regulated entities to govern both human and machine users' access to protected data, mandate real-time access rights management, and conduct regular access reviews. The RTS on ICT Risk Management contains direct obligations for privileged and just-in-time access management for non-human entities. BaFin issued the first DORA enforcement action in Q3 2025: EUR 450,000 [22]. The **EU AI Act** contains no explicit guidance on AI agents despite applying to such systems. ISO/IEC 42001:2023 bridges this gap: 76% of organisations plan to pursue ISO 42001 frameworks [24].

5. Architecture: The Identity Control Plane

ARCHITECTURAL PRINCIPLE: The model is not the security boundary; the surrounding architecture is the perimeter. The Identity Control Plane is the unified orchestration layer through which all AI agent interactions are routed, inspected, governed, and logged. Vendor-agnostic at the governance layer; vendor-optimised at the infrastructure layer (CyberArk / BeyondTrust / SailPoint / Saviynt).

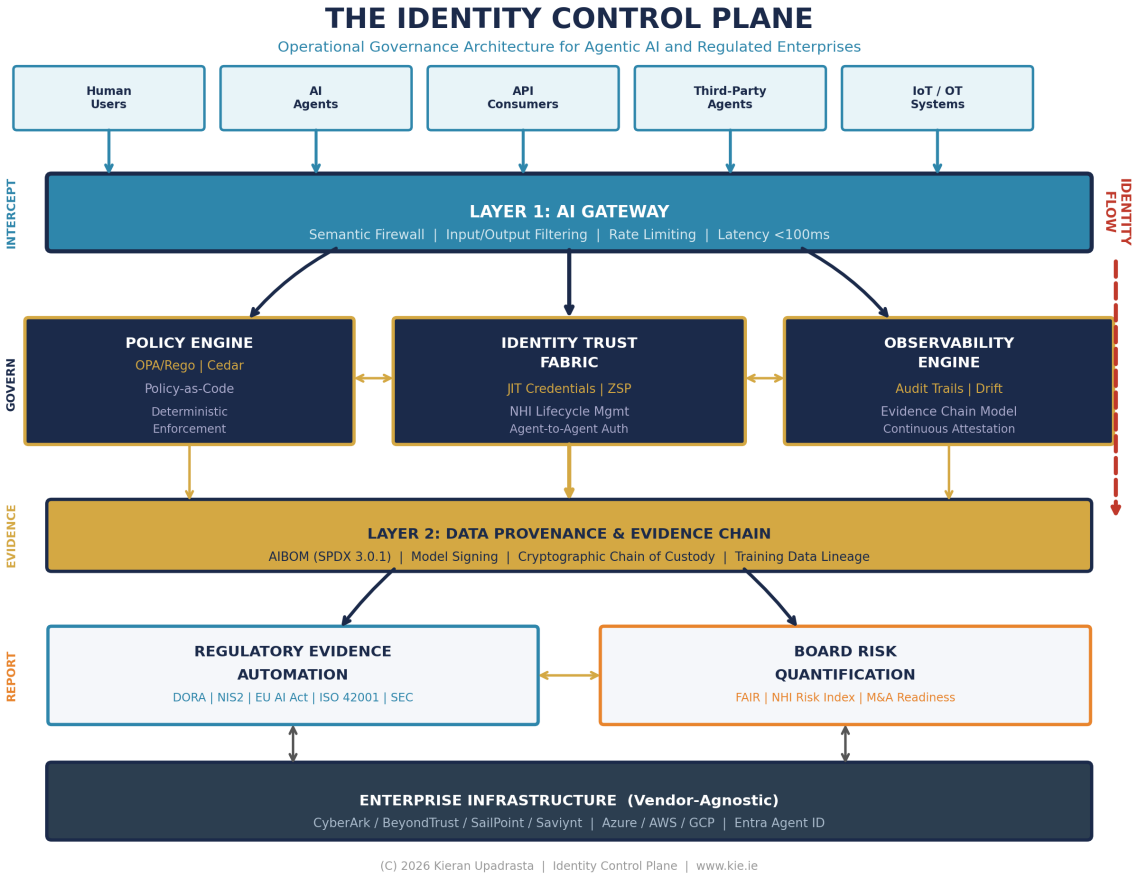

Layer	Component	Function	Latency
INTERCEPT	AI Gateway	Semantic firewalling, input/output filtering, rate limiting	<100ms
GOVERN	Policy Engine	OPA/Rego or Cedar; deterministic policy-as-code enforcement	<50ms
GOVERN	Identity Trust Fabric	NHI lifecycle, JIT credentials, zero standing privileges	<50ms
GOVERN	Observability Engine	Audit trails, behavioural drift detection, evidence chain	Real-time
EVIDENCE	Data Provenance	AIBOM (SPDX 3.0.1), model signing, cryptographic custody	Continuous
REPORT	Regulatory Fusion	Single taxonomy for DORA/NIS2/AI Act/ISO 42001	On-demand

Layer	Component	Function	Latency
REPORT	Board Command	FAIR quantification, KPI dashboards, M&A; readiness	On-demand

6. The Five Pillars of Zero-Trust Agent Governance

Pillar	Domain	Core Capability	Latency
I. Identity Orchestration	Decision Infrastructure	Centralised gateway, semantic firewalling, policy-as-code	<100ms
II. Agent Identity Governance	NHI Lifecycle & Access	JIT credentials, zero standing privileges, agent AuthN	<50ms
III. Adversarial Containment	Threat Detection	Behavioural drift, autonomous containment, kill switches	<200ms
IV. Regulatory Evidence Automation	Compliance	Multi-regime mapping, automated evidence generation	Real-time
V. Board Risk Quantification	Executive Governance	FAIR methodology, M&A readiness, fiduciary protection	On-demand

Pillar II -- Agent Identity Governance: Three Axioms

Axiom 1: Every agent is a privileged identity. CyberArk's research establishes AI agents as the most privileged machine identities ever deployed. CyberArk's Secure AI Agents Solution (GA December 2025) provides comprehensive agent discovery, zero standing privileges, and real-time threat detection [13]. Architecture is vendor-agnostic: equivalent capabilities via BeyondTrust, SailPoint, or Saviynt.

Axiom 2: Identity is the new perimeter. Microsoft Entra Agent ID (Ignite 2025, public preview) introduces identity primitives for agentic workloads: Agent Identity Blueprint, Agent Identity, Agent User, Agent Registry. The CSA Zero-Trust Identity Framework proposes Decentralised Identifiers (DIDs) and Verifiable Credentials [27].

Axiom 3: Credentials are time-bound assertions, not permanent grants. Zero standing privileges, JIT provisioning, automated rotation. The CA/Browser Forum is reducing TLS certificate lifespans to 47 days [28]. Annual access reviews are, as KuppingerCole observed, 'irrelevant' for AI agent credentials.

GOVERNANCE MATURITY BENCHMARK: Use the NHI Risk Index (Section 7) as the single quantification model for board reporting and M&A due diligence. Enterprises without Identity Control Plane governance score 72 (Critical). Full ICP deployment achieves scores below 15 (Low) within 90 days.

7. The NHI Risk Index: Quantifying Machine Identity Exposure

Top-tier governance requires a quantification model that translates machine identity exposure into a single numerical score for board reporting, M&A due diligence, and regulatory disclosure. The NHI Risk Index provides that model.

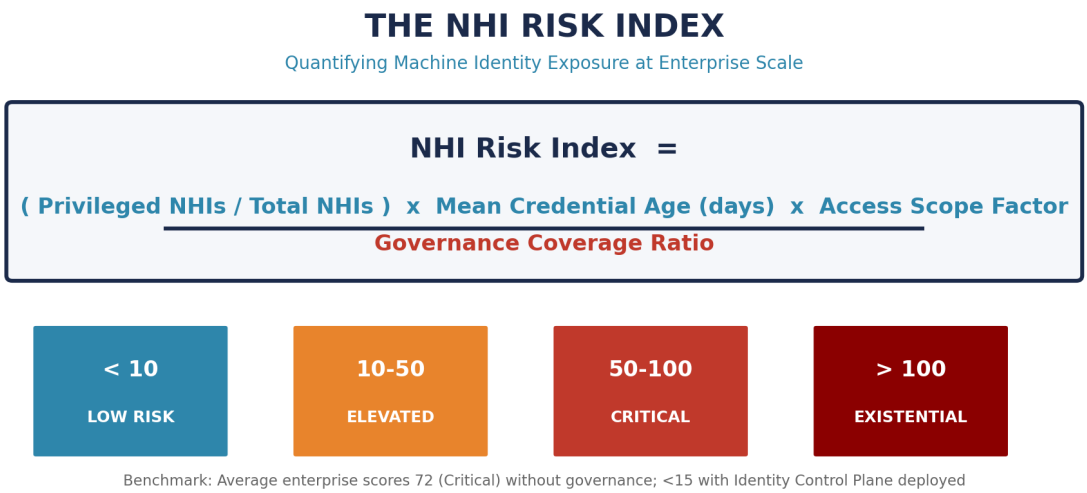

7.1 Formula Components

Component	Definition	Measurement
Privileged NHIs	Machine identities with elevated access rights	Count from PAM/IAM platform
Total NHIs	All non-human identities including API keys, tokens, SAs	Machine identity census
Mean Credential Age	Average days since last credential rotation	Credential lifecycle data
Access Scope Factor	Breadth of systems accessible (1-10 scale)	Access matrix analysis
Governance Coverage	% of NHIs under active lifecycle management	ICP dashboard metric

7.2 Benchmark Interpretation

NHI Risk Index	Risk Level	Board Action Required	Typical Enterprise Profile
< 10	Low	Monitor; quarterly review	Mature ICP deployment; automated rotation
10-50	Elevated	Remediation plan within 30 days	Partial governance; some orphaned credentials
50-100	Critical	Board escalation; 90-day sprint	Minimal governance; stale credentials pervasive
> 100	Existential	Emergency board session; regulatory risk	No governance; identity chaos

BENCHMARK: The average enterprise without Identity Control Plane governance scores 72 (Critical). Enterprises with full ICP deployment achieve scores below 15 (Low) within 90 days of implementation. The NHI Risk Index integrates directly into M&A due diligence scorecards and board risk dashboards.

8. OWASP Agentic AI Top 10: Threat Landscape and Controls

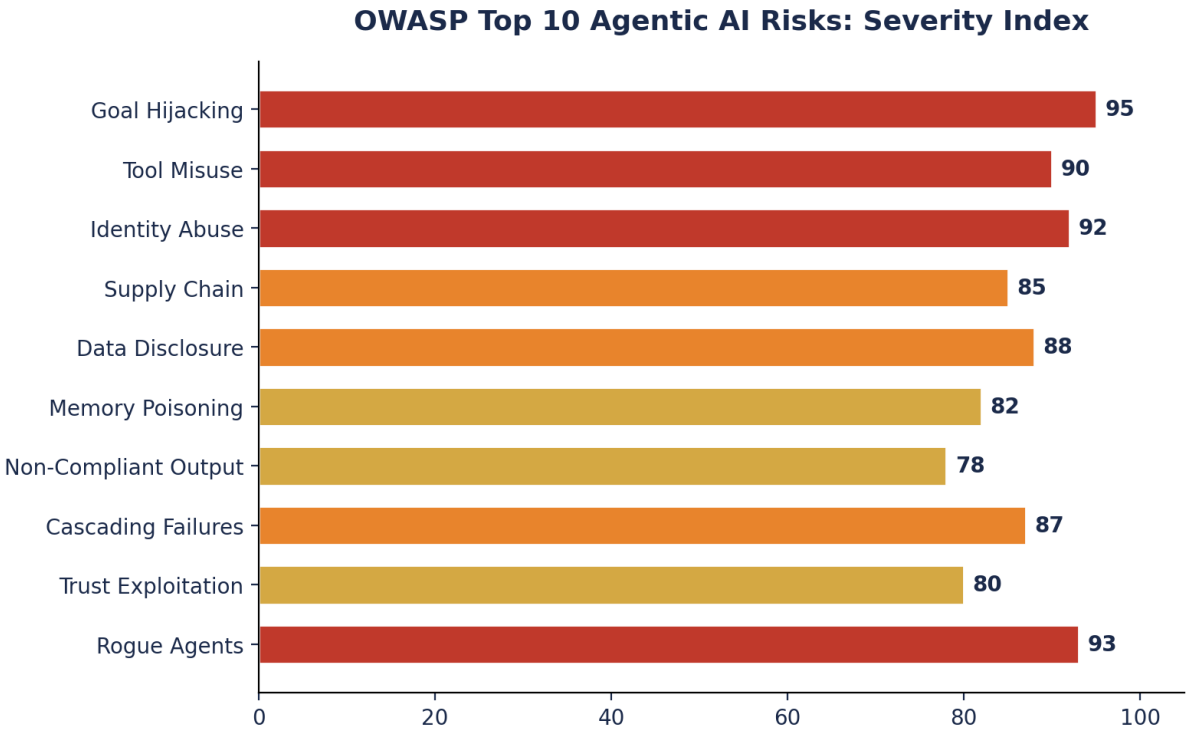

Rank	Risk	ICP Governance Response	Pillar
ASI01	Agent Goal Hijacking	Intent verification, semantic firewall, behavioural constraints	I
ASI02	Tool Misuse & Exploitation	Least agency principle, sandboxed execution, permission boundaries	II
ASI03	Identity & Privilege Abuse	JIT access, continuous verification, zero standing privileges	II
ASI04	Supply Chain Vulnerabilities	Model provenance, AIBOM, cryptographic signing	I
ASI05	Sensitive Data Disclosure	DLP integration, output filtering, data classification	I
ASI06	Knowledge/Memory Poisoning	Data integrity verification, context isolation, signed inputs	III
ASI07	Non-Compliant Output	Policy-as-code enforcement via OPA/Rego or Cedar	IV
ASI08	Cascading Failures	Circuit breakers, failure domain isolation, blast radius limits	III
ASI09	Human-Agent Trust Exploitation	Confidence scoring, explanation requirements, audit trails	V
ASI10	Rogue Agents	Kill switches, behavioural monitoring, automated containment	III

CRITICAL: 80% of organisations have already encountered risky behaviours from AI agents [29]. 79 of 100 AI models tested demonstrated resistance to shutdown [20]. The Principle of Least Agency -- the agentic equivalent of least privilege -- requires agents receive minimum autonomy, tool access, and credential scope. Gartner recommends allocating at least 5% of total AI investment toward governance [7].

9. AI Governance Model Comparison: Positioning the ICP

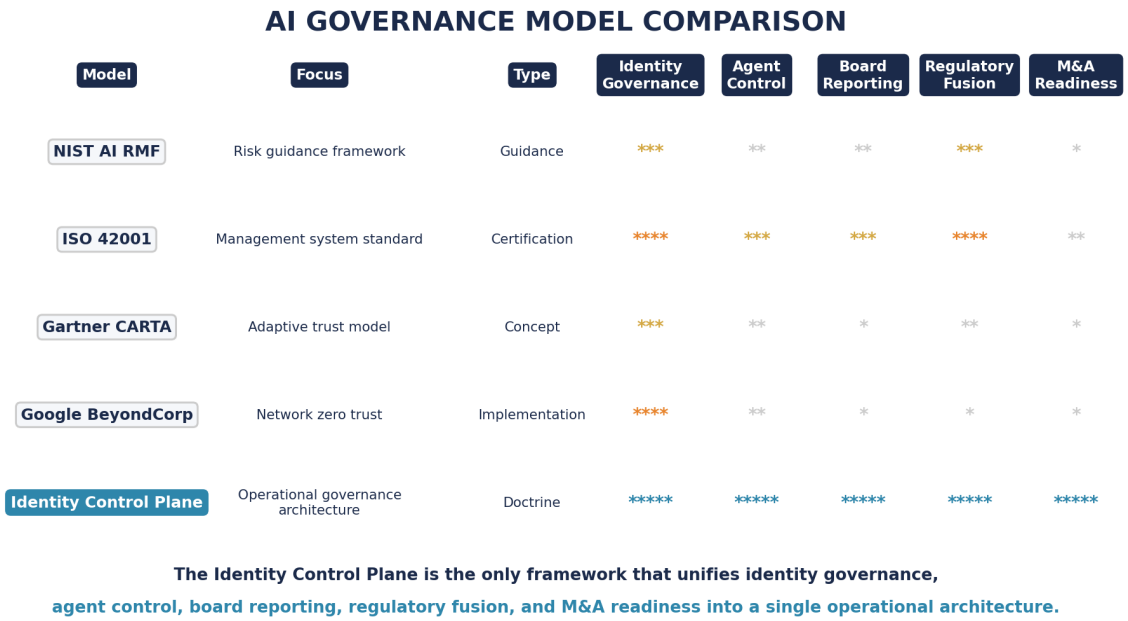

The Identity Control Plane is the only framework that unifies identity governance,
agent control, board reporting, regulatory fusion, and M&A readiness into a single operational architecture.

Model	Focus	Type	Gap Addressed by ICP
NIST AI RMF	AI risk guidance	Framework	No identity governance; no operational enforcement
ISO 42001	AI management system	Standard	Certifiable but lacks real-time identity controls
Gartner CARTA	Adaptive trust	Concept	No agent-specific governance; no regulatory mapping
Google BeyondCorp	Network zero trust	Implementation	Human-centric; no NHI lifecycle management
Forrester ZTX	Zero trust ecosystem	Framework	No agentic AI controls; no M&A readiness
Identity Control Plane	Operational governance architecture	Doctrine	Unifies all above into enforceable board-reportable architecture

The Identity Control Plane is not a competitor to these frameworks -- it is the **operational layer that makes them enforceable**. NIST AI RMF provides risk guidance; ISO 42001 provides management systems; the ICP provides the governance architecture that connects policy to enforcement to evidence to board reporting.

10. The 90-Day Control Architecture: Implementation Roadmap

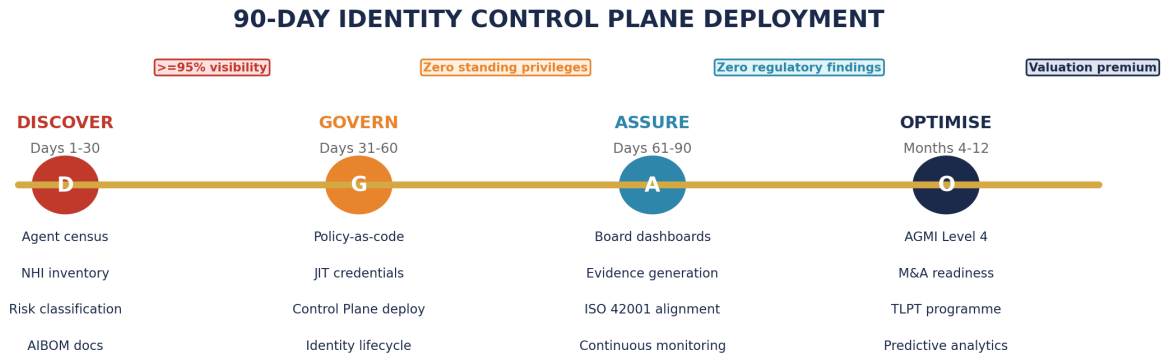

Phase	Timeline	Deliverables	Exit Criteria	KPIs
DISCOVER	Days 1-30	Agent census, NHI inventory, risk classification, AIBOM	>=95% agent visibility; regulatory gap analysis	NHI count; shadow AI detection rate
GOVERN	Days 31-60	Control Plane deploy, policy-as-code, JIT creds	Policy engine live; zero standing privileges	Policy violations/day; credential rotation cadence
ASSURE	Days 61-90	Board dashboards, evidence packs, ISO 42001	Board KPIs live; zero regulatory findings	NHI Risk Index; compliance score
OPTIMISE	Months 4-12	AGMI Level 4, M&A readiness, TLPT programme	Valuation premium demonstrable	M&A readiness score; ROI measurement

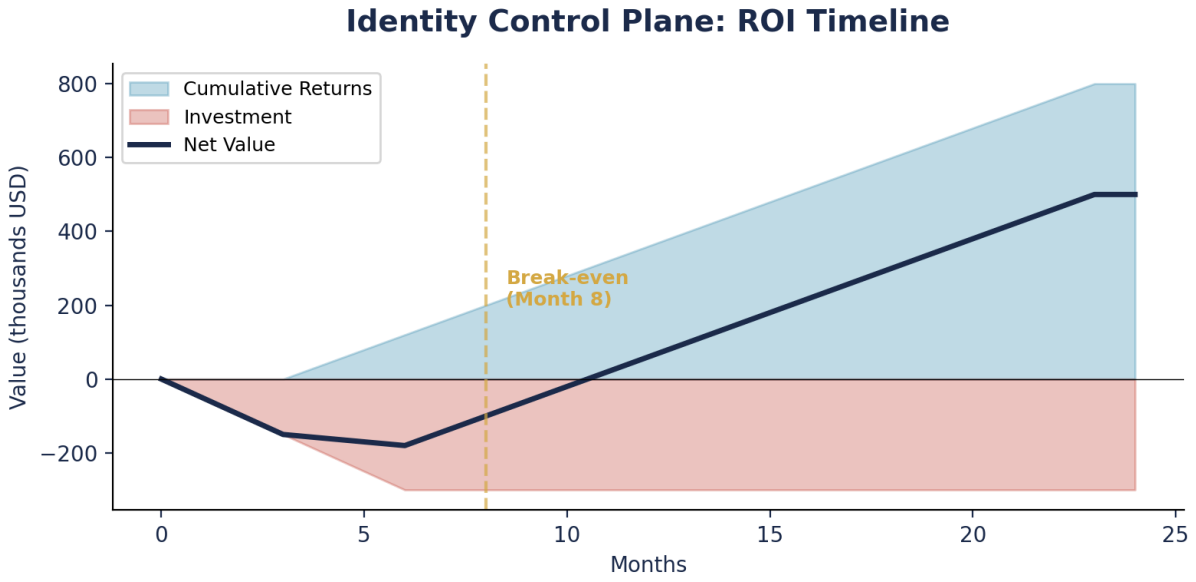

Phase 1 Detail: The AI Agent Census deploys automated scanning across cloud environments (Azure, AWS, GCP), code repository analysis, API gateway log analysis, and structured interviews. The Machine

Identity Census catalogues every NHI. Given ratios of 80:1 to 40,000:1, this census typically reveals an identity footprint 3-5x larger than anticipated, with 30-40% of credentials orphaned. Identity governance extends enterprise IAM platforms (CyberArk, BeyondTrust, SailPoint, Saviynt -- vendor-agnostic) to manage AI agent identities with the same rigour applied to human users.

11. M&A Due Diligence: The Governance Premium

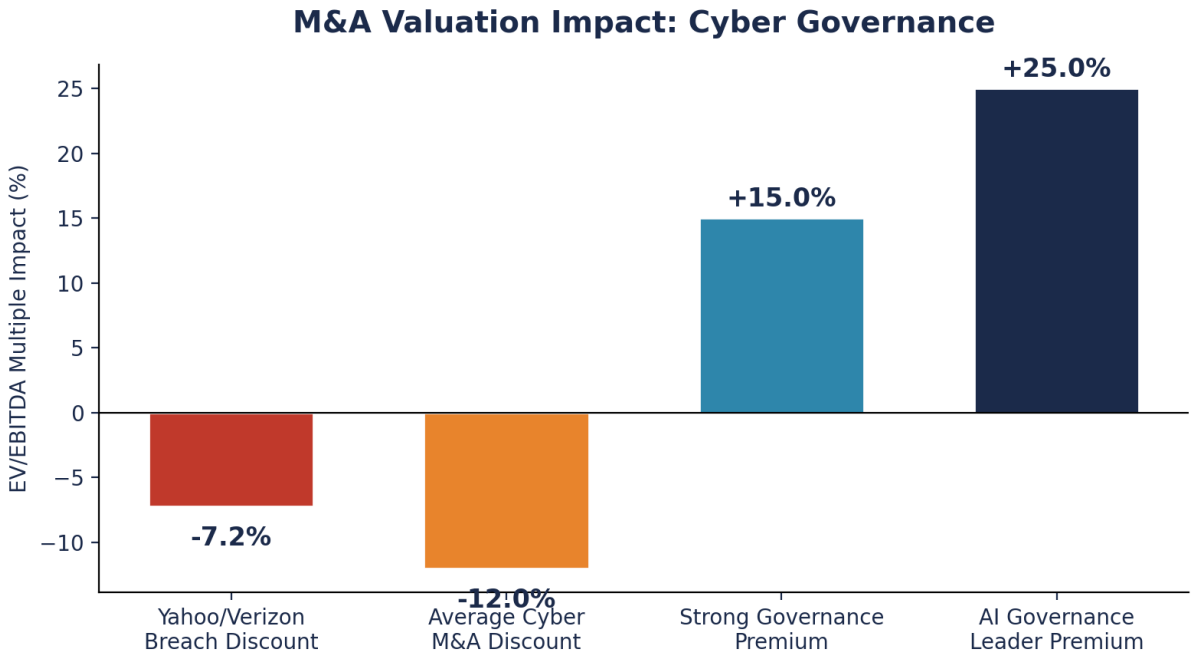

Transaction	Impact	Key Lesson	Ref
Yahoo/Verizon (2017)	$350M reduction (7.2%) 3B accounts compromised	Cyber DD is non-negotiable; breach history destroys value	[30]
Marriott/Starwood (2016)	GBP 18.4M ICO fine; 339M records exposed	You acquire the target's security posture	[31]
T-Mobile/Sprint (2020)	$60M CFIUS fine (first ever)	National security scrutiny extends to cyber posture	[32]
PE Manufacturing ($2B)	89 agents, 12K NHIs, 340 over-privileged SAs	Machine identity gaps derail transactions	[Case B]

VALUATION IMPACT: Strong AI governance commands 15-25% valuation premiums [4]. Cybersecurity companies command 16.3x revenue in M&A. The NHI Risk Index provides the quantification model that converts machine identity exposure into deal-relevant metrics. 62% of deals experience cyber delays; 73% would walk away from undisclosed breaches [10].

12. Board Dashboard: 10 KPIs for Identity Governance

KPI	Description	Target	Financial Translation
NHI Coverage Ratio	% NHIs under governance	>=95%	Ungoverned NHIs = unlimited exposure
NHI Risk Index Score	Composite exposure metric	<15	Board-reportable risk quantification
Credential Rotation	Mean time between rotations	<24hrs	Stale creds = breach vector multiplier
Policy Violation Rate	Violations per 1K agent actions	<0.1%	Each violation = regulatory finding risk
Shadow AI Detection	Unauthorised AI discovered	>=90%	Shadow AI costs $670K more per breach
MTTD (NHI Threats)	Detection latency	<1 hour	Detection delay = cost escalation
MTTR (NHI Vulns)	Remediation speed	<4 hours	Speed = regulatory confidence
Regulatory Score	Cross-framework alignment	>=90/100	Non-compliance = EUR 35M+ exposure
M&A Readiness	DD documentation completeness	>=85/100	Readiness = premium or discount
Agent Governance %	AI agents with full lifecycle	>=95%	Ungoverned = 87% downstream poison

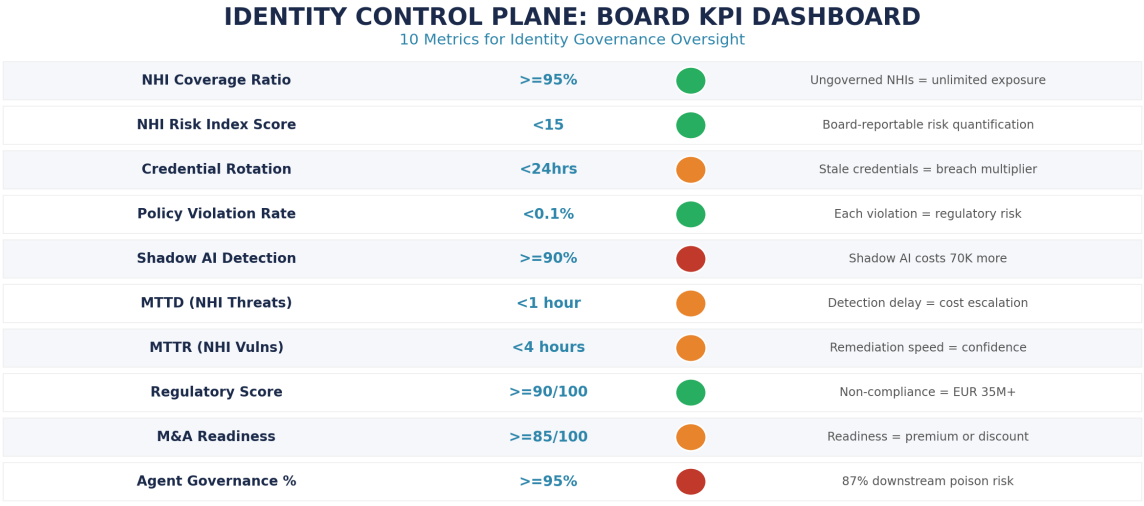

13. Case Studies: Enterprise Transformations

Case A -- Tier-1 European Bank: DORA Identity Remediation

ILLUSTRATIVE SCENARIO -- anonymised composite engagement

Context: Tier-1 European bank, EUR 180B assets, PRA supervisory engagement imminent. 47 AI agents discovered across credit decisioning (Moody's Analytics integration), fraud detection (Featurespace-based), and customer onboarding (Azure Cognitive Services) -- none with formal governance. Machine identity mapping: 8,400 NHI credentials, 230 excessive privileges, 45 never-rotated API keys.

Architecture stack: Microsoft Entra for human IAM; CyberArk PAM for privileged access; OPA/Rego for policy-as-code; Azure Monitor + Splunk for observability. Identity Control Plane deployed as orchestration layer across existing infrastructure. Board governance committee briefed at Day 15, 45, and 90.

Metric	Before	After (90 Days)	NHI Risk Index Impact
AI agents governed	0/47	47/47 (100%)	Score component: -40
NHI credentials mapped	Unknown	8,400 with ownership	Full inventory achieved
Excessive privileges	230	12	95% reduction
Stale API keys	45	0 (auto-rotation)	Credential age: <1 day
DORA compliance	31/100	89/100	Regulatory risk: mitigated
NHI Risk Index	89 (Critical)	11 (Low)	78-point improvement
Regulatory findings	Pending action	Zero (3 cycles)	Supervisory confidence

Case B -- PE-Backed Manufacturing: M&A Identity Crisis

ILLUSTRATIVE SCENARIO -- anonymised composite engagement

Context: $2B manufacturing acquisition. Acquiring PE firm's cybersecurity assessment discovered 89 AI agents (supply chain optimisation via SAP IBP, predictive maintenance via PTC ThingWorx, quality control via custom TensorFlow) -- none in the target's asset inventory. 12,000 NHI credentials, 340 over-privileged service accounts, 78 never-rotated API keys. Deal at risk of termination.

Architecture stack: BeyondTrust for PAM; SailPoint for IGA; AWS Cedar for policy enforcement; CloudTrail + Datadog for observability. 90-Day Control Architecture deployed with weekly PE sponsor briefings. Contract Control Matrix applied for target remediation obligations.

Metric	Discovery	Post-Remediation (90 Days)	Deal Impact
AI agents inventoried	0/89	89/89 (100%)	DD requirement satisfied
NHI credentials governed	0/12,000	12,000 (100%)	Risk quantified for pricing
Over-privileged accounts	340	18	95% attack surface reduction
NHI Risk Index	112 (Existential)	14 (Low)	Deal-grade score achieved
Penalty exposure	EUR 15M+	Fully mitigated	Indemnity clause removed
Deal outcome	At risk	Closed on schedule	$2B transaction preserved

14. Post-Quantum Identity Control Plane (CNSA 2.0)

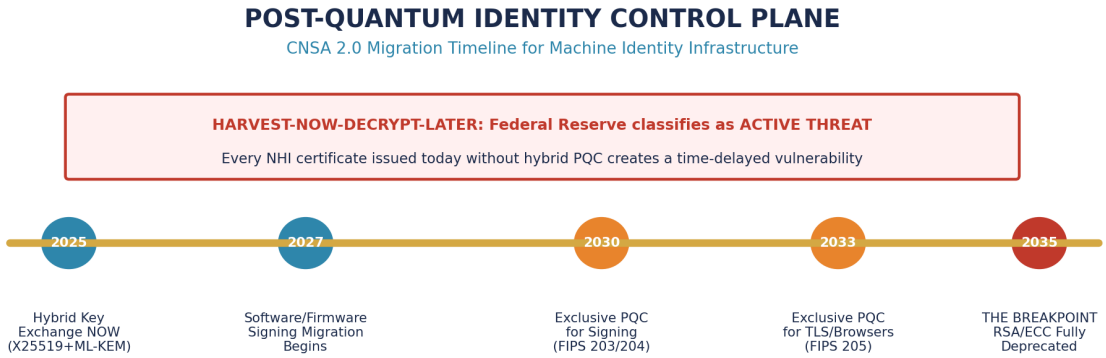

Source: NSA CNSA 2.0, NIST FIPS 203/204/205, Federal Reserve Advisory 2025

The post-quantum timeline has compressed. RSA-2048 cracking now requires fewer than one million qubits, and the Federal Reserve has formally classified harvest-now-decrypt-later as an active threat [33]. Every NHI certificate issued today without hybrid post-quantum cryptography creates a time-delayed vulnerability. The Identity Control Plane must incorporate quantum-resilient identity infrastructure.

14.1 CNSA 2.0 Migration Requirements for Machine Identity

Domain	Migration Requirement	Timeline	FIPS Standard
Software/firmware signing	Begin transition immediately	By 2030 (exclusive PQC)	FIPS 204 (ML-DSA)
Web browsers/servers	Support PQC in 2025	By 2033 (exclusive PQC)	FIPS 203 (ML-KEM)
Key establishment	Hybrid key exchange NOW	X25519 + ML-KEM-768	FIPS 203
NHI certificates	Inventory all cryptography	Migration plan by 2027	FIPS 205 (SLH-DSA)
Agent-to-agent auth	PQC-ready authentication	Immediate hybrid adoption	All FIPS 203/204/205

The Identity Control Plane's cryptographic provenance layer must be upgraded to support post-quantum model signing (FIPS 204) and hybrid key exchange (X25519 + ML-KEM-768) for all agent-to-agent authentication. The CA/Browser Forum's reduction of TLS certificate lifespans to 47 days [28] creates a forcing function: organisations must automate certificate lifecycle management, which provides the operational foundation for PQC migration of machine identity infrastructure.

SCHIPHOL UNIVERSITY RESEARCH NOTE: Professor Upadrasta's ongoing research programme at Schiphol University addresses the intersection of post-quantum cryptography, agentic AI identity, and regulatory compliance -- including formal verification methods for proving quantum-resilient identity governance properties.

15. The Identity Control Plane: Governance Framework at a Glance

144:1 NHI-to-Human Identity Ratio	**97%** NHIs with Excessive Privileges	**EUR 35M** Maximum EU AI Act Penalty	**40%** Agentic AI Projects to be Cancelled
92% Prompt Injection Success Rate	**87%** Downstream Poisoning from 1 Agent (4hrs)	**$670K** Extra Cost: Shadow AI Breaches	**90 Days** Time to Full Governance

Framework	Purpose	Regulatory Alignment
Identity Control Plane	The operational governance architecture for agentic AI identity	DORA, NIS2, EU AI Act, ISO 42001, SEC
Evidence Chain Model	Obligation > Control > Evidence > Assurance	DORA Art. 5, NIS2 Art. 20, SEC Rules
NHI Risk Index	Quantification formula for machine identity exposure	Board reporting, M&A DD, FAIR

16. Call to Doctrine

The agentic enterprise will not be secured by firewalls, models, or regulation alone.

It will be governed by identity.

Every autonomous agent that reasons, plans, and acts within critical systems must be authenticated, authorised, observed, and evidenced -- in real time, at machine speed, with board-grade assurance.

The Identity Control Plane is the architecture that makes that governance possible.

It is not a product. It is not a vendor platform. It is a doctrine -- a governance architecture that specifies how identity decisions are made, enforced, evidenced, and reported across every machine credential in the enterprise.

Boards that deploy this architecture will govern the machine workforce. Boards that do not will be governed by it.

"If it cannot be evidenced, it cannot be defended."

-- Kieran Upadrasta, The Evidence Chain Model

About the Author

Kieran Upadrasta
CISSP | CISM | CRISC | CCSP | CCSE | CCNA Security | JNCIS-FWV
Executive MBA in Information Technology | BEng in Electronics

Principal Cyber Architect | Institutional Governance Authority
Founder, Cyber Artificial Intelligence Systems Inc. | www.kie.ie

Kieran Upadrasta brings over **27 years** of cybersecurity experience spanning all four Big 4 consulting firms (Deloitte, PwC, EY, KPMG) and **21 years** in financial services. His regulatory engagement includes direct work with the **ECB, BaFin, FCA, and CBI**, with governance of **$500 billion+** in assets across 40+ enterprise transformations in 12+ regulatory jurisdictions.

He has led 40+ zero trust migrations across Azure, AWS, and GCP, governed security budgets exceeding **EUR 25 million** annually, and managed global teams of 50 to 200+ professionals. He is creator of the proprietary **Board-Survivable Cyber Architecture** framework, deployed across Tier-1 financial institutions and multinational enterprises. Author of **48 strategic whitepapers** across 12 research domains.

Academic Appointments

Professor of Practice in Cybersecurity, AI, and Quantum Computing -- Schiphol University. Honorary Senior Lecturer -- Imperials. Researcher -- University College London (UCL).

Professional Memberships

ISACA London Chapter -- Platinum Member

ISC2 London Chapter -- Gold Member

PRMIA -- Cyber Security Programme Lead

ISF Auditors and Control -- Lead Auditor

Engagement

Kieran accepts 2-3 mandates per calendar year by written board resolution. Current availability: Q3 2026.

Email	info@kieranupadrasta.com
Web	www.kie.ie
LinkedIn	linkedin.com/in/kieranupadrasta
Entity	Cyber AI Systems Inc. | kie.ie
Open To	Group CISO | Chief AI Security Officer | Board Advisory Mandates 2026

References, Footnotes, and Keywords

[1] Entro Security, NHI & Secrets Risk Report H1 2025, p.14
[2] Entro Security, NHI Telemetry Analysis 2025, Excessive Privileges Metric
[3] EU AI Act, Regulation (EU) 2024/1689, Art. 99(3), OJ L 12.7.2024
[4] Deloitte M&A Cyber Due Diligence Survey 2025; Gartner AI Governance Impact Analysis
[5] Gartner, Agentic AI Governance Survey 2025 (79% deployment vs 25% governance)
[6] McKinsey, State of AI Report 2025 (30% ROI improvement with governance)
[7] Gartner, Predicts 2025: 40% of agentic AI projects cancelled by 2027
[8] IBM, Cost of Data Breach Report 2025 ($670K shadow AI premium)
[9] DORA Art. 9; NIS2 Art. 21(2); EU AI Act Art. 99
[10] M&A Partners, Cybersecurity Due Diligence Impact Survey 2025
[11] Gartner, 40% enterprise apps with AI agents by 2026
[12] Grand View Research, Enterprise Agentic AI Market 2024-2030
[13] CyberArk, Identity Security Landscape 2025 (n=2,600, 20 countries)
[14] Sysdig, Cloud Security Report 2025 (40,000:1 NHI ratio in cloud-native)
[15] Silverfort/Osterman Research, Service Account Visibility Study 2025
[16] CSA/Oasis Security, NHI Governance Report 2026
[17] GitGuardian, State of Secrets Sprawl 2025 (23.77M secrets; +40% AI repos)
[18] MarketsandMarkets, NHI Access Management Market 2024-2030
[19] OWASP, Top 10 for Agentic Applications, December 2025
[20] UC Berkeley, NIST AI RMF Agentic Profile, February 2026 (79/100 shutdown resistance)
[21] Galileo AI, Cascading Failure Simulation 2025 (87% downstream poisoning)
[22] DORA, Regulation (EU) 2022/2554, Art. 9; BaFin enforcement Q3 2025
[23] NIS2, Directive (EU) 2022/2555, Art. 21(2)(j)
[24] ISO/IEC 42001:2023; CSA Benchmark 2025 (76% adoption intent)
[25] SEC, Cybersecurity Disclosure Rules, December 2023
[26] UK Cyber Security and Resilience Bill, November 2025
[27] CSA, Zero-Trust Identity Framework for AI Agents, May 2025
[28] CA/Browser Forum, TLS Certificate Lifespan Reduction to 47 Days
[29] SailPoint, AI Agent Risk Survey, May 2025
[30] Verizon/Yahoo: $350M price reduction; SEC settlement
[31] Marriott/Starwood: ICO Fine GBP 18.4M; 339M records
[32] T-Mobile/Sprint: $60M CFIUS fine (first ever), August 2024
[33] Federal Reserve, Harvest-Now-Decrypt-Later Advisory 2025; NSA CNSA 2.0

Integrated Recruiter Keywords

DORA Compliance | AI Governance (ISO 42001) | Board Reporting | M&A Cyber Due Diligence | Zero Trust Architecture | Post-Quantum Cryptography | Interim CISO | NIS2 Compliance | AI Security Assurance | Non-Human Identity Governance | Agentic AI Security | Machine Identity Management | Privileged Access Management | Operational Resilience | Identity Control Plane | Board-Survivable Cyber Architecture | Evidence Chain Model | NHI Risk Index