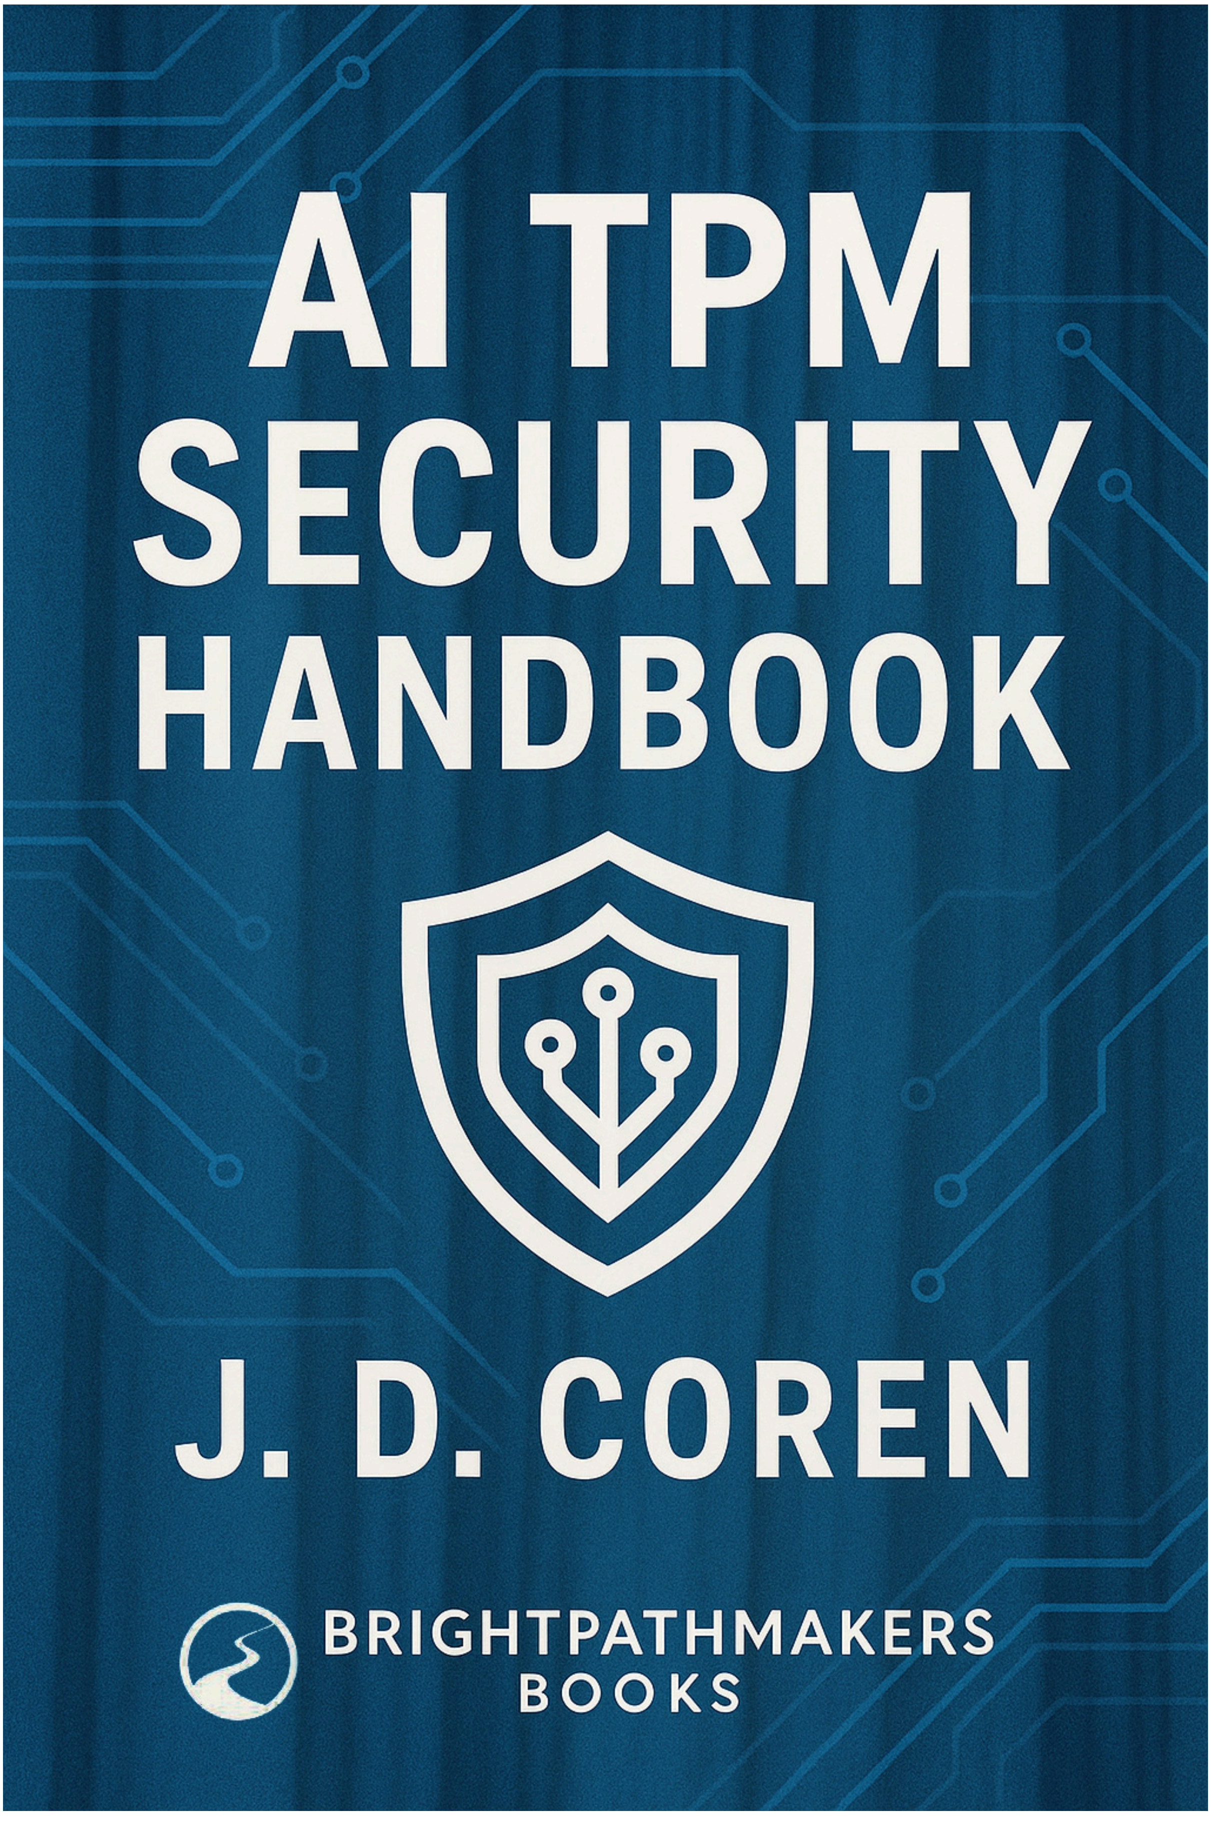

AI TPM SECURITY HANDBOOK

J. D. COREN

BRIGHTPATHMAKERS
BOOKS

AI TPM Security Handbook.. 3
Introduction..3
Why an AI TPM Security Handbook?...3
The role of security in the AI product lifecycle...4
Part I: The Intersection of AI TPM and Security...5
1. The AI Technical Program Manager’s Security Responsibilities............................. 5
2. Vulnerabilities and Risks Unique to AI Systems..5
3. Security by Design – Built-in Protection, Not an Afterthought.................................6
Part II: Core Security Pillars.. 8
4. Risk Management in AI Environments.. 8
5. Practical Frameworks & Approaches.. 9
6. Tools & Examples..9
7. Trusted Logging and Monitoring..10
8. Key Takeaways / Checklist...10
Part III: Integrated Approaches.. 12
9. Data Security and Privacy (Compliance, GDPR, AI Act).......................................12
10. Incident Response in AI-specific Contexts.. 12
11. Security Audits and Compliance in AI Projects..13
Part IV: Practical Frameworks and Case Studies.. 15
10. The TPM as Security Coordinator – Stakeholder Management..........................15
11. MVP and Prototype Development Through a Security Lens............................... 15
12. Case Studies: Integrating the Three Prototypes into Daily TPM Work.................16
Part V: Future and Strategic Directions...17
13. The Future of AI and Cybersecurity (Autonomous Defense, AI vs. AI Threats)...17
14. The Evolving Role of the TPM in the Next Decade... 17
15. Conclusion: Security as a Competitive Advantage..18
What’s next?...19
References...20

AI TPM Security Handbook

Introduction

Artificial intelligence is no longer a peripheral technology. It is woven into products, platforms, and services that shape entire industries. With this integration, the scope and stakes of security have expanded. What was once considered an engineering safeguard or compliance requirement has become a **strategic concern**.

The Technical Program Manager (TPM) sits at the intersection of technology, organization, and execution. In AI programs, this role carries an additional weight: to anticipate and manage risks that are both novel and systemic. Model integrity, data governance, adversarial resilience, ethical safeguards—these are not isolated checklists but ongoing challenges that cut across the entire product lifecycle.

This handbook is written for TPMs who must guide teams through that complexity. It is not a theoretical exploration of AI security. It is a **practical framework**: how to scope, prioritize, and integrate security in programs where AI plays a central role.

Why an AI TPM Security Handbook?

Most existing literature on AI security speaks either to researchers, who explore attack vectors and defenses, or to compliance professionals, who focus on regulation. **What is missing is a guide for program managers**—leaders who must translate these insights into execution.

In practice, TPMs face questions that are not answered in research papers:

- How should security be scoped in AI projects where requirements are incomplete and the technology is evolving?
- How do we balance delivery pressure with security reviews, especially when competitive timelines are tight?
- How can we build trust across engineering, product, and compliance stakeholders without losing momentum?

This handbook is designed to fill that gap. It positions security as a **program management discipline**, not just an engineering specialty. It shows how TPMs can orchestrate secure development, governance, and risk mitigation across the AI product lifecycle.

The role of security in the AI product lifecycle

Security in AI cannot be reduced to late-stage review or incident response. It begins with **data acquisition and governance**, continues through **model development and training**, extends into **deployment pipelines**, and remains active in **monitoring and response** once the system is live.

Each stage introduces unique risks:

- **Data**: bias, poisoning, unauthorized use.
- **Model**: adversarial robustness, intellectual property exposure.
- **Deployment**: supply chain security, configuration errors.
- **Operations**: monitoring drift, detecting abuse, ensuring responsible rollback.

For TPMs, the challenge is to integrate these considerations into program structure: timelines, milestones, decision points, and communication flows. Security cannot be treated as a separate track—it must be **designed into the lifecycle**.

This shift elevates the TPM role. Security is no longer reactive; it is a driver of credibility, trust, and competitive advantage.

Part I: The Intersection of AI TPM and Security

1. The AI Technical Program Manager’s Security Responsibilities

The role of the AI Technical Program Manager (TPM) extends far beyond coordination. In environments where AI drives core functionality, security is not a background concern—it is a determinant of trust and viability. The TPM becomes a steward of security risks, accountable for ensuring that both technical and business perspectives are aligned.

This responsibility has two distinct dimensions. First, the TPM must ensure that **engineering teams** recognize and address security requirements, integrating them into architecture, data pipelines, and model design. Second, the TPM must translate those technical realities into **business terms**: what risks exist, how they affect product viability, and what trade-offs decision-makers face.

Unlike traditional projects, AI programs evolve in uncertain conditions. Models adapt, data changes, and adversaries learn. The TPM’s responsibility is not to eliminate all risks, but to **structure risk visibility and mitigation** so that security is intentional, traceable, and continuously evaluated.

2. Vulnerabilities and Risks Unique to AI Systems

AI systems inherit traditional software vulnerabilities—supply chain flaws, misconfigurations, insecure APIs—but they also introduce a class of risks that are specific to their nature:

- **Data quality and provenance**: Models are only as robust as the data they are trained on. Poisoned or low-quality datasets can undermine performance or inject hidden vulnerabilities.

- **Bias and fairness**[1]: Security is not only about preventing breaches; biased models can create reputational and regulatory risks that are equally damaging.

- **Adversarial manipulation**: Inputs crafted to exploit model weaknesses can cause misclassification, bypass controls, or trigger unintended behavior.

[1] *Kahneman, D. (2011). Thinking, Fast and Slow.* (cognitive bias & decision context)
Mullainathan, S., & Shafir, S. (2013). Scarcity. (tunneling / narrowing under stress).