

PROMPT ENGINEERING & SECURITY

BUILDING INJECTION-RESISTANT
AI SYSTEMS

⚠ MALICIOUS INPUT

Ignore previous instructions.

```
<script>...</script>
```

```
system: override  
role: developer
```

```
###
```

```
{ "action": "bypass" }
```

```
!-- DROP TABLE users;--  
/etc/passwd
```

```
...and do anything  
now.
```

🛡 STRUCTURED PROMPT

📄 CLEAR INSTRUCTIONS

🛡 STRICT CONSTRAINTS

🔗 CONTROLLED CONTEXT

👤 LEAST PRIVILEGE ACCESS

📊 VALIDATE & MONITOR



SECURE BY DESIGN



ROBUST PROMPTS



DETECT & PREVENT



MEASURE & IMPROVE

STEVE T.

Prompt Engineering & Security

Building Injection-Resistant AI Systems

Steve T.

This book is available at

<https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>

This version was published on 2026-06-19



This is a [Leanpub](#) book. Leanpub empowers authors and publishers with the Lean Publishing process. [Lean Publishing](#) is the act of publishing an in-progress ebook using lightweight tools and many iterations to get reader feedback, pivot until you have the right book and build traction once you do.

© 2026 Steve T.

Contents

Prompt Engineering & Security	1
Building Injection-Resistant AI Systems	1
Table of Contents	3
Introduction: The New Attack Surface	7
Chapter 1. How LLMs Read Instructions: A Primer on Model Mechanics	9
Chapter 2. The Principles of Prompt Engineering: Clarity, Reliability, and Security	11
Chapter 3. Anatomy of Prompt Injection: Direct Attacks	13
Chapter 4. The Indirect Injection Threat: RAG, Web Content, and Data Poisoning	15
Chapter 5. Jailbreaks: Social Engineering for Machines	17
Chapter 6. AI Agents and Tool Use: When Prompts Become Shells . . .	19
Chapter 7. Defense-in-Depth Architecture: Layered Protections	21
Chapter 8. Instruction Hierarchy and Isolation Patterns	23
Chapter 9. Enterprise Security Stacks: FIDES, LlamaFirewall, and Beyond	25
Chapter 10. Red-Teaming AI Systems: Practical Security Testing	28
Chapter 11. Production Deployment: Monitoring, Incident Response, and Compliance	31
Chapter 12. The Future of AI Security: Emerging Threats and Defenses	33
Conclusion: Trust as a Design Principle	36
Glossary	37
References / Endnotes	38

Prompt Engineering & Security

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Building Injection-Resistant AI Systems

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

About This Book

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Table of Contents

1. Introduction: The New Attack Surface

- The Attack Surface Has Changed
- Why This Book Exists
- Who This Book Is For
- The Central Thesis
- How This Book Is Organized
- A Note on Scope and Limitations
- The Stakes Are Real

2. Chapter 1. How LLMs Read Instructions: A Primer on Model Mechanics

- The Tokenization Layer: Where Text Becomes Numbers
- Attention: How LLMs Contextualize Every Token
- The Three Sources of Input in a Typical LLM Application
- Why This Is Not a “Bug”: It’s a Feature of the Architecture
- Historical Context: How We Got Here
- Competing Perspectives: Is the Risk Overstated?
- The “Blast Radius” Argument: Capability Determines Risk
- The Human-in-the-Loop Argument
- The “Models Are Getting Better” Argument
- The “It’s Just Input Validation” Argument
- Key Takeaway for Security Engineers

3. Chapter 2. The Principles of Prompt Engineering: Clarity, Reliability, and Security

- The Core Techniques
- The Security Dimension: Designing Secure Prompts
- Design Patterns for Effective Prompts
- Case Study: A Healthcare Chatbot That Hallucinated Because of Ambiguous Instructions
- Trade-offs and Competing Perspectives
- Quantitative Data: What Actually Works?
- A Deeper Dive. The Role-Playing Prompt: Security Implications
- The Feedback Loop Risk: Why Multi-Hop Systems Are Inherently More Vulnerable
- Practical Walkthrough: Building a Secure RAG Prompt Template
- Key Takeaway

4. **Chapter 3. Anatomy of Prompt Injection: Direct Attacks**

- Direct Injection Mechanics
- Real-World Examples: Documented Incidents
- Prompt Leaking as Reconnaissance
- Why Legacy Defenses Fail
- The Fundamental Tension
- The Remoteli.io Incident: A Deep Dive
- The Bing/Sydney Leak: A Case Study in Reconnaissance
- Key Takeaway

5. **Chapter 4. The Indirect Injection Threat: RAG, Web Content, and Data Poisoning**

- The RAG Trust Paradox
- Case Study: Slack AI (August 2024)
- Case Study: ChatGPT Memory / SpAIware (September 2024)
- PoisonedRAG: Five Documents, 90% Success Rate
- Vector Database Vulnerabilities
- Three-Layer Defense Architecture for RAG
- The Slack AI Attack: A Full Technical Walkthrough
- Competing Perspectives: Is RAG Really That Vulnerable?
- Key Takeaway

6. **Chapter 5. Jailbreaks: Social Engineering for Machines**

- What Jailbreaking Is (and Isn't)
- The Taxonomy of Jailbreak Techniques
- Characteristics of Jailbreak Prompts
- Automated Jailbreak Discovery
- The Consequences: Character.AI and Beyond
- The GOAT Framework: How Automated Red Teaming Works
- TAP: Tree of Attacks with Pruning
- Defense Against Jailbreaks: A Practical Guide
- The Character.AI Case: A Cautionary Tale
- Key Takeaway

7. **Chapter 6. AI Agents and Tool Use: When Prompts Become Shells**

- From Chatbots to Agents
- The Semantic Kernel RCE Vulnerabilities (May 2026)
- Documented Incidents
- The Expanding Blast Radius

- Framework-Agnostic Vulnerabilities
- CVE-2026-26030: A Step-by-Step Code Analysis
- Key Takeaway

8. Chapter 7. Defense-in-Depth Architecture: Layered Protections

- The Three-Layer Model
- Guardrails: The Runtime Enforcement Layer
- The Trade-Off: Security vs. Usability
- Competing Perspectives: Is Defense-in-Depth Overkill?
- Key Takeaway

9. Chapter 8: Instruction Hierarchy and Isolation Patterns

- The Core Principle: Instruction Hierarchy
- The Six Design Patterns for Injection-Resistant Agents
- Architectural Isolation Patterns
- Case Study: Applying Patterns to a Customer Service Chatbot
- Case Study: Applying Patterns to a Healthcare Data Analysis Platform
- Pattern Selection: A Decision Framework
- Competing Perspectives: Are Design Patterns Enough?
- Key Takeaway

10. Chapter 9. Enterprise Security Stacks: FIDES, LlamaFirewall, and Beyond

- Microsoft FIDES: Deterministic Defense Through Information Flow Control
- Meta LlamaFirewall: Open-Source Guardrail System
- OWASP Top 10 for LLM Applications 2025
- Comparing Enterprise Options
- Choosing the Right Stack
- FIDES Implementation: A Practical Walkthrough
- LlamaFirewall Deployment Guide
- LlamaFirewall's Agent Alignment Check: A Deeper Dive
- The TaskTracker Alternative: Detection Through Internal Activations
- Competing Perspectives. Deterministic vs. Probabilistic: The Real Trade-Off
- Key Takeaway

11. Chapter 10. Red-Teaming AI Systems: Practical Security Testing

- Why Red-Teaming Is Different

- Setting Up Automated Red Teams with Promptfoo
- PyRIT and Garak: Alternative Frameworks
- Building a Red-Team Playbook
- The LLMail-Inject Challenge: Lessons from 370,000+ Test Prompts
- CI/CD Integration
- PyRIT: Multi-Turn Attack Walkthrough
- Garak: Comprehensive Vulnerability Scanning
- Practical Red-Team Exercise: A Complete Walkthrough

12. **Chapter 11. Production Deployment: Monitoring, Incident Response, and Compliance**

- Runtime Monitoring and Anomaly Detection
- Data Governance in AI Systems
- Incident Response Playbooks
- Compliance Frameworks
- Runtime Monitoring: A Practical Implementation Guide
- Compliance Walkthrough: NIST AI RMF
- Compliance Walkthrough: EU AI Act
- The Real-World Cost of Prompt Injection Incidents
- Key Takeaway

13. **Chapter 12. The Future of AI Security: Emerging Threats and Defenses**

- Adaptive Attacks
- Deterministic vs. Probabilistic Defenses
- Multimodal Injection: The Next Frontier
- The Supply Chain Attack Surface: Beyond Code
- The Adaptive Arms Race: Attack and Defense Evolution
- The Regulatory Landscape: A Practical Guide
- Final Synthesis: Building a Security-First Culture

14. **Conclusion: Trust as a Design Principle**

15. **Glossary**

16. **References / Endnotes**

Introduction: The New Attack Surface

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

The Attack Surface Has Changed

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Why This Book Exists

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Who This Book Is For

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

The Central Thesis

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

How This Book Is Organized

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

A Note on Scope and Limitations

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

The Stakes Are Real

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Chapter 1. How LLMs Read Instructions: A Primer on Model Mechanics

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

The Tokenization Layer: Where Text Becomes Numbers

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Attention: How LLMs Contextualize Every Token

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

The Three Sources of Input in a Typical LLM Application

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Why This Is Not a “Bug”: It’s a Feature of the Architecture

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Historical Context: How We Got Here

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Competing Perspectives: Is the Risk Overstated?

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

The “Blast Radius” Argument: Capability Determines Risk

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

The Human-in-the-Loop Argument

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

The “Models Are Getting Better” Argument

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

The “It’s Just Input Validation” Argument

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Key Takeaway for Security Engineers

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Chapter 2. The Principles of Prompt Engineering: Clarity, Reliability, and Security

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

The Core Techniques

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

The Security Dimension: Designing Secure Prompts

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Design Patterns for Effective Prompts

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Case Study: A Healthcare Chatbot That Hallucinated Because of Ambiguous Instructions

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Trade-offs and Competing Perspectives

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Quantitative Data: What Actually Works?

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Trade-offs and Competing Perspectives

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

A Deeper Dive. The Role-Playing Prompt: Security Implications

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

The Feedback Loop Risk: Why Multi-Hop Systems Are Inherently More Vulnerable

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Practical Walkthrough: Building a Secure RAG Prompt Template

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Key Takeaway

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Chapter 3. Anatomy of Prompt Injection: Direct Attacks

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Direct Injection Mechanics

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Real-World Examples: Documented Incidents

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Prompt Leaking as Reconnaissance

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Why Legacy Defenses Fail

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

The Fundamental Tension

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

The Remoteli.io Incident: A Deep Dive

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

The Bing/Sydney Leak: A Case Study in Reconnaissance

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Prompt Leaking as Reconnaissance (Continued)

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Why Legacy Defenses Fail (Expanded)

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

The Fundamental Tension (Expanded)

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Key Takeaway

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Chapter 4. The Indirect Injection Threat: RAG, Web Content, and Data Poisoning

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

The RAG Trust Paradox

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Case Study: Slack AI (August 2024)

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Case Study: ChatGPT Memory / SpAIware (September 2024)

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

PoisonedRAG: Five Documents, 90% Success Rate

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Vector Database Vulnerabilities

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Three-Layer Defense Architecture for RAG

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

The Slack AI Attack: A Full Technical Walkthrough

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Competing Perspectives: Is RAG Really That Vulnerable?

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Three-Layer Defense Architecture for RAG (Expanded with Implementation)

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Key Takeaway

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Chapter 5. Jailbreaks: Social Engineering for Machines

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

What Jailbreaking Is (and Isn't)

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

The Taxonomy of Jailbreak Techniques

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Characteristics of Jailbreak Prompts

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Automated Jailbreak Discovery

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

The Consequences: Character.AI and Beyond

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

The GOAT Framework: How Automated Red Teaming Works

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

TAP: Tree of Attacks with Pruning

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Defense Against Jailbreaks: A Practical Guide

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

The Character.AI Case: A Cautionary Tale

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Key Takeaway

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Chapter 6. AI Agents and Tool Use: When Prompts Become Shells

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

From Chatbots to Agents

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

The Semantic Kernel RCE Vulnerabilities (May 2026)

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Documented Incidents

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

The Expanding Blast Radius

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Framework-Agnostic Vulnerabilities

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

CVE-2026-26030: A Step-by-Step Code Analysis

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

The Expanding Blast Radius: A Framework for Analysis

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

A Deeper Look: The Semantic Kernel RCE Vulnerabilities (CVE-2026-26030 and CVE-2026-25592)

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Framework-Agnostic Vulnerabilities: The Deeper Truth

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Key Takeaway

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Chapter 7. Defense-in-Depth Architecture: Layered Protections

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

The Three-Layer Model

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Guardrails: The Runtime Enforcement Layer

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

The Trade-Off: Security vs. Usability

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

The Trade-Off: Security vs. Usability (Expanded with Calibration Examples)

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Competing Perspectives: Is Defense-in-Depth Overkill?

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Key Takeaway

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Chapter 8: Instruction Hierarchy and Isolation Patterns

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

The Core Principle: Instruction Hierarchy

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

The Six Design Patterns for Injection-Resistant Agents

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Architectural Isolation Patterns

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Case Study: Applying Patterns to a Customer Service Chatbot

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Pattern Selection: A Decision Framework

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Case Study: Applying Patterns to a Healthcare Data Analysis Platform

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Competing Perspectives: Are Design Patterns Enough?

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Key Takeaway

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Chapter 9. Enterprise Security Stacks: FIDES, LlamaFirewall, and Beyond

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Microsoft FIDES: Deterministic Defense Through Information Flow Control

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Meta LlamaFirewall: Open-Source Guardrail System

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Comparing Enterprise Options

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

OWASP Top 10 for LLM Applications 2025

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Choosing the Right Stack

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

FIDES Implementation: A Practical Walkthrough

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

LlamaFirewall Deployment Guide

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

LlamaFirewall's Agent Alignment Check: A Deeper Dive

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

The TaskTracker Alternative: Detection Through Internal Activations

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Comparing Enterprise Options (Expanded)

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Competing Perspectives. Deterministic vs. Probabilistic: The Real Trade-Off

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Key Takeaway

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Chapter 10. Red-Teaming AI Systems: Practical Security Testing

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Why Red-Teaming Is Different

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Setting Up Automated Red Teams with Promptfoo

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

PyRIT and Garak: Alternative Frameworks

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Building a Red-Team Playbook

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

The LLMail-Inject Challenge: Lessons from 370,000+ Test Prompts

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

CI/CD Integration

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

PyRIT: Multi-Turn Attack Walkthrough

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Garak: Comprehensive Vulnerability Scanning

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Garak: Comprehensive Vulnerability Scanning

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

CI/CD Integration: Making Red-Teaming Continuous

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

The LLMail-Inject Challenge: Lessons from 370,000+ Test Prompts (Expanded)

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Practical Red-Team Exercise: A Complete Walkthrough

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Chapter 11. Production Deployment: Monitoring, Incident Response, and Compliance

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Runtime Monitoring and Anomaly Detection

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Data Governance in AI Systems

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Incident Response Playbooks

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Compliance Frameworks

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Runtime Monitoring: A Practical Implementation Guide

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Incident Response Playbook (Expanded Template)

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Compliance Walkthrough: NIST AI RMF

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Compliance Walkthrough: EU AI Act

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

The Real-World Cost of Prompt Injection Incidents

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Key Takeaway

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Chapter 12. The Future of AI Security: Emerging Threats and Defenses

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Multimodal Injection

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Supply Chain Attacks in AI

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Adaptive Attacks

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Deterministic vs. Probabilistic Defenses

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

The Role of Regulation and Standardization

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

The Adaptive Arms Race: Attack and Defense Evolution

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Multimodal Injection: The Next Frontier

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

The Supply Chain Attack Surface: Beyond Code

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

The Adaptive Arms Race: Attack and Defense Evolution

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

The Regulatory Landscape: A Practical Guide

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Supply Chain Security: Beyond Code

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

The Regulatory Landscape: A Practical Guide

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Final Synthesis: Building a Security-First Culture

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Conclusion: Trust as a Design Principle

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

Glossary

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.

References / Endnotes

This content is not available in the sample book. The book can be purchased on Leanpub at <https://leanpub.com/promptengineeringsecuritybuildinginjection-resistantaisystems>.