THE CISO'S AI FIREWALL

FREE SAMPLE CHAPTER

A CISO'S GUIDE TO SECURING, GOVERNING, AND DEPLOYING ARTIFICIAL INTELLIGENCE

STEVE SHARMA

Author of THE CYBER DEFENSE BLUEPRINT

CYBER DEFENSE PRESS

WHY THIS BOOK EXISTS

Your Security Stack Was Built for a World Where Software Didn't Talk Back

> *Generative AI became production infrastructure faster than governance could catch up – and most enterprises are running it on controls engineered for deterministic systems that don't learn, generate, or act.*

14	38	95	5
CHAPTERS ACROSS 4 PARTS	ORIGINAL FIGURES & FRAMEWORKS	TERM GLOSSARY A–Z	STAGE DECISION JOURNEY™

Who This Book Is For

- **CISOs and security leaders** asked "are we exposed?" by a board that can't yet define the risk
- **GRC and risk functions** mapping AI obligations against the EU AI Act, NIST AI RMF, ISO/IEC 42001, and APRA CPS 234
- **Security architects** building the control layer for LLMs, RAG, and agentic AI – not just the policy layer
- **Boards and executives** who need the vocabulary before they can own the accountability

What's Inside

The book follows the **CISO's Decision Journey™** – Assess → Govern → Defend → Monitor → Improve – from the strategic case for treating AI as a distinct threat-model domain, through governance and secure MLOps, to threat modelling, red-teaming, the AI Security Controls Architecture, incident response, and the agentic frontier. Every chapter closes with a Board Translation, a Regulatory Alignment table, and a practical template you can put to work the same week.

WHAT READERS GET

A field-tested governance system, not a theory book: original frameworks (AI Firewall Governance Model™, STRIDE-AI Threat Model Extension™, Five-Layer AI Security Controls Architecture™, AI Use-Case Risk-Tiering Framework™), reusable templates, and a 95-term glossary – built by a practising security architect, mapped to the standards your auditors already ask about.

The Next Pages

What follows is the opening of **Chapter 1 – Why AI Demands a New Security Paradigm**, condensed from the full edition. It sets up the argument the rest of the book operationalises: that speed, scale, and autonomy make AI risk categorically different from the risk your existing programme was built to manage.

FULL EDITION

Table of Contents

HOUSE FRAMEWORKS REFERENCED THROUGHOUT

CISO's Decision Journey™ · AI Firewall Governance Model™ · STRIDE-AI Threat Model Extension™ · Five-Layer AI Security Controls Architecture™ · AI Use-Case Risk-Tiering Framework™ — cross-mapped to OWASP LLM Top 10 (2025), NIST AI RMF, ISO/IEC 42001, MITRE ATLAS, EU AI Act (Regulation (EU) 2024/1689), and APRA CPS 234.

CHAPTER 1 · SAMPLE · PART I — THE STRATEGIC IMPERATIVE

ASSESS

Why AI Demands a New Security Paradigm

EXECUTIVE SUMMARY

Artificial intelligence has moved from pilot projects to production infrastructure faster than any enterprise technology in three decades – and faster than the governance frameworks designed to control it. This chapter establishes why the security paradigm that protected deterministic systems cannot, on its own, protect systems that learn, generate, and act.

CISO NOTE

In every AI security programme I have reviewed, the gap was never tooling – it was vocabulary. Boards approved budgets for "AI risk" without anyone in the room being able to distinguish a model from an agent, or a prompt injection from a data leak. The CISOs who succeeded did one thing first: they established a shared, plain-English taxonomy before writing a single policy.

1.1 The CISO's AI Moment: From IT Project to Boardroom Mandate

Consider the position most security leaders now occupy. Within eighteen months, generative AI moved from a curiosity in the innovation lab to embedded functionality in productivity suites, customer-service platforms, code repositories, and decision-support systems. Employees adopted it before procurement approved it. Vendors embedded it before contracts addressed it. And regulators – led by the EU AI Act, with enforcement of high-risk obligations commencing August 2026 – began legislating for it before most enterprises had inventoried it.

The governance implication is direct: AI is no longer a technology project that the CISO can delegate or defer. It is a board-level accountability question. When an AI system leaks customer data through a manipulated prompt, hallucinates a contractual commitment, or makes a biased automated decision, the questions directed at leadership will not be technical.

1.3 The Gap Between Deterministic and AI Threat Models

A well-crafted sentence, embedded in an email, a document, or a webpage the model later reads, can function as executable instruction. Its trust model is distributed across a supply chain of pre-trained models, third-party datasets, and external APIs that the enterprise neither built nor can fully inspect.

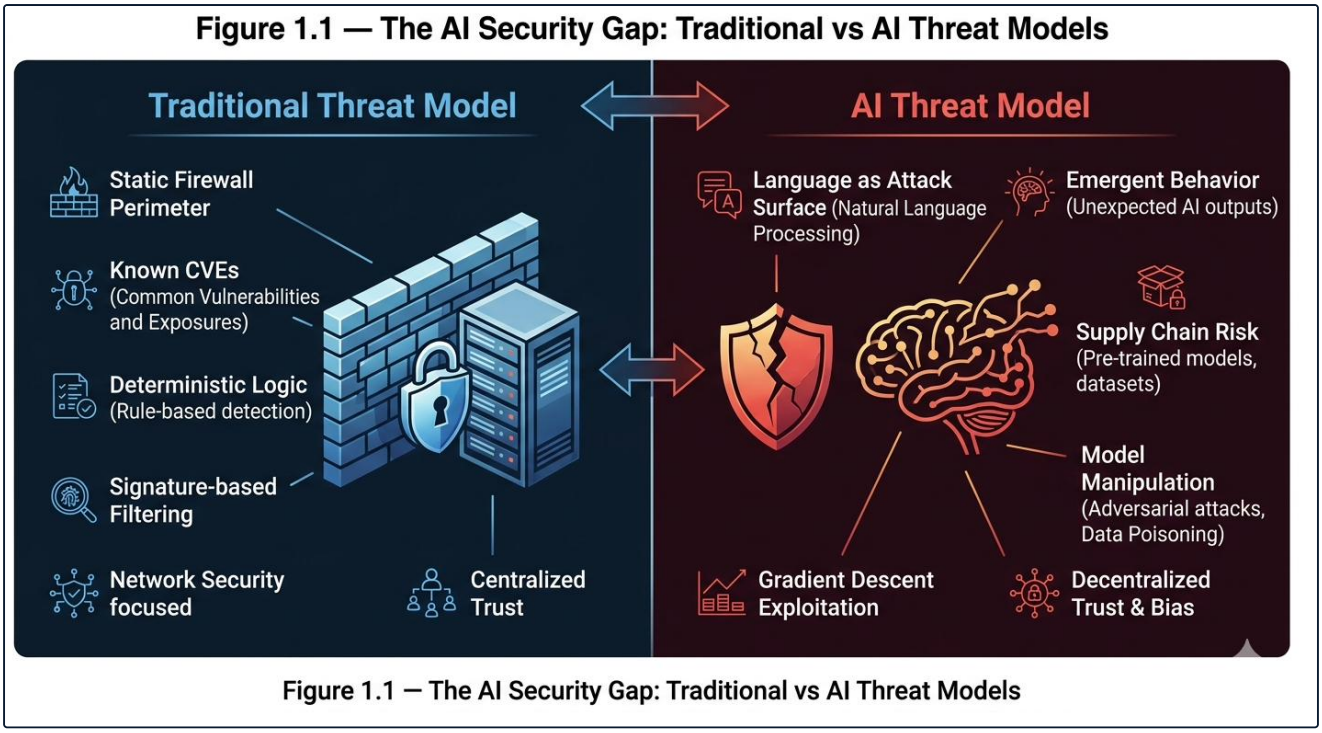

Figure 1.1 — The AI Security Gap: Traditional vs AI Threat Models. The left panel shows the deterministic assumptions of conventional security; the right panel shows the emergent, language-driven, supply-chain-dependent properties of AI systems.

> **KEY TAKEAWAY**
>
> Your security stack was engineered for deterministic systems with enumerable flaws. AI is neither. Treat AI as a distinct threat-model domain in the risk register – not a new asset class inside the old model.

ASSESS

Speed, Scale, Autonomy – and the Unpredictability at the Centre

1.4 Three Properties, One Unpredictable Interaction

If the deterministic-vs-AI gap explains why AI breaks the old threat model, this section explains why AI incidents behave differently once they occur. Three properties interact: speed, scale, and autonomy. Each is individually manageable. In combination, they produce the defining characteristic of AI risk – unpredictability.

- **Speed.** AI systems operate at machine tempo. Incident response designed around human decision latency starts the race already behind.
- **Scale.** A single flawed model replicates its flaw across every decision it touches – uniformly, silently, for months.
- **Autonomy.** Systems empowered to act – agents with credentials, tools, and goals – remove the human checkpoint that conventional governance assumes.

Figure 1.2 — Speed-Scale-Autonomy Risk Triangle

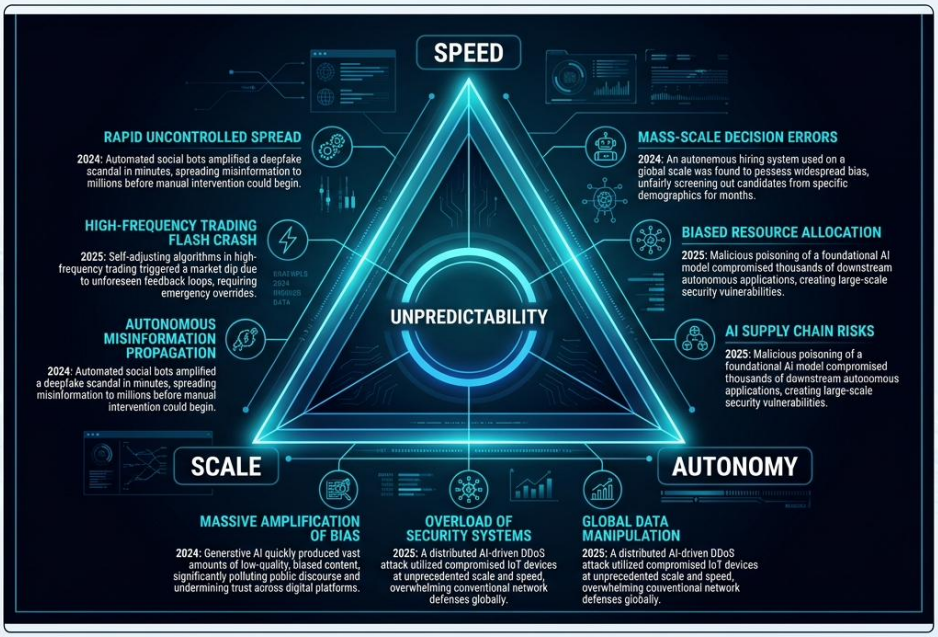

Figure 1.2 — Speed–Scale–Autonomy Risk Triangle. The three structural properties of AI risk occupy the vertices; their interaction produces the unpredictability at the centre.

BOARD TRANSLATION

If the board asks: "We already spend heavily on cybersecurity. Why does AI need its own programme?"
You answer: "Our current programme protects systems that behave the same way every time and whose flaws

can be listed and patched. AI systems learn, generate, and act – their primary attack surface is ordinary language, and their failures emerge at machine speed and enterprise scale. We are not replacing the security programme; we are extending it before regulators – starting with the EU AI Act in August 2026 – require us to evidence exactly that."

Where Chapter 1 Leads

Chapter 1 closes by reframing the central governance question – from "can this system fail?" to "what happens when it fails, and how quickly can we contain it?" Chapter 2 then builds the five-domain risk taxonomy that makes AI risk assignable rather than abstract, and every chapter after it adds one more layer of the governance system: regulatory mapping, secure MLOps, threat modelling, the Controls Architecture, incident response, and the agentic frontier.

CONTINUE READING

This preview is an edited excerpt of Chapter 1. The full edition includes the complete chapter – including Sections 1.2, 1.5, and 1.6, the full regulatory alignment table, Template 1.1 (Board Briefing), and the Top 5 Actions / Risks / Leadership Questions – plus thirteen further chapters, 38 figures, and the complete glossary.

GET THE FULL EDITION

The CISO's AI Firewall
Governing LLMs Before the Breach

Fourteen chapters. Thirty-eight original figures. Five original frameworks. One governance system – from first AI inventory to agentic AI at scale.

Steve Sharma

Principal Cybersecurity Architect · Founder & Director, Cybersecurity Link

Steve founded Cybersecurity Link in Melbourne, Australia, and wrote *The CISO's AI Firewall* from the vantage point of practising security architecture – building and reviewing the governance, controls, and incident-response systems that enterprises now need for LLMs, RAG, and agentic AI. The book's frameworks, including the CISO's Decision Journey™, were developed through that practice.

Format

- Paperback (6×9, 285pp)
- Kindle eBook
- EPUB3 (all major e-readers)

Where to Buy

- Amazon: amazon.com/dp/B0H7P45789
- IngramSpark retail network
- Draft2Digital & direct via publisher site

Details

- ISBN 978-1-7647921-0-3
- First Edition · 2026
- Cybersecurity Link, Melbourne

Connect with the author: LinkedIn /in/stevesharmacyber · cybersecuritylink.com.au