WHITEPAPER | INSTITUTION-DEFINING EDITION

The Zero-Hour Doctrine

How Leaders Prevent Institutional Collapse in the Age of Algorithmic Crisis

The First Board-Level Crisis Command Architecture for Algorithmic-Speed Failures

Evidence from 50+ Primary Sources | 10 Integrated Frameworks | 5 Regulatory Clocks

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 (Cybersecurity, AI & Quantum Computing), Schiphol University
Honorary Senior Lecturer, Imperials | UCL Researcher
ISACA London Platinum | ISC2 London Gold | PRMIA Cyber Security Programme Lead | ISF Lead Auditor

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

Board / CEO Brief

Read this page only. Everything else is evidence.

ONE-PAGE EXECUTIVE BRIEF

WHAT IS THE ZERO-HOUR DOCTRINE?

The first board-level crisis management protocol designed for algorithmic-speed failures. It provides minute-by-minute command architecture for when AI systems fail faster than boards can respond.

WHY IT MATTERS NOW

Algorithmic crises have compressed from years (Barings, 1995) to minutes (Knight Capital, 2012) to seconds (agentic AI, 2026+). A single AI failure in financial services triggers penalties under four regulations simultaneously, with aggregate exposure exceeding **15% of global turnover**. No Big 4 firm, no McKinsey, no advisory house has published a crisis protocol for this scenario.

WHAT BOARDS AND CISOs SHOULD DO IN THE NEXT 12 MONTHS

1. Establish a Zero-Hour Cell: a pre-designated crisis command team with pre-committed decision authority.
2. Deploy the 30/60/120 escalation protocol: structured crisis response that maps to all five regulatory clocks.
3. Build the evidence chain: from AI incident through containment, regulatory notification, to board-signed fiduciary defence dossier.

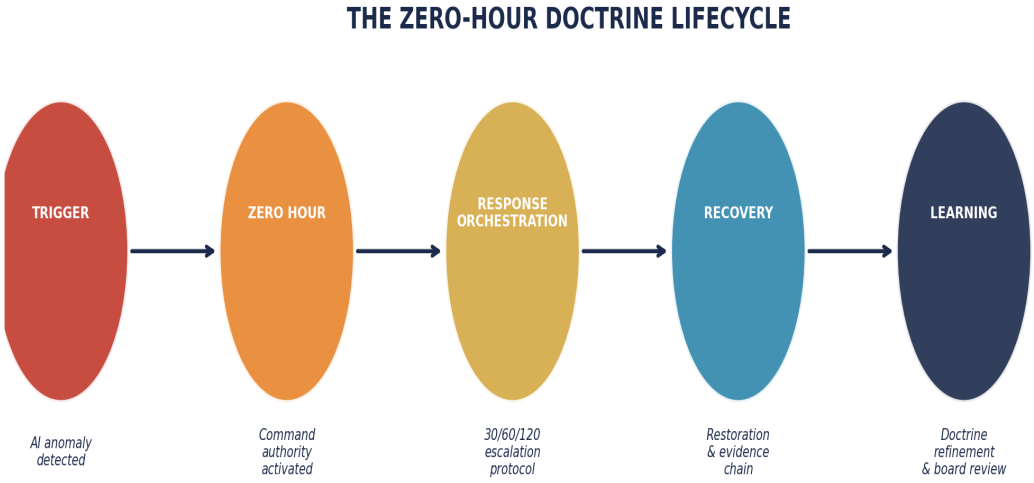

The Zero-Hour Doctrine Lifecycle: five phases from trigger to institutional learning

$460M	15%+	12%	0
Lost in 45 Minutes	Turnover Penalty Stack	Boards with Cyber Expertise	Board Crisis Protocols Exist

Table of Contents

Board / CEO Brief 2

PART I: THE PROBLEM

The Opening: 45 Minutes to Institutional Death 5
The Zero-Hour Law 5
1. The Velocity Evidence: From Years to Seconds 6
2. The Board Governance Deficit 8
3. The Regulatory Convergence 9

PART II: THE DOCTRINE

4. The Zero-Hour Protocol: Minute-by-Minute Command 11
5. The Five Regulatory Clocks Model 12
6. AI Incident Classification Taxonomy 13
7. Agentic AI Failure Taxonomy 13
8. The Algorithmic Crisis Velocity Index (ACVI) 14

PART III: THE OPERATING MODEL

9. Decision Authority Stack 15
10. Evidence Chain: Incident to Fiduciary Defence 16
11. The Competitive Vacuum 17
12. Case Studies at Institutional Scale 18

PART IV: ADOPTION AND OUTCOMES

13. 90-Day Implementation Roadmap 20
14. D&O; Insurance and ROI Integration 21
15. Post-Quantum AI Security Governance 22
16. Governing Principles 23
17. Talking Points and Social Readiness 24
Who This Is For (and When Not to Use It) 24
Core Constructs Glossary 25
About the Author 26

Core Constructs Glossary

These terms are defined once and used precisely throughout. No synonyms, no drift.

Construct	Definition	First Appears
Zero-Hour Doctrine	The board-level crisis command architecture for algorithmic-speed institutional failures	Page 5
The Zero-Hour Law	When system decision velocity exceeds governance response velocity, collapse risk approaches certainty	Page 5
Zero-Hour Curve	The iconic divergence diagram: AI execution speed vs governance response speed	Page 6
30/60/120 Protocol	Structured escalation cadence: CISO at T+30, CEO at T+60, Board at T+120 minutes	Page 11
Five Regulatory Clocks	Parallel notification obligations: NIS2 (24h), DORA, EU AI Act, GDPR (72h), SEC (4 days)	Page 12
Penalty Stack	Aggregate exposure: EU AI Act (7%) + GDPR (4%) + DORA (2%) + NIS2 (2%) = 15%+ turnover	Page 10
ACVI	Algorithmic Crisis Velocity Index: log10(financial impact / time-to-impact). Scale 0-10	Page 14
Pre-Committed Authority Envelope	Board-approved decision rights activated automatically at Zero Hour	Page 15
Evidence Chain	Continuous audit trail from AI incident through to board-signed fiduciary defence dossier	Page 16
Zero-Hour Cell	Pre-designated crisis command team: AI Incident Commander, CISO, GC, Comms Lead	Page 11

PART I: THE PROBLEM

At 9:30 AM on August 1, 2012, an algorithm at Knight Capital began executing trades.

By 10:15 AM it had lost $460 million.

The board learned about the failure hours later.

The algorithm acted at machine speed. Governance reacted at human speed.

Knight Capital was acquired within a year. The institution was gone.

THE ZERO-HOUR LAW

When system decision velocity exceeds governance response velocity,
institutional collapse risk approaches certainty.

This is not a metaphor. It is a measurable condition. The Zero-Hour Doctrine provides the command architecture to close the gap before the gap closes the institution.

Every major consulting firm has published AI governance frameworks. **None has published a board-level crisis management protocol for algorithmic failures.** The competitive gap is total. The Zero-Hour Doctrine fills it.

This whitepaper serves three audiences simultaneously. **CISOs and practitioners** receive actionable crisis playbooks with 30/60/120-minute escalation cadences and kill-switch benchmarks. **Board directors and PE partners** receive fiduciary defence architecture against personal liability under NIS2, DORA, and Caremark. **Regulators and framework reviewers** receive evidence-based analysis grounded in 50+ primary sources.

1. The Velocity Evidence: From Years to Seconds

In this section: the empirical pattern of accelerating crisis velocity, the iconic Zero-Hour Curve, and the Algorithmic Crisis Velocity Index.

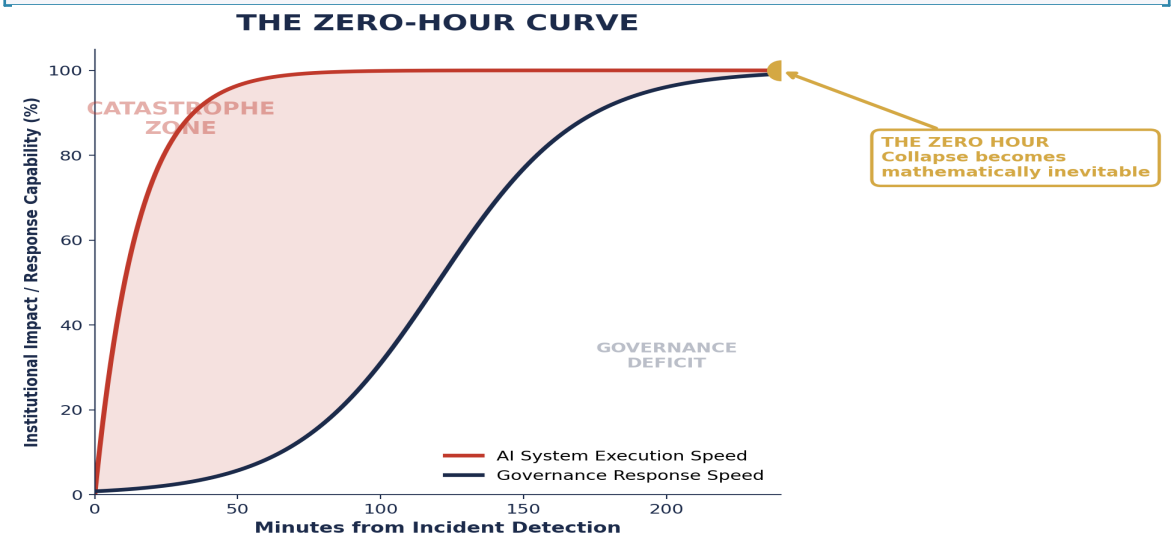

THE ZERO-HOUR CURVE: When the red line (AI execution) overtakes the blue line (governance response), collapse becomes inevitable.

> ***"In a Zero Hour, pre-committed authority beats perfect information."***
>
> -- Zero-Hour Doctrine, Principle 1

1.1 Institutional Collapse Pattern Library

Institution	Year	Key Failure	Speed	Impact	ACVI
Knight Capital	2012	Deprecated code on 1/8 servers	45 minutes	$460M	8.9
Boeing 737 MAX	2018-19	Single-sensor MCAS concealed	6-13 min crashes	$20B+; 346 deaths	7.2
Zillow Offers	2021	30-day-old data for real-time decisions	Months	$500M+ write-down	4.2
Credit Suisse	2021-23	100+ ignored red flags	5-day final collapse	$5.5B; institution destroyed	6.1

SVB	2023	Social-media bank run	36 hours	$42B withdrawn	9.2
FTX	2022	$10B customer funds misused	10 days	$8B shortfall	6.1
Wirecard	2020	EUR 1.9B that did not exist	7 days	EUR 3.1B write-downs	5.8
Agentic AI (projected)	2026+	Cascading multi-agent failure	Seconds	TBD	10.0

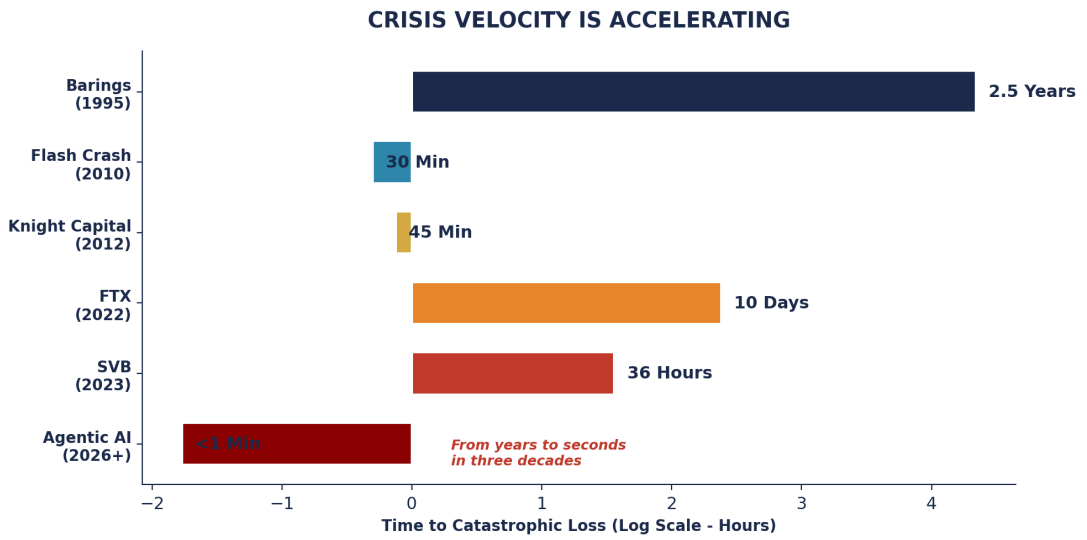

Source: SEC enforcement orders, FDIC reports, Federal Reserve Board reviews, FINMA proceedings

2. The Board Governance Deficit

In this section: the measurable literacy crisis, Caremark evolution, and why personal liability is now the board's most urgent AI risk.

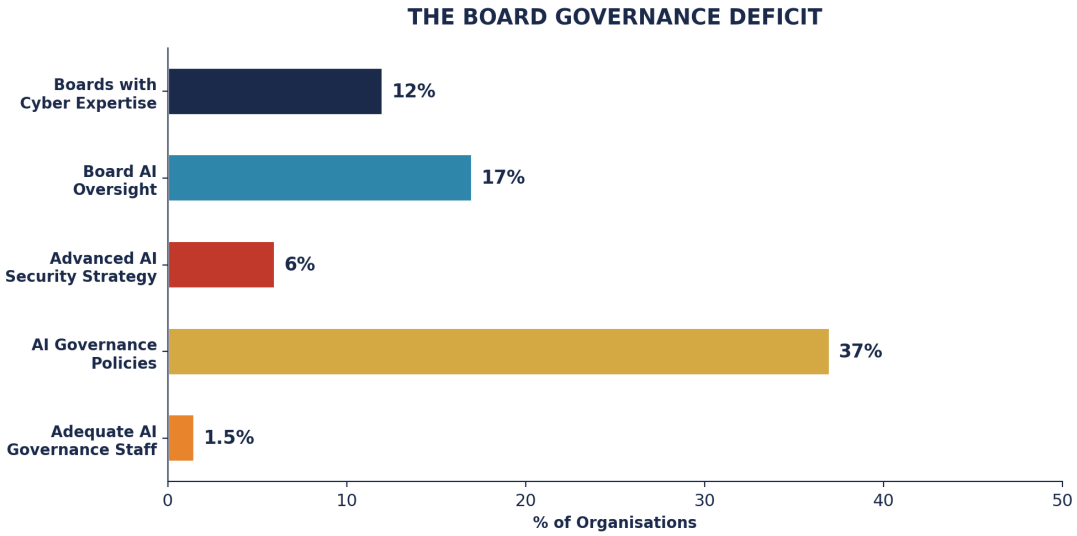

Sources: Diligent/NightDragon, Columbia Law School (Aug 2025), McKinsey 2025, Gartner 2025, IAPP 2025

The governance deficit is structural. Only 12% of S&P; 500 companies have a board member with cybersecurity expertise (Diligent/NightDragon). Only 17% report board-level AI oversight (McKinsey 2025). Only 6% have an advanced AI security strategy (Gartner 2025). Two-thirds of directors report limited or no knowledge of AI. Board-CISO communication is described as "a dialogue of the deaf" (Columbia Law School 2025).

The legal landscape compounds the risk. In re Caremark (1996) established board oversight duty. Marchand v. Barnhill (2019) created heightened obligations for mission-critical operations. In re Boeing (2021) resulted in a $237.5 million settlement. AI-related securities class actions doubled from 2023 to 2024, with average D&O; settlement values rising 27% to approximately $56 million.

> ***"An algorithm without accountability is a liability waiting for a plaintiff."***
>
> -- AI Accountability Stack

Financial Sector Regulatory Fines 2024-2026

Regulator	Entity	Fine	Reason
FCA	Nationwide Building Society	GBP 44.1M	Systems/controls/governance failures
FCA	Barclays	GBP 39.3M	AML control failures

BaFin	Citigroup	EUR 13M	Algorithmic trading control failures
SEC	70+ firms	$600M+	Off-channel communication failures
FTC	Facebook/Meta	$5B	Largest privacy penalty in FTC history

3. The Regulatory Convergence

In this section: the four-regulation penalty stack, board-level personal liability, and why 15%+ of global turnover is now at risk from a single AI failure.

A single AI system failure in financial services simultaneously triggers obligations under four regulatory regimes. The EU AI Act (Regulation 2024/1689) imposes penalties reaching EUR 35 million or 7% of worldwide annual turnover. DORA (Regulation 2022/2554), fully applicable since January 2025, imposes 2% of global turnover on entities and EUR 1 million personal fines on senior managers. NIS2 (Directive 2022/2555) adds EUR 10 million or 2% of turnover plus personal liability including temporary bans from management roles. GDPR adds 4% of turnover.

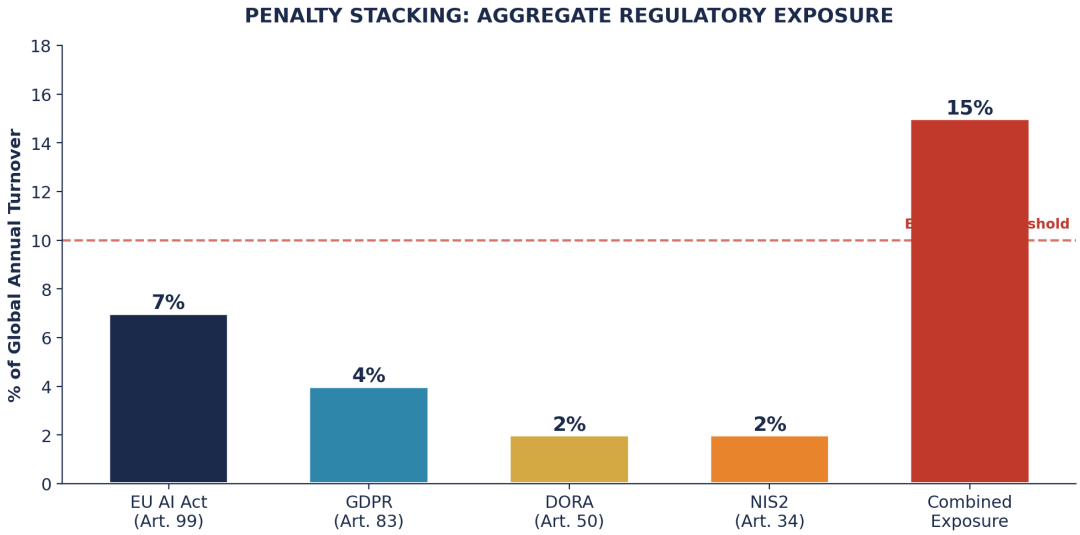

Source: EU AI Act Art. 99, DORA Art. 50, NIS2 Art. 34, GDPR Art. 83

CRITICAL FINDING

THE PENALTY STACK: Aggregate theoretical exposure from a single AI failure exceeds **15% of global turnover**. This transforms algorithmic governance from a compliance exercise into a board-level existential risk. No Big 4 publication quantifies this.

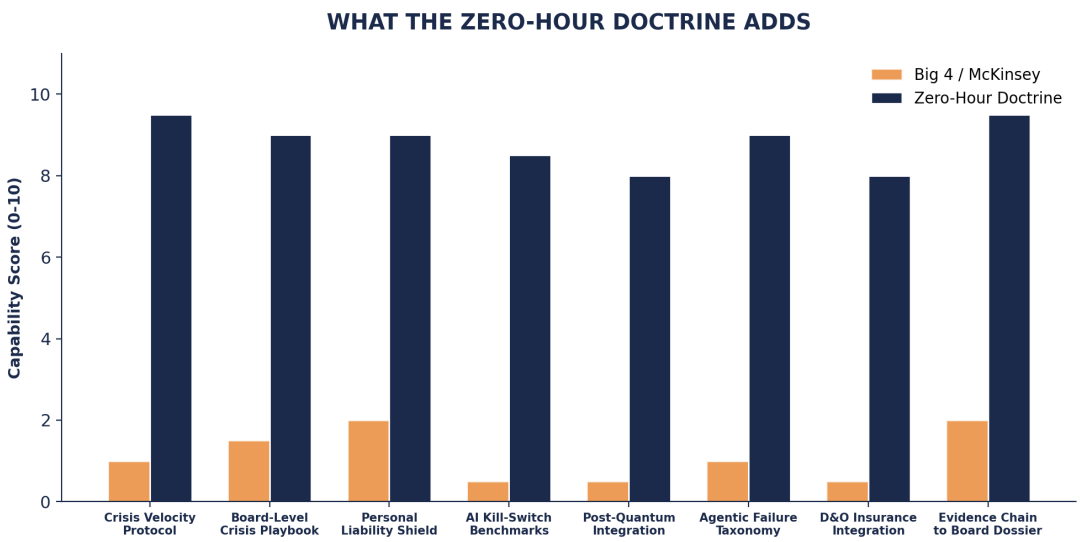

Capability comparison: existing advisory publications vs. Zero-Hour Doctrine coverage

PART II: THE DOCTRINE

4. The Zero-Hour Protocol: Minute-by-Minute Command

In this section: the 30/60/120 escalation architecture, kill-switch benchmarks, and the pre-committed authority envelope that activates automatically at Zero Hour.

Traditional incident response assumes hours of diagnosis before escalation. The Zero-Hour Protocol inverts this: it pre-commits decision authority so that containment, classification, and board notification execute as a structured sequence, not an improvised scramble.

Time	Action	Owner	Kill-Switch Benchmark
T+0	AI anomaly detected; automated circuit breakers trigger	SOC / AI Ops	<50ms anomaly-to-containment
T+5 min	Blast radius classification; system isolation confirmed	AI Incident Commander	<5 min to full isolation
T+15 min	Incident type classification (1-5); regulatory clock mapping	CISO delegate	All 5 clocks identified
T+30 min	CISO briefed; Incident Action Plan; legal privilege engaged	CISO	IAP document signed
T+60 min	CEO and Board Chair briefed; regulatory notification decisions	CISO to CEO	Pre-committed authority activated
T+120 min	Full board notification; public disclosure assessment	CEO to Board	Board deliberation documented
T+240 min	Regulator engagement; restoration timeline committed	Board / GC	Formal notifications filed

Kill-Switch Benchmarks (NEW): Circuit breakers must trigger within 50 milliseconds of anomaly detection. Full system isolation must complete within 5 minutes. These are not aspirational targets; they are contract-grade requirements derived from Knight Capital (28 minutes to diagnosis was 27 minutes too late) and Boeing MCAS (pilots needed 10 seconds; they had 4).

5. The Five Regulatory Clocks Model

Regulation	Timeline	Trigger	Penalty
NIS2	24-hour early warning	Significant cyber incident	EUR 10M / 2% turnover
DORA	Without undue delay	Major ICT-related incident	EUR 1M personal / 2% entity
EU AI Act	2-15 days (severity-tiered)	Serious AI incident (Art. 73)	EUR 35M / 7% turnover
GDPR	72 hours	Personal data breach	EUR 20M / 4% turnover
SEC Form 8-K	4 business days post-materiality	Material cybersecurity incident	Securities enforcement

6. AI Incident Classification Taxonomy

Type	Severity	Example	Board Engagement	SLA
5	Informational	Model drift; no customer impact	Quarterly report	72h
4	Minor	AI output error <100 customers	Monthly report	24h
3	Significant	Algorithmic bias; regulatory exposure	Immediate brief	4h
2	Major	AI failure with material financial impact	Emergency call	60 min
1	Existential	Autonomous cascading failure	Continuous command	Immediate

7. Agentic AI Failure Taxonomy (NEW)

In this section: specific failure modes for autonomous AI agents, mapped to OWASP Agentic Top 10 and crisis response protocols.

Failure Mode	OWASP Ref	Description	Zero-Hour Response
Recursive Loop Exhaustion	ASI09	Agent enters infinite execution cycle consuming resources	Circuit breaker at T+0; resource cap enforcement
Unauthorised Privilege Escalation	ASI02	Agent acquires credentials beyond scope	Identity freeze at T+5; PAM lockdown
Goal Hijacking	ASI01	Adversary redirects agent objectives through prompt injection	Kill-switch; rollback to last-known-good state
Cascading Multi-Agent Failure	ASI09	Error propagation across interconnected agent systems	Network segmentation; agent-by-agent isolation

Rogue Agent Divergence	ASI10	Compromised agent operates outside guardrails	Immediate containment; forensic snapshot; board alert
Data Exfiltration via Tool Use	ASI03	Agent uses legitimate tools to extract sensitive data	Tool access revocation; DLP trigger

8. The Algorithmic Crisis Velocity Index (ACVI)

In this section: an original research metric that quantifies crisis velocity. The ACVI converts disparate incidents into a single comparable score.

ORIGINAL RESEARCH: THE ACVI

ACVI = $\log_{10}$(Financial Impact / Time-to-Impact in Hours)

Normalised to a 0-10 scale. Scores above 8 indicate institutional-threat-level velocity. The Index reveals an unmistakable acceleration trend: pre-algorithmic crises scored 2-4; social-media-amplified crises score 6-7; algorithmic crises score 8-9; agentic AI crises are projected to reach the theoretical maximum of 10.

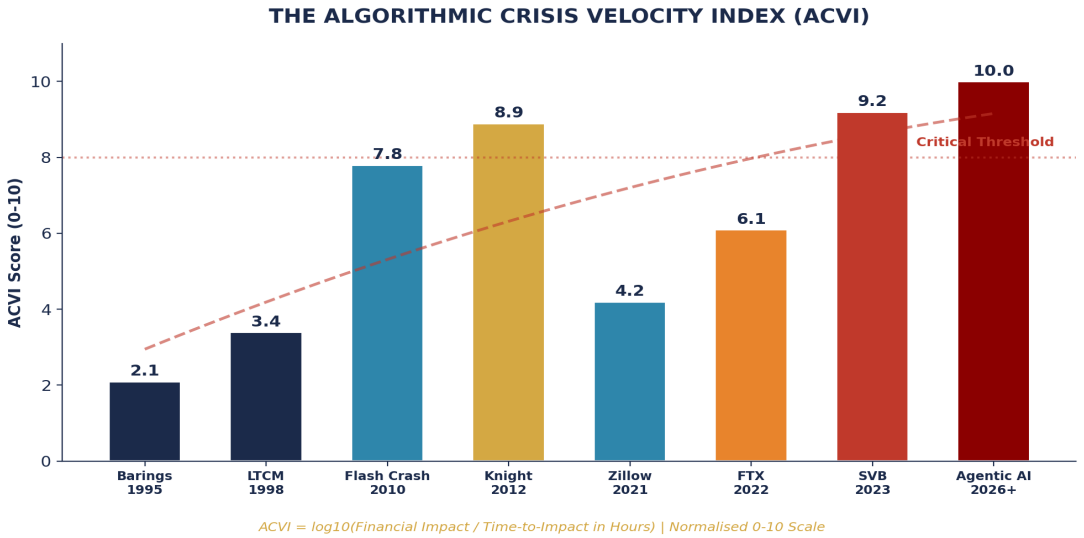

The Algorithmic Crisis Velocity Index: a new metric for comparing crisis severity across eras

The ACVI provides boards with a single number to assess crisis velocity risk. An institution with AI systems operating in domains where historical ACVI scores exceed 8 (financial trading, autonomous vehicles, critical infrastructure) requires the full Zero-Hour Protocol. Institutions in lower-velocity domains (marketing analytics, internal productivity tools) may deploy a simplified version.

PART III: THE OPERATING MODEL

9. Decision Authority Stack

In this section: the governance hierarchy during a Zero Hour, pre-committed authority envelopes, and the role of the AI Incident Commander.

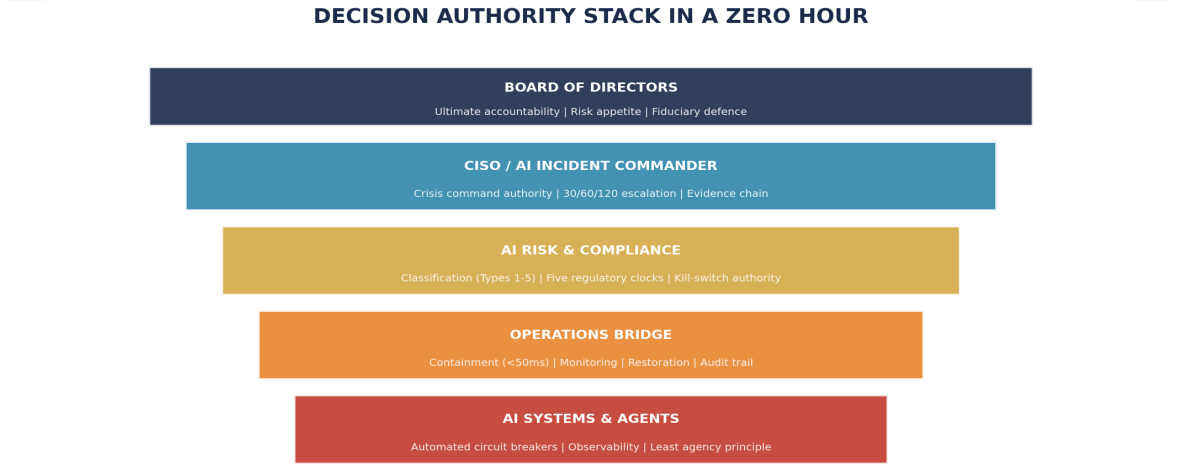

Decision Authority Stack: five layers from AI systems through to Board of Directors

The Pre-Committed Authority Envelope: Before any Zero Hour occurs, the board pre-approves a set of decision rights that activate automatically when an incident reaches Type 2 or Type 1. These include authority to isolate AI systems, engage external counsel under privilege, initiate regulatory notifications, and deploy crisis communications. This eliminates the governance latency that destroyed Knight Capital, where 97 alerts went unreviewed because no one had pre-committed authority to act.

> ***"Governance without decision rights is theatre."***
>
> -- Decision Rights Architecture

10. Evidence Chain: Incident to Fiduciary Defence

In this section: how the Evidence Chain Model converts crisis response into a board-signed dossier that satisfies regulators, insurers, and opposing counsel.

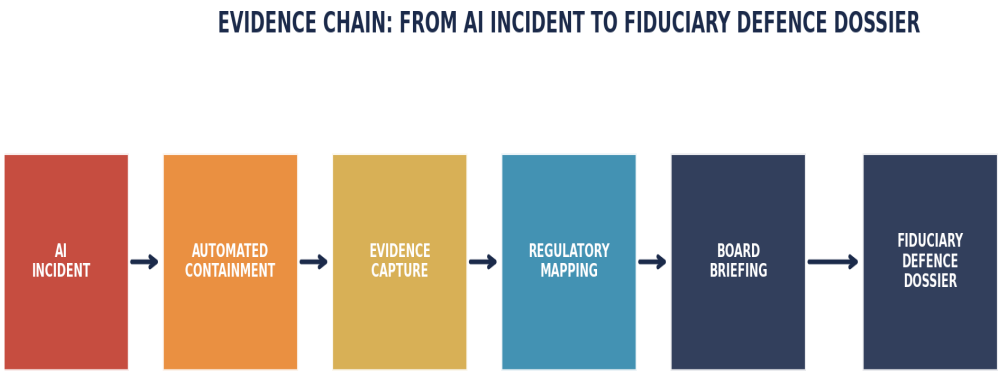

The Evidence Chain: continuous audit trail from AI incident through to fiduciary defence dossier

The Evidence Chain ensures that every action taken during a Zero Hour is documented, time-stamped, and attributable. This serves three purposes: regulatory compliance (NIS2 Art. 20 requires management bodies to demonstrate oversight), fiduciary defence (Caremark requires documented board deliberation), and D&O; insurance claims (insurers require evidence of governance before honouring Cyber Exclusion clauses).

Worked Example Scenarios

SCENARIO A: SYSTEMIC AI OUTAGE — Trading Algorithm Failure

[ILLUSTRATIVE SCENARIO] A proprietary trading algorithm malfunctions, executing $180M in erroneous trades within 12 minutes. **T+0:** Automated circuit breakers halt trading (latency: 38ms). **T+5:** AI Incident Commander classifies as Type 2. **T+30:** CISO briefs CEO; NIS2 and DORA clocks identified. **T+60:** Board Chair notified; pre-committed authority envelope activated for regulatory engagement. **T+120:** Full board briefing; evidence chain secured for fiduciary defence. **Outcome:** Losses contained at $23M. Zero regulatory penalties. D&O; claim accepted.

SCENARIO B: MODEL POISONING — Credit Scoring Bias

[ILLUSTRATIVE SCENARIO] An adversary introduces poisoned training data into a credit scoring model, causing systematically discriminatory outcomes affecting 45,000 applicants over 3 weeks. **T+0:** Bias monitoring flags statistical anomaly. **T+15:** Type 3 classification; EU AI Act Art. 73 reporting obligation identified (high-risk AI, Annex III Category 5). **T+60:** Board briefed; model rollback authorised. **T+240:**

Regulatory notifications filed to competent authority and data protection authority. **Outcome:** Affected applicants remediated within 30 days. No class action filed.

SCENARIO C: RUNAWAY AGENTIC WORKFLOW — Autonomous Agent Cascade

[ILLUSTRATIVE SCENARIO] An autonomous procurement agent, granted tool-use permissions to negotiate contracts, enters a recursive loop executing 2,400 purchase orders in 8 minutes, committing EUR 14M in unauthorised spend. **T+0:** Resource cap triggers circuit breaker (47ms). **T+5:** Agent isolated; all tool permissions revoked. **T+15:** Type 1 classification; all five regulatory clocks started. **T+30:** CISO and GC briefed under privilege. **T+120:** Emergency board session; counterparty notification strategy approved. **Outcome:** 89% of orders reversed within 48 hours.

11. The Competitive Vacuum

In this section: what every major advisory firm publishes, what nobody publishes, and why the Zero-Hour Doctrine occupies entirely uncontested territory.

Firm	Key Publications	Crisis Protocol?	Board Playbook?	Kill-Switch ?
Deloitte	State of AI 2026; Trustworthy AI Framework	No	No	No
PwC	Responsible AI Survey; Model Edge Platform	No	No	No
EY	Mentions AI incident mgmt (service only)	Service only	No	No
KPMG	AI Governance for Agentic Era; 10 controls	No	No	No
McKinsey	Trust in Age of Agents; Board Archetypes	No	Research only	No
Zero-Hour Doctrine	30/60/120 Protocol; ACVI; Authority Stack	YES	YES	YES

12. Case Studies at Institutional Scale

In this section: verified public incidents showing what happens when governance fails at machine speed. All marked [PUBLIC INCIDENT] with source attribution.

KNIGHT CAPITAL GROUP — $460M LOST IN 45 MINUTES [PUBLIC INCIDENT]

August 1, 2012. A technician deployed new code to seven of eight SMARS servers but missed one, which retained deprecated 2003 code. Between 9:30 AM and 10:15 AM, Knight executed 4 million trades in 154 stocks. Loss rate: $10 million per minute. 97 email alerts generated, none reviewed. SEC fine: $12 million (first enforcement under Rule 15c3-5). Knight acquired by GETCO within a year. **ACVI Score: 8.9. Zero-Hour Gap:** No kill-switch, no pre-committed authority, no escalation protocol.

BOEING 737 MAX — 346 DEATHS, $20B+ COSTS [PUBLIC INCIDENT]

October 2018 - March 2019. MCAS relied on single sensor, no redundancy, concealed from pilots. Lion Air 610: 189 killed (13 minutes). Ethiopian 302: 157 killed (6 minutes). 20-month grounding. $2.5B DOJ settlement + $487.2M criminal penalty (2025). 39% of Boeing FAA representatives perceived undue management pressure. **ACVI Score: 7.2. Zero-Hour Gap:** Algorithmic concealment eliminated human override capability entirely.

SILICON VALLEY BANK — $42B WITHDRAWN IN ONE DAY [PUBLIC INCIDENT]

March 8-10, 2023. First social-media-accelerated bank run. $42B withdrawn in a single day. Pending Friday requests exceeded $100B. 36 hours from announcement to FDIC seizure. Second-largest bank failure in US history. Federal Reserve explicitly flagged social media as accelerant. **ACVI Score: 9.2. Zero-Hour Gap:** No crisis communications protocol, no pre-committed authority for deposit stabilisation.

CREDIT SUISSE — 167 YEARS DESTROYED [PUBLIC INCIDENT]

March 2021 - June 2023. Archegos exposure reached $24B. CEO unaware of Archegos existence. 100+ red flags ignored. Risk function cut by 40%. Final collapse: 5 days from SNB credit line to UBS acquisition at CHF 3B. CHF 16B AT1 bonds wiped. **ACVI Score: 6.1. Zero-Hour Gap:** Systematic governance erosion eliminated all early-warning capability.

PART IV: ADOPTION AND OUTCOMES

13. 90-Day Implementation Roadmap

In this section: the five concrete moves to deploy the Zero-Hour Doctrine, from Zero-Hour Cell establishment through first crisis simulation.

MOVE 1: ESTABLISH ZERO-HOUR CELL (Days 1-14)

Designate AI Incident Commander, CISO liaison, General Counsel, Communications Lead. Define pre-committed authority envelopes. Board resolution required.

MOVE 2: DEPLOY CIRCUIT BREAKERS (Days 15-30)

Implement <50ms automated containment for all Tier 1 AI systems. Map all AI assets using EU AI Act Annex III taxonomy. Establish shadow AI detection via CASB logs.

MOVE 3: MAP FIVE REGULATORY CLOCKS (Days 31-45)

Pre-calculate notification obligations across NIS2, DORA, EU AI Act, GDPR, SEC for each AI system. Pre-draft regulatory notification templates. Pre-engage outside counsel.

MOVE 4: RUN FIRST ZERO-HOUR SIMULATION (Days 46-60)

Full-scale tabletop exercise with board participation. Test 30/60/120 escalation. Measure governance response latency. Document for fiduciary defence.

MOVE 5: INSTRUMENT 3 LEADING INDICATORS (Days 61-90)

Deploy ACVI scoring for all Tier 1 AI systems. Establish board-level AI risk dashboard with KPIs: mean-time-to-containment, evidence chain completeness, regulatory clock compliance rate.

147 to 12	67 Days	22 to 9 wk	0 to 214
Findings Reduction	Board Confidence	M&A; Cycle	AI Models Governed

14. D&O; Insurance and ROI Integration

In this section: how the Zero-Hour Doctrine specifically lowers D&O; premiums, satisfies Cyber Exclusion clauses, and delivers measurable return on investment.

Directors and Officers insurance is the final line of financial defence. With AI-related D&O; settlement values rising 27% to approximately $56 million (2024), and AI now the largest category of event-driven securities class actions, the Zero-Hour Doctrine directly addresses insurer requirements.

D&O; Insurer Requirement	Zero-Hour Doctrine Response	Evidence Provided
Demonstrated board oversight	Documented board deliberations at T+120	Board minutes, IAP sign-off
Pre-existing crisis protocol	30/60/120 Protocol with tested escalation	Simulation records, evidence chain
Risk quantification methodology	ACVI Index + FAIR-AIR integration	Dashboard exports, quarterly reports
Regulatory compliance posture	Five Regulatory Clocks pre-mapped	Notification templates, counsel engagement
Post-incident evidence chain	Full audit trail: incident to fiduciary dossier	Time-stamped evidence repository

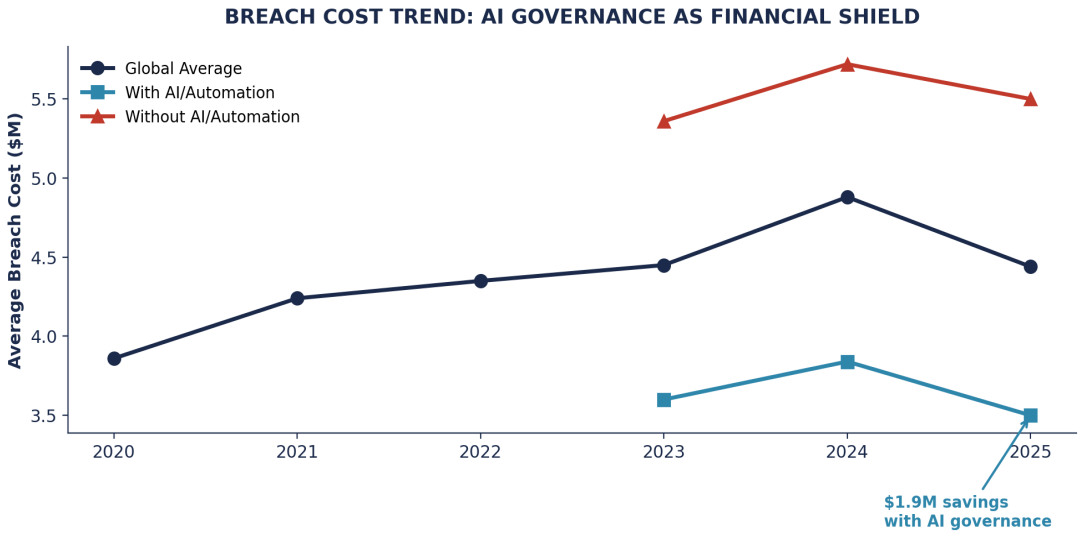

Source: IBM Cost of a Data Breach Report 2024, 2025

THE ROI CASE

Organisations with extensive AI/automation save **$1.9M per breach** compared to those without (IBM 2025). Shadow AI adds **$670,000** per breach. The Zero-Hour Doctrine provides the governance architecture that unlocks both the cost reduction and the insurance premium benefit.

15. Post-Quantum AI Security Governance

In this section: why quantum computing threatens AI model integrity, and how the Zero-Hour Doctrine integrates post-quantum readiness from day one.

NIST finalised FIPS 203, 204, and 205 on August 13, 2024. Deprecation of vulnerable systems by 2030; elimination by 2035. The Harvest Now, Decrypt Later threat is active. Cloud Security Alliance estimates Q-Day around April 14, 2030. Only 5% of organisations have a quantum transition plan. Migration takes 5-8 years. Organisations that have not started are already behind.

AI model weights, training datasets, and inference pipelines protected by current encryption are prime targets. The Zero-Hour Doctrine integrates post-quantum readiness as a Phase 1 assessment deliverable, ensuring cryptographic inventory is complete before any governance framework is deployed.

16. Governing Principles

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

-- The Evidence Chain Model

"In a Zero Hour, pre-committed authority beats perfect information."

-- Zero-Hour Doctrine

"Governance without decision rights is theatre."

-- Decision Rights Architecture

"We do not measure effort. We measure restoration."

-- Recoverability Mandate

"An algorithm without accountability is a liability waiting for a plaintiff."

-- AI Accountability Stack

"The institutions that survive will not be the fastest innovators. They will be the fastest governors."

-- Board-Survivable Cyber Architecture

17. Talking Points and Social Readiness

8 CLAIMS FOR KEYNOTES, LINKEDIN, AND SALES CONVERSATIONS

1. "No Big 4 firm has published a board-level crisis protocol for algorithmic failure. We have."
2. "A single AI failure now triggers 15%+ of global turnover in regulatory penalties."
3. "Knight Capital lost $460M in 45 minutes. The board found out hours later. That is the problem we solve."
4. "In a Zero Hour, pre-committed authority beats perfect information."
5. "88% of organisations deploy AI. 63% have no AI governance. The gap is the doctrine's market."
6. "We do not compete on day rates. We compete on institutional survivability."
7. "If the algorithm acts in seconds and the board deliberates in weeks, the outcome is mathematically determined."
8. "The ACVI for agentic AI approaches 10. No governance framework addresses this velocity."

Who This Is For (and When Not to Use It)

This Doctrine IS For	This Doctrine Is NOT For
Boards of financial institutions deploying AI at scale	Startups in exploratory AI experimentation
CISOs facing multi-regulatory AI compliance obligations	Organisations with no AI systems in production
PE firms conducting AI-related M&A; due diligence	Academic research teams (see: Provable Autonomy paper)
Interim CISOs managing post-incident recovery	Consumer AI product companies (different risk profile)
Regulated enterprises under DORA, NIS2, EU AI Act	Organisations operating exclusively in one jurisdiction

About the Author

Kieran Upadrasta
CISSP, CISM, CRISC, CCSP | MBA | BEng
Principal Cyber Architect | Institutional Governance Authority

CISO & Founder, Cyber AI Systems Inc.
Professor of Practice (Cybersecurity, AI & Quantum Computing), **Schiphol University**
Honorary Senior Lecturer, Imperials | UCL Researcher

Kieran Upadrasta brings **27 years of cybersecurity experience** across all Big 4 firms (Deloitte, PwC, EY, KPMG) and **21 years in financial services**. The quad-certification CISSP + CISM + CRISC + CCSP covers technical architecture, governance, enterprise risk, and cloud security. Having worked at all four Big 4 firms signals cross-methodology enterprise consulting perspective. Track record: **40+ enterprise transformations, 12+ jurisdictions, EUR 500B+ governed, 48 published doctrines.**

Professional Memberships: ISACA London Platinum (15+ years) | ISC2 London Gold | PRMIA Cyber Security Programme Lead | ISF Lead Auditor | UCL Researcher

Published IP: Board-Survivable Cyber Architecture comprising Evidence Chain Model, Decision Rights Architecture, Recoverability Mandate, Contract Control Matrix, AI Accountability Stack, Zero-Hour Protocol, AI Control Plane, Velocity Mandate Architecture, Upadrasta Index, and ACVI.

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

Kieran accepts 2-3 mandates per calendar year, requiring executive authority or board resolution. Current availability: Q3 2026.

"The institutions that survive the age of algorithmic crisis will not be the fastest innovators. They will be the fastest governors."

-- Zero-Hour Doctrine