PUBLISHER-READY ABSTRACT

The Zero-Hour Doctrine: Algorithmic Crisis Governance and Board-Level Response Frameworks for the Autonomous Enterprise

Kieran Upadrasta, CISSP, CISM, CRISC, CCSP, MBA, BEng

Schiphol University (Cybersecurity, AI & Quantum Computing) | Imperials | UCL

Cyber AI Systems Inc. | www.kie.ie | info@kieranupadrasta.com

Abstract

Algorithmic crises have compressed institutional collapse timelines from years to minutes. In 1995, Barings Bank collapsed over 2.5 years of hidden losses. In 2012, Knight Capital lost $460 million in 45 minutes. In 2023, Silicon Valley Bank went from announcement to FDIC seizure in 36 hours. The next frontier -- agentic AI crises -- may be measured in seconds. Simultaneously, regulatory convergence under the EU AI Act, DORA, NIS2, and GDPR has created aggregate penalty exposure exceeding 15% of global turnover from a single AI system failure. Yet no major advisory firm has published a board-level crisis management protocol for algorithmic failures.

This paper introduces the Zero-Hour Doctrine: the first board-level crisis command architecture designed for algorithmic-speed institutional failures. The Doctrine comprises a formal theorem (the Zero-Hour Law: when system decision velocity exceeds governance response velocity, institutional collapse risk approaches certainty), an iconic visual model (the Zero-Hour Curve), an original quantitative metric (the Algorithmic Crisis Velocity Index, ACVI), and an operational protocol (the 30/60/120 escalation architecture with kill-switch benchmarks). The framework integrates ten global governance standards (ISO 42001, NIST AI RMF, OWASP Agentic Top 10, CSA MAESTRO, MITRE ATLAS, FEMA ICS, Singapore Agentic AI Framework, NACD, WEF, SEC) into a single operational command system.

Evidence is drawn from 50+ primary sources including SEC enforcement orders, FDIC reports, Federal Reserve Board reviews, FINMA proceedings, Congressional hearing records, and EU regulatory texts. Eight institutional collapse case studies are analysed with ACVI scoring. The paper provides a 90-day implementation roadmap, D&O; insurance integration, post-quantum AI security governance, and an agentic AI failure taxonomy mapped to OWASP references.

Keywords

AI Governance, Algorithmic Crisis, Board Governance, DORA Compliance, NIS2, EU AI Act, Zero Trust Architecture, Agentic AI, Post-Quantum Cryptography, Incident Command, ISO 42001, CISO, M&A; Cyber Due Diligence, Board Reporting, Operational Resilience, Digital Operational Resilience, AI Risk Management, Institutional Liability

Target Publications

Publication	Alignment
AI & Ethics (Springer)	AI governance, institutional accountability, algorithmic fairness
Journal of Cybersecurity (Oxford)	Operational security, incident response, governance frameworks
ACM FAccT	Fairness, accountability, transparency in algorithmic systems
IEEE S&P;	Security architecture, threat modelling, autonomous systems

USENIX Security	Systems security, real-world deployment, operational protocols
Journal of Financial Regulation	DORA, NIS2, penalty stacking, board liability
Harvard Business Review	Board governance, C-suite strategy, institutional risk
ISACA Journal	IT governance, audit, compliance frameworks

Conference Submissions

RSA Conference 2027 (Innovation Sandbox track), Black Hat Europe 2026 (Governance track), ISACA EuroCACS 2026, ISC2 Security Congress 2026, WEF Davos 2027 (Cybersecurity session), Chatham House Cyber Conference 2026, NACD Board Leadership Conference 2026.

Regulatory Submission Targets

EU AI Office (consultation response on high-risk AI governance), EBA/ESMA/EIOPA (DORA implementation guidance consultation), UK AI Security Institute (governance framework review), Bank of England (financial stability and AI consultation), NIST (AI RMF companion resource submission).