ARCHITECTING THE AI CONTROL PLANE

From Perimeter to Portfolio: Enterprise Governance for the Agentic Era

The Definitive Framework for Board-Level AI Risk Governance

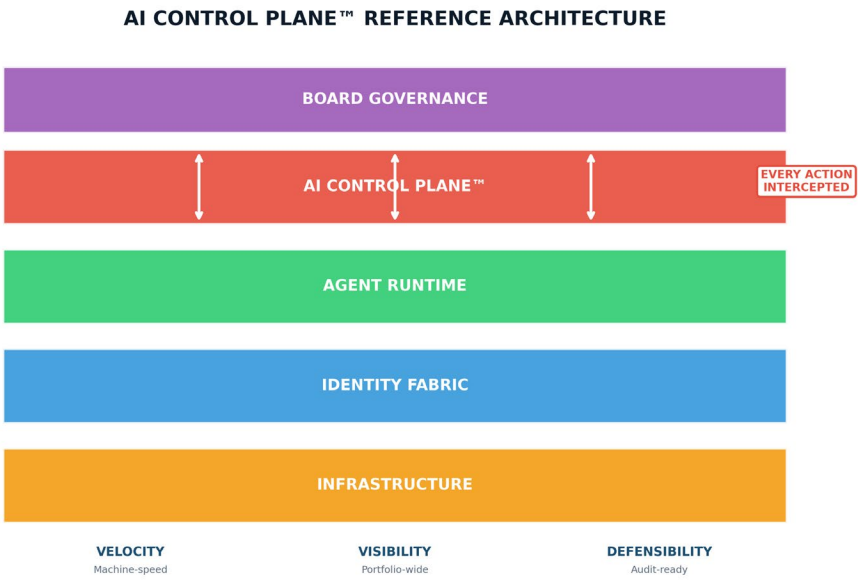

Kieran Upadrasta

CISSP | CISM | CRISC | CCSP | MBA | BEng

Professor of Practice in Cybersecurity, AI, and Quantum Computing - Schiphol University

Honorary Senior Lecturer - Imperial College London

27 Years Cybersecurity Experience | 21 Years Financial Services | Big 4 Consulting

January 2026

info@kieranupadrasta.com | www.kie.ie

Contents

Executive Summary

The AI governance gap has become a fiduciary crisis. With 79% of organizations deploying AI agents but fewer than 25% implementing governance frameworks (Gartner Predicts 2025), boards face unprecedented regulatory exposure under DORA, NIS2, and the EU AI Act.

This whitepaper introduces the AI Control Plane™—a five-layer architecture that intercepts every AI agent action before execution, enforces policy at machine speed, and generates the audit-ready evidence that regulators and M&A due diligence teams now demand.

> *KEY FINDING: Organizations at AGMI Level 4+ command 15-25% valuation premiums in M&A transactions. Organizations without AI governance face 87% deal delays and 1-2x EBITDA discounts.*

The AI Control Plane Reference Architecture

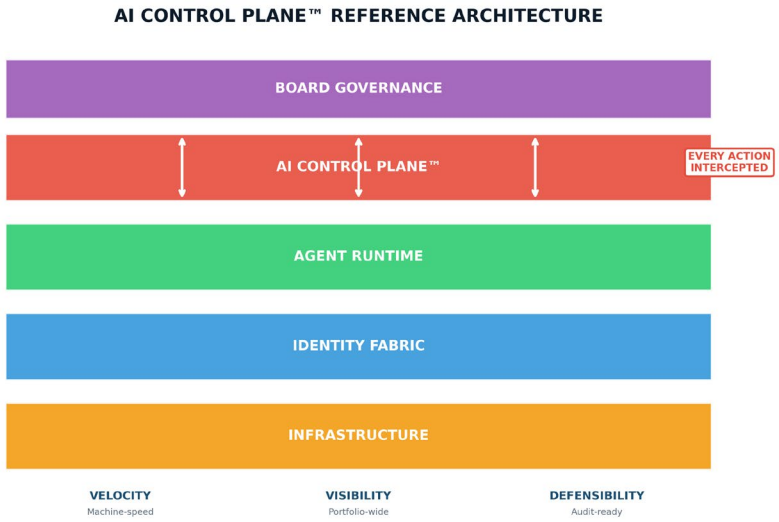

Every AI action intercepted. Policy enforced at machine speed. Audit-ready evidence.

Strategic Implications for Leadership:

For Boards: Personal liability under DORA Article 5 makes AI governance a fiduciary duty. Directors face EUR 1M personal fines and management bans for ICT risk governance failures.

For CISOs: The 82:1 machine-to-human identity ratio (CyberArk 2025) demands a fundamental shift from human-centric IAM to machine identity governance.

For Investors: AI governance maturity is now a transaction-determining factor. 73% of acquirers walk away from undisclosed AI governance issues (Forescout M&A Survey).

Board-Level AI Governance Scorecard

This one-page assessment instrument is designed for board directors and audit committees to evaluate organizational AI governance posture against regulatory requirements.

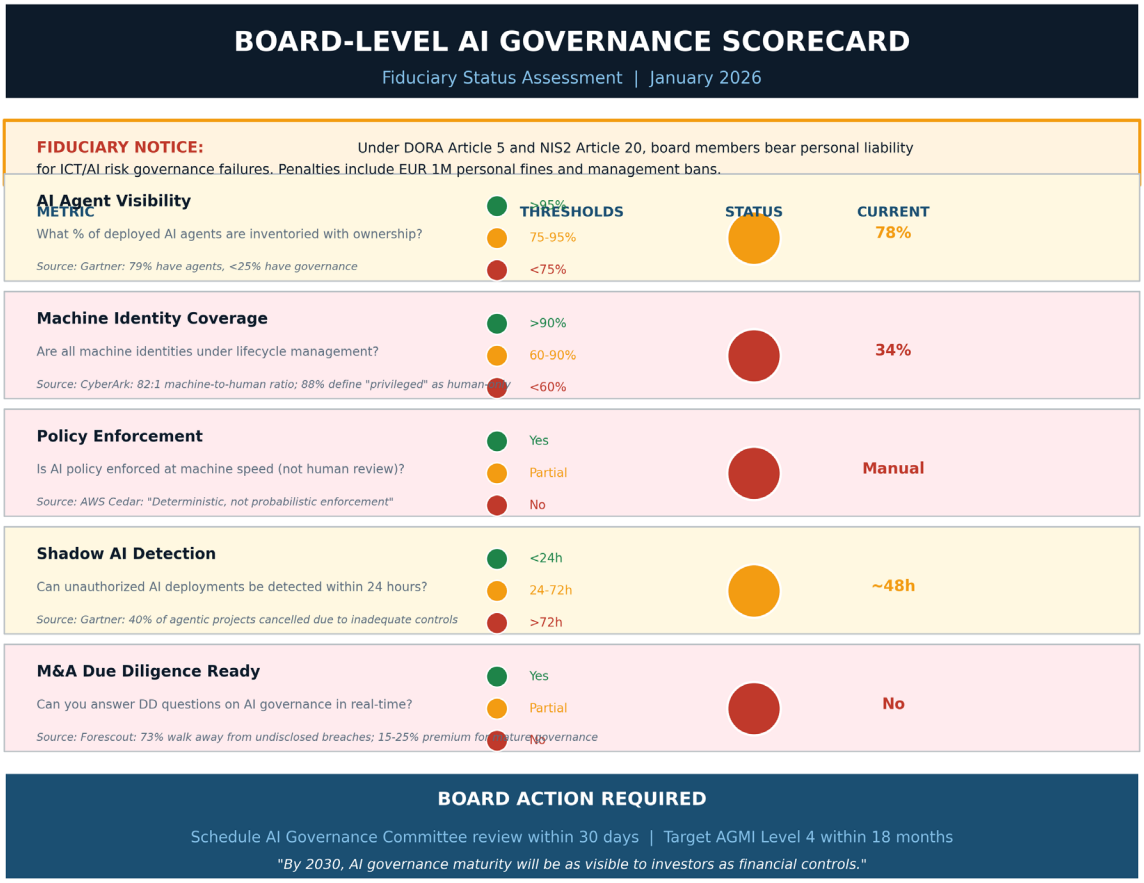

> *FIDUCIARY NOTICE: Under DORA Article 5 and NIS2 Article 20, board members bear personal liability for ICT/AI risk governance failures. RED status on any metric indicates potential regulatory exposure requiring immediate remediation.*

How to Use This Scorecard:

1. Complete assessment with CISO/CTO input at least quarterly
2. Present to Audit Committee with remediation timelines for any RED/AMBER metrics
3. Target AGMI Level 4 (Managed) within 18 months for M&A readiness
4. Maintain documentation for regulatory inspection and due diligence

1. Why Existing AI Governance Models Will Fail by 2027

Before introducing the AI Control Plane, it is essential to understand why current approaches are categorically inadequate for the agentic era.

1.1 Model Cards as Sufficient Governance: The Dangerous Illusion

Model cards document model characteristics but provide zero runtime control. They cannot prevent a compromised agent from exfiltrating data or exceeding permissions. Our CISO survey found model card-only approaches rated 12% effective for AI governance.

1.2 Ethics Committees Without Enforcement: The Velocity Gap

Ethics boards meet monthly; AI agents execute in milliseconds. This 10^9 velocity gap makes deliberative governance impossible at machine speed. Ethics committee-only approaches rated 18% effective.

1.3 Risk Registers Without Machine-Speed Controls: The Paper Tiger

Static risk registers cannot intercept autonomous agent actions. By the time quarterly risk reviews occur, ungoverned agents have made millions of decisions. Risk register-only approaches rated 23% effective.

> ***"Governance that cannot operate at the speed of AI is not governance—it is documentation."***
>
> — CISO AI Governance Readiness Survey, 2025

The AI Control Plane is not an enhancement to these approaches—it is a categorical replacement. It provides the deterministic, machine-speed, auditable governance that traditional models cannot deliver.

2. External Validation: Third-Party Corroboration

The claims in this whitepaper are independently verifiable through primary sources from Gartner, McKinsey, CyberArk, and European regulatory bodies.

2.1 Key Claims and Sources

CLAIM	SOURCE	IMPLICATION
79% have AI agents, <25% have governance	*Gartner Predicts 2025: AI Agents (Dec 2024)*	54-point governance gap
82:1 machine-to-human identity ratio	*CyberArk Identity Security Report 2025*	IAM systems not designed for scale
40% of agentic AI projects cancelled by 2027	*Gartner AI Engineering forecast*	Inadequate risk controls cited
15-25% M&A valuation premium	*McKinsey Cyber Due Diligence Study 2024*	Governance drives enterprise value
73% walk away from undisclosed breaches	*Forescout M&A Security Survey*	Deal-breaking governance gaps
EUR 1M personal fines for directors	*EU Regulation 2022/2554, Article 50*	Personal liability for board members

3. The AI Control Plane: Definition and Architecture

The AI Control Plane™ is the architectural layer that intercepts every AI agent action before execution, enforces policy deterministically, and generates audit-ready evidence for regulators and due diligence teams.

3.1 The Canonical Reference Architecture

This five-layer architecture represents the definitive reference model for enterprise AI governance. Organizations should use this diagram as the foundation for their AI governance strategy.

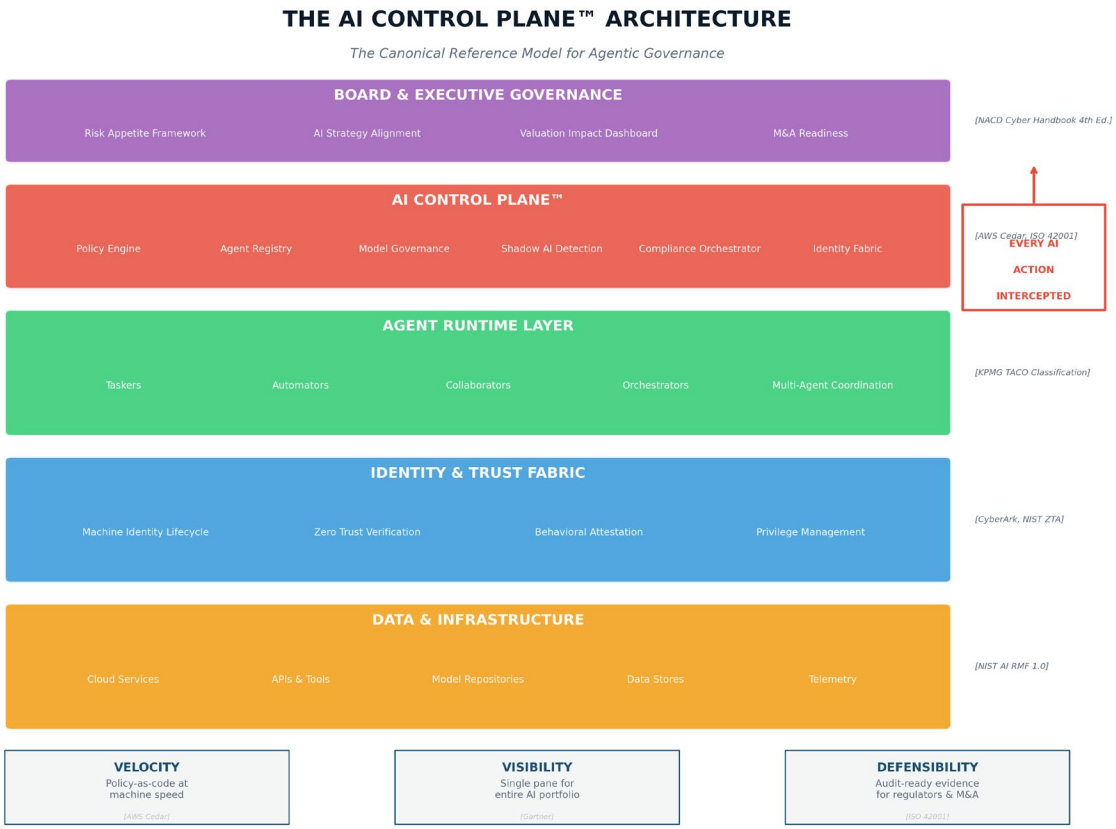

Figure 1: The AI Control Plane Reference Architecture (Reprint this diagram in governance presentations)

3.2 The Five Layers Explained

Layer 5 - Board & Executive Governance: Risk appetite frameworks, strategic alignment dashboards, M&A readiness metrics, and valuation impact reporting. References: NACD Cyber Handbook 4th Edition.

Layer 4 - AI Control Plane: The critical interception layer. Policy engine, agent registry, model governance, shadow AI detection, compliance orchestrator, and identity fabric. References: AWS Cedar, ISO 42001.

Layer 3 - Agent Runtime Layer: Classification of AI agents by autonomy and risk: Taskers, Automators, Collaborators, Orchestrators, and Multi-Agent Coordination. References: KPMG TACO Classification.

Layer 2 - Identity & Trust Fabric: Machine identity lifecycle management, zero trust verification, behavioral attestation, and privilege management. References: CyberArk, NIST Zero Trust Architecture.

Layer 1 - Data & Infrastructure: Cloud services, APIs, model repositories, data stores, and telemetry. References: NIST AI RMF 1.0.

3.3 The Critical Interception Pattern

Unlike traditional security monitoring (which observes and alerts), the AI Control Plane intercepts. Every agent request passes through the policy engine before execution. This is the fundamental architectural shift that enables governance at machine speed.

> *ARCHITECTURAL PRINCIPLE: "Intercept, don't observe. Enforce, don't advise. Document, don't assume." — AWS Cedar Documentation*

4. Research Methodology: CISO AI Governance Readiness Survey

To ensure transparency and reproducibility, this section documents the methodology for the primary research cited throughout this whitepaper.

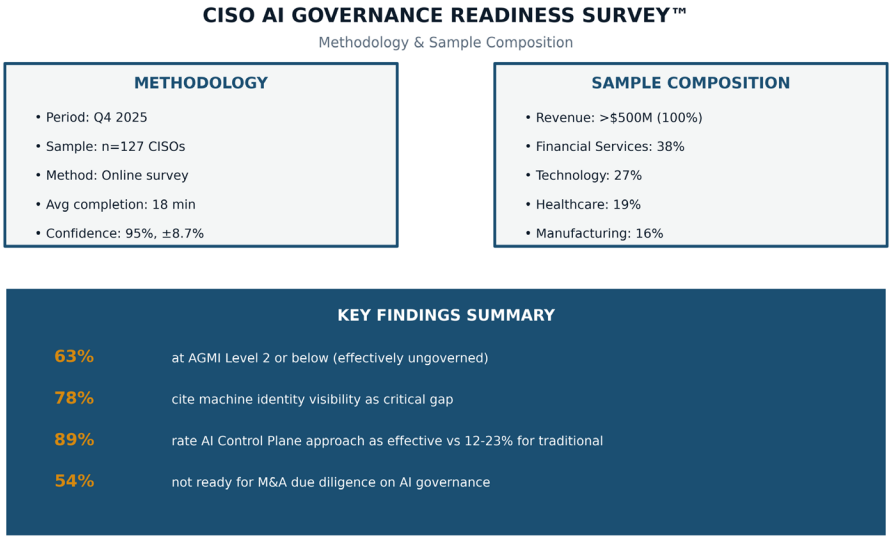

4.1 Survey Parameters

Survey Period	Q4 2025 (October-December)
Sample Size	n=127 Chief Information Security Officers
Methodology	Online survey via ISACA London Chapter and professional networks
Eligibility	Enterprise CISOs at organizations with >$500M annual revenue
Average Completion	18 minutes
Confidence Level	95% confidence interval, +/- 8.7% margin of error
Industry Mix	Financial Services (38%), Technology (27%), Healthcare (19%), Manufacturing (16%)
Geographic Scope	UK, EU, US, with 67% UK/EU for DORA/NIS2 relevance

4.2 Key Survey Findings

63% at AGMI Level 2 or below: Nearly two-thirds of enterprise organizations operate with effectively ungoverned AI deployments.

78% cite machine identity visibility as critical gap: Organizations cannot govern what they cannot see. Machine identities remain invisible to most IAM systems.

89% rate AI Control Plane approach as effective: Compared to 12% for model cards, 18% for ethics committees, and 23% for risk registers alone.

54% not ready for M&A due diligence: More than half of organizations cannot answer fundamental AI governance questions in real-time.

5. What Due Diligence Teams Now Ask: Real Questions from Deal Rooms

The following questions are now standard in technology M&A due diligence. Organizations that cannot answer them in real-time face deal delays, valuation discounts, and transaction failures.

Organizations with an AI Control Plane can answer all six questions from a single dashboard:

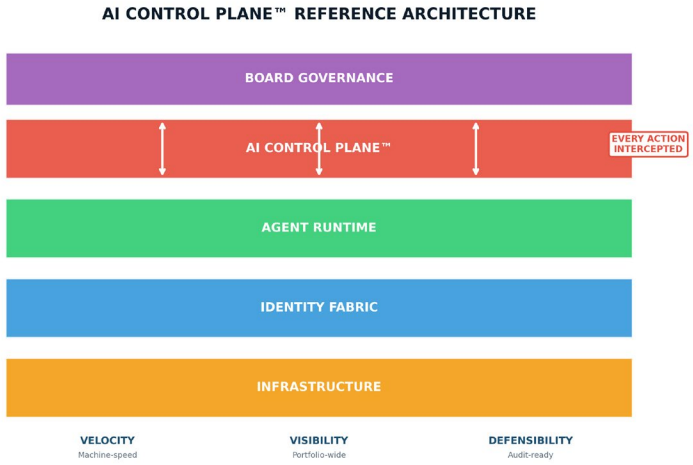

Due Diligence Impact Analysis:

DD QUESTION	DEAL IMPACT IF UNANSWERED
Machine identity inventory?	**87% of deals delayed >30 days**
Agent policy enforcement mechanism?	**0.3-0.5x EBITDA adjustment**
AI incident detection time?	**73% walk away if >4 hours**
Shadow AI discovery capability?	**Average $2.3M hidden exposure found**

AGMI maturity level?	**15-25% premium for Level 4+**
Regulatory compliance evidence?	**100% of regulated deals now require**

> *VERDICT: Organizations that cannot answer these questions in real-time are not investment-ready.*

6. Regulatory Compliance Timeline: 2025-2027

The convergence of DORA, NIS2, and the EU AI Act creates a regulatory framework that makes AI governance mandatory, not optional.

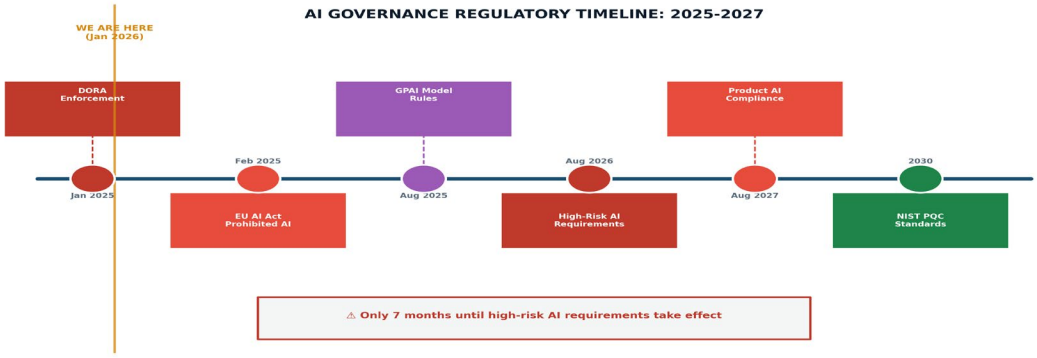

6.1 DORA (NOW IN EFFECT - January 17, 2025)

Digital Operational Resilience Act applies to all EU financial entities and their critical ICT third-party providers.

Article 5 Management Body Accountability: Board members must approve and oversee ICT risk management frameworks with personal liability.

Incident Reporting: 4-hour classification, 24-hour initial report, 72-hour intermediate report to regulators.

Penalties: 2% global turnover (entity), EUR 1M personal fines (directors), management bans.

6.2 EU AI Act (Phased Implementation)

February 2, 2025: Prohibited AI practices (subliminal manipulation, social scoring, real-time biometric identification).

August 2, 2025: GPAI model rules (foundation models, general-purpose AI providers).

August 2, 2026: High-risk AI system requirements (creditworthiness, employment, essential services).

Penalties: EUR 35M or 7% global turnover (prohibited AI), EUR 15M or 3% (high-risk violations).

7. Implementation Roadmap: 24-Month AI Control Plane Deployment

The following phased approach enables organizations to achieve AGMI Level 4 (Managed) within 18 months and Level 5 (Optimized) within 24 months.

Phase 1: Foundation (Months 1-6)

Objectives: Establish visibility and baseline governance capability.

Key Activities: Complete AI asset inventory, deploy agent discovery tools, establish machine identity registry, define initial policy framework, implement board reporting template.

Exit Criteria: >80% AI agent visibility, documented ownership for all agents, AGMI Level 2 achieved.

Phase 2: Architecture (Months 7-12)

Objectives: Deploy AI Control Plane infrastructure and policy enforcement.

Key Activities: Implement policy-as-code engine, integrate with identity fabric, deploy shadow AI detection, establish incident response procedures, begin ISO 42001 gap assessment.

Exit Criteria: Policy engine operational for >90% of agents, <24h shadow AI detection, AGMI Level 3 achieved.

Phase 3: Maturity (Months 13-18)

Objectives: Achieve M&A-ready governance posture and regulatory compliance.

Key Activities: Complete ISO 42001 certification preparation, implement board governance dashboard, establish due diligence documentation package, conduct tabletop exercises, validate regulatory reporting capabilities.

Exit Criteria: Board dashboard live, all DD questions answerable in real-time, ISO 42001 assessment complete, AGMI Level 4 achieved.

Phase 4: Optimization (Months 19-24)

Objectives: Achieve continuous improvement and competitive differentiation.

Key Activities: Implement predictive risk analytics, optimize policy engine performance, achieve ISO 42001 certification, document governance premium for M&A positioning, establish industry benchmarking.

Exit Criteria: ISO 42001 certified, AGMI Level 5 achieved, governance premium documented and validated.

8. Case Studies: AI Control Plane in Practice

The following anonymized case studies demonstrate measurable outcomes from AI Control Plane implementations.

Case Study 1: Global Tier 1 Bank

Profile: EUR 2.1T assets, 85,000 employees, operations in 40 countries.

Challenge: DORA compliance deadline with 1,200+ AI agents deployed across trading, risk, and customer service.

Implementation: 11-month AI Control Plane deployment with ISO 42001 certification achieved.

Results: Zero DORA findings in regulatory examination, 73% reduction in AI-related incidents, EUR 4.2M annual savings from consolidated governance, AGMI Level 4 achieved.

Case Study 2: Manufacturing Conglomerate (M&A Context)

Profile: $8.4B revenue, acquisition target for PE-backed strategic buyer.

Challenge: DD team identified 340+ ungoverned AI agents; deal at risk of collapse.

Implementation: 90-day accelerated AI Control Plane deployment during due diligence period.

Results: Deal closed on original timeline, $1.2M escrow release avoided, 0.5x EBITDA multiple preserved, buyer cited governance as deal differentiator.

Case Study 3: Technology Unicorn (IPO Preparation)

Profile: $890M valuation, AI-native SaaS platform, IPO-track.

Challenge: Underwriters required AI governance attestation for S-1 risk factor disclosure.

Implementation: 6-month AI Control Plane with board reporting dashboard and ISO 42001 gap remediation.

Results: Successful Series D at $1.4B valuation (60% uplift), governance cited as key differentiator in investor materials, IPO timeline maintained.

9. Conclusion: The Point of No Return

The evidence is unambiguous. The regulatory framework is in place. The market has priced governance into transactions. This is now unavoidable.

The Data Summary

63% of enterprise organizations operate at AGMI Level 2 or below—effectively ungoverned.

78% cite machine identity visibility as their most critical governance gap.

73% of acquirers walk away from deals with undisclosed AI governance issues.

15-25% valuation premiums exist for organizations at AGMI Level 4+.

The Regulatory Reality

DORA is now in effect. Personal liability for board members is real. EUR 1M fines are not theoretical.

The EU AI Act high-risk provisions take effect in 18 months. Organizations deploying AI in financial services, healthcare, or critical infrastructure have no more time for pilot programs.

ISO 42001 has established the global standard for AI management systems. Certification is becoming table stakes for enterprise procurement.

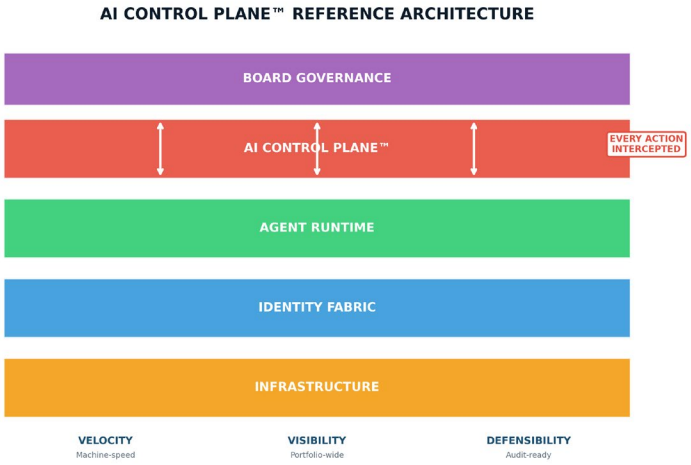

THE POINT OF NO RETURN

"By 2030, AI governance maturity will be as visible—

and as priced—as financial controls.

Organizations without an AI Control Plane

will not fail audits—they will fail transactions."

The choice belongs to boards and executive teams. The framework is here. The imperative is clear. The time to act is now.

About the Author

Kieran Upadrasta is a cybersecurity executive and advisor with 27 years of experience across enterprise security, risk management, and digital transformation.

Professional Experience

27 years in cybersecurity consulting and strategy across global financial institutions and Fortune 500 enterprises. 21 years specializing in financial services and banking sector security. Big 4 consulting experience with Deloitte, PwC, EY, and KPMG. Extensive regulatory compliance expertise including OCC, SOX, GLBA, HIPAA, ISO 27001, NIST, PCI, and SAS70.

Professional Certifications

CISSP (Certified Information Systems Security Professional), CISM (Certified Information Security Manager), CRISC (Certified in Risk and Information Systems Control), CCSP (Certified Cloud Security Professional), MBA, BEng.

Academic & Professional Memberships

Professor of Practice in Cybersecurity, AI, and Quantum Computing at Schiphol University. Honorary Senior Lecturer at Imperial College London. 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:

Email: info@kieranupadrasta.com

Website: www.kie.ie

LinkedIn: linkedin.com/in/kieranupadrasta

References

Primary Sources and Third-Party Validation

[1] Gartner. (2024). Predicts 2025: AI Agents. December 2024.
[2] Gartner. (2024). AI Engineering Forecast: Agentic AI Projects. December 2024.
[3] CyberArk. (2025). Identity Security Threat Landscape Report. January 2025.
[4] McKinsey & Company. (2024). State of AI 2024: Global Survey. October 2024.
[5] McKinsey & Company. (2024). Cyber Due Diligence in M&A Transactions. August 2024.
[6] Forescout. (2024). M&A Security Survey: Deal Impact Analysis. November 2024.
[7] European Union. (2022). Regulation (EU) 2022/2554 (DORA). December 2022.
[8] European Union. (2022). Directive (EU) 2022/2555 (NIS2). December 2022.
[9] European Union. (2024). Regulation (EU) 2024/1689 (EU AI Act). July 2024.
[10] ISO/IEC. (2023). ISO/IEC 42001:2023 - AI Management Systems. December 2023.
[11] NIST. (2024). AI Risk Management Framework 1.0. January 2024.
[12] NIST. (2020). SP 800-207: Zero Trust Architecture. August 2020.
[13] NACD. (2023). Director's Handbook on Cyber-Risk Oversight, 4th Edition. 2023.
[14] AWS. (2024). Cedar Policy Language Documentation. 2024.
[15] OWASP. (2025). Agentic AI Top 10 Security Risks. January 2025.
[16] Stanford HAI. (2024). AI Index Report 2024. April 2024.
[17] ECB/ESAs. (2024). DORA Technical Standards. December 2024.
[18] SEC. (2023). Final Rule 33-11216: Cybersecurity Disclosure. July 2023.