AI SECURITY GOVERNANCE AND ASSURANCE

A COMPREHENSIVE GUIDE TO GOVERNING, SECURING, AND ASSURING ARTIFICIAL INTELLIGENCE SYSTEMS

GOVERN
Establish strategy, policies, and accountability

SECURE
Protect AI systems across the entire lifecycle

ASSURE
Implement assurance activities and manage risk

COMPLY
Demonstrate compliance with confidence

NIST AI RMF

ISO/IEC 42001

EU AI ACT

OWASP LLM SECURITY GUIDANCE

INDUSTRY BEST PRACTICES

NICHOLAS M. DONOVAN

AI Security Governance and Assurance

A Comprehensive Guide to Governing, Securing, and Assuring Artificial Intelligence Systems

Steve Publications

This book is available at
https://leanpub.com/aisecuritygovernanceandassurance

This version was published on 2026-09-04

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.

Contents

Introduction **1**

Chapter 1: Why AI Security Governance and Assurance Is Different . . **4**
Why AI Security Is Different 4
The Stakes: Real-World Incidents and Failures 5
From Cybersecurity to AI Security: Evolution of the Discipline 7
What This Book Covers and How to Use It 8
The Integrated Approach: Governance, Risk, and Assurance 9

Chapter 2: Foundations of AI Security and Risk **11**
Core Terminology and Concepts 11
Understanding AI Risk: Categories and Characteristics 11
AI-Specific Failure Modes 11
The Security-AI-Privacy Triangle 11
Trust, Safety, and Reliability: Beyond Traditional Security 11
Accountability and Responsibility in AI Systems 11

Chapter 3: The AI Threat Landscape and Attack Surfaces **13**
Prompt Injection and Indirect Prompt Injection 13
Data Poisoning and Model Poisoning 13
Adversarial Examples and Evasion Attacks 13
Model Extraction, Inversion, and Membership Inference 13
Insecure AI Supply Chains and Model Dependencies 13
Excessive Agency and Tool Abuse in Agentic AI 13
Unauthorized Access, API Abuse, and Training Data Risks 14
AI-Generated Misinformation and Emergent Threats 14

Chapter 4: AI Security Governance Frameworks and Standards **15**
NIST AI Risk Management Framework (AI RMF) 15
NIST Cybersecurity Framework (CSF) and AI Security 15
ISO/IEC 42001: AI Management System 15

ISO/IEC 27001 and AI Security Controls 15
ISO/IEC 23894 and ISO/IEC 38507 15
OWASP LLM Top 10 and Generative AI Security 15
SOC 2 and AI Services 16
EU AI Act and Regulatory Landscape 16
Framework Selection and Integration Strategy 16

Chapter 5: AI Security Architecture 17
AI Security Architecture Principles 17
Governance and Organizational Architecture 17
Data Security Architecture for AI 17
Model and ML Infrastructure Security 17
Application and API Security for AI Systems 17
Identity, Access, and Secrets Management 17
Network Security, Isolation, and Sandboxing 18
CI/CD, MLOps, and LLMOps Security 18
Runtime Protection and Observability 18
Third-Party and Cloud AI Services 18

Chapter 6: Securing the AI Lifecycle 19
Strategy and Requirements Phase Security 19
Data Acquisition and Preparation Security 19
Model Development and Training Security 19
Testing, Validation, and Pre-Deployment Assurance 19
Deployment and Integration Security 19
Operational Monitoring and Maintenance 19
Model Drift, Retraining, and Change Management 20
Retirement, Disposal, and Decommissioning 20

Chapter 7: AI Governance Programs and Organizational Models 21
Establishing an AI Governance Program 21
AI Security Program Design and Structure 21
Roles, Responsibilities, and Accountability 21
AI Governance Committees and Decision Bodies 21
Policies, Standards, and Operating Procedures 21
AI Classification Schemes and Risk Tiers 21
Risk Ownership and Acceptance Processes 22
Organizational Models: Centralized, Federated, and Hybrid 22

Chapter 8: AI Risk Management and Control Frameworks 23

Model Risk Management: Principles and Practice 23
AI Control Frameworks and Control Objectives 23
AI and Model Inventories 23
Risk Registers and Risk Assessment Methods 23
Security Baselines and Control Matrices 23
Assurance Criteria and Evaluation Standards 23
Managing Risk Appetite and Risk Tolerance 24
Exception Management and Risk Acceptance 24

Chapter 9: AI Assurance, Testing, and Validation 25
Principles of AI Assurance 25
AI Red-Team Programs and Methodology 25
Penetration Testing for AI Systems 25
Adversarial Testing and Robustness Evaluation 25
Model Evaluation and Validation Frameworks 25
Independent Review and Audit Programs 25
Bias, Fairness, and Safety Testing 26
Assurance Evidence and Documentation 26

Chapter 10: AI Security Operations and Continuous Monitoring 27
AI Security Operations Centers and Capabilities 27
Runtime Monitoring and Anomaly Detection 27
Prompt and Response Logging 27
Alerting, Detection, and Incident Response 27
AI Incident Classification and Playbooks 27
Business Continuity and Resilience for AI Systems 27
Third-Party AI Service Monitoring 28

Chapter 11: Integrating AI Security with Enterprise Functions 29
Integration with Enterprise Cybersecurity 29
Integration with Privacy and Data Governance 29
Integration with GRC and Risk Management 29
Integration with Internal Audit and External Assurance 29
Integration with Legal and Regulatory Functions 29
Integration with Procurement and Vendor Management 29
Integration with Software Development and DevSecOps 30
Integration with Business Continuity and IT Operations 30

Chapter 12: Advanced Topics - Agentic AI, Cloud AI, and Emerging Risks 31
Security Challenges of Agentic and Autonomous AI 31

Cloud AI Services: Security and Shared Responsibility 31
Foundation Model Security and Governance 31
RAG Architecture Security 31
Multimodal AI Security Considerations 31
Frontier AI and Advanced Model Risks 31
Open Source AI Security 32
Emerging Threats and Future Directions 32

Chapter 13: Measuring Maturity, Compliance, and Continuous Improvement 33
AI Security Maturity Models 33
KPIs, KRIs, and Performance Metrics 33
Demonstrating Regulatory Compliance 33
Audit Programs and Evidence Requirements 33
Traceability and Documentation 33
Lessons Learned and Continuous Improvement 33
Board Reporting and Executive Communication 34
Building an AI Security Culture 34

Chapter 14: Conclusion - The Future of AI Security Governance 35
Key Takeaways: Principles That Endure 35
The Evolving Threat Landscape 35
Technology Trends and Their Security Implications 35
The Role of Standards and Regulation 35
A Call for Integrated, Proactive AI Security 35
Where to Go from Here 35

References 37

Introduction

In February 2026, Google detected and blocked a coordinated campaign that sent over 100,000 suspicious prompts to Gemini in a systematic attempt to replicate its proprietary reasoning capabilities through model extraction [19]. Around the same time, attackers demonstrated that autonomous AI agents could be weaponized at scale without substantial human intervention [16]. These events were not theoretical exercises. They were real attacks against production AI systems by actors who understood that AI security was still catching up to AI capability.

The rise of artificial intelligence has created a new class of assets that existing cybersecurity frameworks were never designed to protect. AI systems are not just software to be hardened against exploitation. They are dynamic, probabilistic systems that learn from data, generate outputs that can influence real-world decisions, interact with external tools and APIs, and evolve over time. They introduce threats that have no equivalent in traditional IT: prompt injection attacks that subvert model behavior, data poisoning that corrupts learning, adversarial examples that evade detection, model extraction that steals intellectual property, and agentic AI that can take autonomous actions across interconnected systems.

This book exists because organizations are deploying AI faster than they are securing it. The gap between AI adoption and AI security maturity is the single largest source of risk facing enterprises today. Closing that gap requires more than a security checklist or a set of point solutions. It requires an integrated approach that treats AI security as a distinct discipline with its own threat landscape, architectural requirements, governance structures, assurance methodologies, and operational practices, while connecting it seamlessly to existing cybersecurity, privacy, risk management, compliance, and business continuity programs.

Effective AI security governance and assurance is not about slowing down innovation. It is about making innovation sustainable. It is about ensuring that the organizations that deploy AI are not the first to be exploited by it. It is about building trust with customers, regulators, and stakeholders who will not accept AI systems that are convenient but unsafe.

This book addresses the following core questions:

- What does it mean to govern, secure, and assure an AI system?
- How do AI-specific threats differ from traditional cybersecurity risks, and what does that mean for architecture and controls?
- Which governance and security frameworks apply to AI, how do they relate to each other, and how should they be combined in practice?
- What does secure AI lifecycle management look like, from strategy through retirement?
- How do you build an AI security program that scales, integrates with existing enterprise functions, and stands up to audit and regulation?
- How do you test, validate, and continuously monitor AI systems to ensure they remain secure and trustworthy over time?
- How do you measure your AI security maturity, demonstrate compliance, and continuously improve?

The book is organized to take you from foundations to implementation to continuous improvement. Chapter 2 establishes the conceptual foundations: core terminology, the nature of AI risk, and the security-AI-privacy triangle that frames everything that follows. Chapter 3 provides exhaustive coverage of the AI threat landscape: prompt injection, data poisoning, adversarial attacks, model extraction, supply chain risks, agentic AI vulnerabilities, and emerging threats. Chapters 4 and 5 address governance and architecture, covering the full ecosystem of frameworks and standards and detailing how to design a comprehensive AI security architecture. Chapters 6 through 9 cover the lifecycle, the program, the controls, and the assurance: securing each stage of the AI system lifecycle, establishing governance and security programs, building control frameworks and risk management processes, and implementing testing and validation functions. Chapters 10 through 13 address operations, integration, advanced topics, and measurement: security operations and continuous monitoring, integrating AI security with existing enterprise functions, advanced topics including agentic AI and cloud AI services, and measuring maturity and demonstrating compliance. Chapter 14 synthesizes the book and looks forward.

The audience for this book is broad. Security architects will find detailed architecture patterns and control implementations. CISOs and security leaders will find strategies for building and scaling AI security programs. AI and ML engineers will find practical guidance for securing models, data, and

pipelines. Governance, risk, and compliance professionals will find authoritative coverage of frameworks and regulatory requirements. Auditors will find assurance criteria and evidence requirements. Enterprise architects will find integration patterns that connect AI security to the broader technology landscape. DevSecOps and MLOps practitioners will find operational procedures and configurations.

The guidance in this book is grounded in three principles:

- Risk-based. Not all AI systems carry the same risk. Effective AI security focuses effort where it matters most, using risk classification and tiering to drive proportional governance.
- Defensible. Decisions about AI security must withstand scrutiny from auditors, regulators, boards, and incident investigators. Documentation, traceability, and clear rationale are non-negotiable.
- Maintainable. AI security programs that cannot be sustained in production are failures in waiting. Every control, process, and architecture pattern presented here considers operational impact and long-term maintainability.

No single standard, framework, or tool can secure your AI systems. The organizations that succeed will be those that build integrated AI security programs grounded in sound principles, aligned with relevant frameworks, tailored to their specific risk profile, and executed with operational rigor. This book is designed to help you do exactly that.

Chapter 1: Why AI Security Governance and Assurance Is Different

Artificial intelligence is transforming enterprise operations at unprecedented speed. AI systems are being deployed to accelerate customer service, enhance cybersecurity, automate operations, power search and recommendation engines, generate code, analyze documents, support medical decisions, optimize supply chains, and automate financial processes. The speed and breadth of this adoption far outpaces traditional technology transitions. But speed alone does not explain the urgency of AI security. The nature of AI risk is fundamentally different from the risks that established cybersecurity frameworks were designed to address.

Why AI Security Is Different

Traditional cybersecurity is built on well-understood principles: protect the confidentiality, integrity, and availability of information and systems through access controls, encryption, network segmentation, monitoring, and incident response. The attack surface is known. The attacker models are documented. The controls are standardized. The metrics are measurable.

AI introduces a fundamentally different class of risk for four reasons:

First, AI systems are probabilistic rather than deterministic. A traditional application takes an input, applies known logic, and produces a predictable output. If you test it thoroughly, you know how it will behave. AI models, especially generative models and deep neural networks, produce outputs based on statistical patterns learned from training data. The relationship between input and output is not fixed by explicit logic. It emerges from patterns that may not be fully understood even by the engineers who built the model. This introduces uncertainty into security testing and assurance. You cannot simply verify that an AI system is safe by checking its code. You must test its behavior across a vast space of possible inputs and observe its responses over time.

Second, AI systems have novel failure modes that do not map cleanly to traditional vulnerability categories. A buffer overflow is a known issue with well-understood mitigations. A prompt injection attack exploits the fundamental way that language models process instructions. An adversarial example fools a neural network by adding imperceptible perturbations to its input. Data poisoning corrupts a model by contaminating its training data. Model extraction reconstructs a proprietary model by querying it. These attacks target the learning and reasoning processes themselves, not just the infrastructure that hosts them [4, 6, 10, 19]. Traditional security tools and tests are blind to many of these threats.

Third, AI systems are deeply entangled with data, infrastructure, and human decision-making. The security of an AI system depends not only on the model itself but on the quality and integrity of its training data, the security of its training and inference infrastructure, the controls on its APIs and integrations, the governance of its outputs, and the human processes that oversee its use. A compromise at any layer can cascade through the entire system. A poisoned dataset produces a flawed model. A compromised API key enables model extraction. An insecure integration allows prompt injection through indirect channels. A lack of human oversight allows harmful outputs to reach production [5, 14].

Fourth, AI systems create new vectors for harm. A compromised traditional system typically results in data breaches, service disruption, or unauthorized access. A compromised AI system can do all of those and more. It can generate convincing misinformation at scale, impersonate authorized users, manipulate financial markets, subvert safety systems, exfiltrate training data, make biased or discriminatory decisions, take autonomous actions that cause real-world harm, and create a feedback loop that amplifies its own flaws over time [5, 6, 8, 19]. The impact of AI security failures extends beyond IT to safety, compliance, ethics, and reputation in ways that traditional cybersecurity incidents rarely match.

The Stakes: Real-World Incidents and Failures

AI security is not a theoretical concern. Real incidents are already occurring at increasing frequency and severity. The following examples illustrate the range and seriousness of AI-specific security failures:

In February 2026, Google detected and blocked a coordinated campaign that sent over 100,000 suspicious prompts to Gemini in a systematic attempt to extract its proprietary reasoning capabilities through model extraction. OpenAI separately informed U.S. lawmakers that a competing AI company had used obfuscated methods to extract results from American AI models [19]. These incidents demonstrate that model extraction is already being attempted at scale by well-resourced actors.

NIST's own red-team research found that novel attack techniques targeting AI agents achieved an 81 percent task-hijacking success rate compared to 11 percent for the strongest known baseline attacks, demonstrating that agent-specific offensive research can dramatically outperform defenses calibrated to known attack taxonomies [2]. This gap between attack and defense capability is widening.

In 2024 and 2025, researchers and security companies documented prompt injection vulnerabilities in enterprise AI deployments including Slack AI, Microsoft 365 Copilot, Cursor, GitHub MCP, and AI coding assistants, showing that the threat shifted from chatbot trick to enterprise risk within a single year [7]. Prompt injection has evolved into one of the most common and impactful AI security threats.

A JFrog report analyzing 18.2 billion artifacts and surveying 1,508 security professionals found that while 97 percent of organizations claim some form of AI model governance, 53 percent self-host models sourced from public registries where 495 malicious AI models were identified [7]. The AI supply chain is vast, opaque, and largely unaudited.

Analysis of 10,000 open-source AI and ML repositories revealed that 70 percent have critical or high-severity vulnerabilities in their build workflows, making them prone to attacks including code injection, credential theft, and repository takeover via malicious pull requests [9]. The foundation of the AI ecosystem is riddled with exploitable weaknesses.

AI security failures are not limited to external attacks. Model drift has undermined fraud detection, intrusion detection, and behavioral analysis systems when live data diverged from training distributions without timely detection or remediation [4, 5]. Overreliance on AI recommendations without adequate validation has led to erroneous decisions in financial, legal, and healthcare contexts. AI systems deployed without proper access controls have exposed sensitive training data and internal documents through simple queries. These

are not edge cases. They are the baseline risks of deploying AI without comprehensive security and governance.

From Cybersecurity to AI Security: Evolution of the Discipline

Cybersecurity as a discipline matured over decades through cycles of attack, defense, standardization, and improvement. The OWASP Top 10 for web applications has been refined through multiple iterations. NIST SP 800-53 provides hundreds of vetted security controls. ISO/IEC 27001 has defined the baseline for information security management for over two decades. These frameworks are robust because they were built on extensive real-world experience.

AI security is still building that foundation. It stands on three pillars:

The first pillar is traditional cybersecurity applied to AI infrastructure. The servers, databases, networks, APIs, identity systems, and software that underpin AI systems must be secured using well-established practices. Access controls, encryption, logging, monitoring, vulnerability management, secure configuration, and incident response are all necessary. Standards like ISO/IEC 27001 and the NIST Cybersecurity Framework remain fully applicable and essential [4, 11].

The second pillar is AI-specific security that addresses threats unique to machine learning and generative AI systems. These include prompt injection, adversarial attacks, data poisoning, model extraction, membership inference, and the security challenges of autonomous agents. Frameworks like the OWASP Top 10 for LLM Applications, MITRE ATLAS, and emerging guidance from NIST and others provide growing guidance in this domain [2, 6, 10, 19].

The third pillar is AI governance and assurance that encompasses not just security but also trustworthiness, safety, ethics, fairness, accountability, transparency, privacy, compliance, and regulatory obligations. Frameworks like the NIST AI Risk Management Framework, ISO/IEC 42001, ISO/IEC 23894, ISO/IEC 38507, and regulations like the EU AI Act address this broader domain [1, 3, 5, 13].

These three pillars are not separate. They must be integrated. AI-specific security that sits on insecure infrastructure will fail. Traditional cybersecurity that ignores AI-specific threats is incomplete. Security that operates without

governance and assurance will lack accountability and sustainability. The organizations that succeed will be those that build an integrated approach that covers all three pillars and connects them to existing enterprise risk management, compliance, audit, and business continuity functions.

What This Book Covers and How to Use It

This book is designed to be both a comprehensive reference and a practical guide. It covers the full spectrum of AI security governance and assurance, from foundational concepts through detailed architecture, controls, operations, and measurement. The content is organized so that readers can use it in different ways depending on their role and needs.

For readers building an AI security program from scratch, the book should be read in sequence. The early chapters establish the conceptual foundation and threat landscape. The middle chapters cover governance frameworks, architecture, program design, and controls. The later chapters address operations, integration, advanced topics, and measurement. Following the sequence will provide a coherent understanding of how all the pieces fit together.

For readers with existing responsibilities who need specific guidance, each chapter is self-contained enough to be used as a reference. Security architects can jump to Chapter 5 for architecture patterns and Chapter 6 for lifecycle controls. CISOs and program leaders can focus on Chapters 4, 7, and 8 for frameworks, governance programs, and risk management. AI engineers can use Chapter 3 for threat understanding, Chapter 6 for lifecycle security, and Chapter 9 for testing guidance. GRC and compliance teams can use Chapters 4, 7, and 13 for regulatory requirements, governance programs, and compliance evidence.

Throughout the book, guidance is categorized by type:

- Mandatory requirements refer to obligations imposed by regulations or standards that are legally or contractually binding. These must be implemented.
- Strong recommendations refer to practices that are essential for effective AI security and should be implemented unless there is a compelling, documented justification not to.

- Recommended practices refer to good practices that add value but may not be feasible or necessary in all contexts. They should be implemented where practical.

The book also distinguishes between:

- General guidance that applies to AI systems broadly.
- Environment-specific guidance that applies to particular deployment models such as cloud AI services, on-premises models, foundation models, RAG systems, or autonomous agents.
- Risk-tier-specific guidance that applies differently based on the risk classification of the AI system.

Every chapter includes concrete examples, configuration artifacts where applicable, and references to frameworks, standards, and primary sources. Technical and regulatory claims are carefully sourced. Where evidence is contested or evolving, the book notes the uncertainty and presents competing perspectives fairly.

The Integrated Approach: Governance, Risk, and Assurance

The title of this book contains three terms: governance, security, and assurance. Each represents a distinct dimension of AI risk management, and all three are necessary.

AI governance is the system of decision-making, oversight, accountability, policies, and processes that determines how an organization develops, deploys, and uses AI. It defines who is responsible, what decisions require approval, how risk is assessed, what standards must be followed, and how compliance is maintained. Governance without security is theater. Governance without assurance is self-certification. But security and assurance without governance lack authority, direction, and sustainability.

AI security is the set of controls, technologies, architectures, and operational practices that protect AI systems and their associated data, infrastructure, and processes from unauthorized access, misuse, disruption, and exploitation. It addresses both traditional cyber threats and AI-specific threats.

Security without governance lacks mandate and scope. Security without assurance lacks proof of effectiveness.

AI assurance is the systematic evaluation, testing, monitoring, and verification that AI systems meet their security, safety, reliability, fairness, and compliance requirements throughout their lifecycle. It includes red teaming, penetration testing, adversarial testing, model evaluation, independent review, audit, and continuous monitoring. Assurance without governance lacks mandate. Assurance without security has nothing to validate.

Together, governance, security, and assurance form an integrated system. Governance sets the direction and accountability. Security implements the controls. Assurance provides the evidence. Each enables and reinforces the others. This book treats all three as equally essential and covers how they interconnect at every level.

The journey ahead is demanding but achievable. Organizations that invest now in integrated AI security governance and assurance will not only protect themselves from emerging threats but will also build the trust that will differentiate them in an AI-driven marketplace. The organizations that delay will find themselves scrambling to catch up while their systems, reputation, and compliance posture are on the line.

This book is designed to help you do the work right, the first time, with rigor, practicality, and depth. Let us begin.

Chapter 2: Foundations of AI Security and Risk

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Core Terminology and Concepts

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Understanding AI Risk: Categories and Characteristics

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

AI-Specific Failure Modes

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

The Security-AI-Privacy Triangle

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Trust, Safety, and Reliability: Beyond Traditional Security

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Accountability and Responsibility in AI Systems

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Chapter 3: The AI Threat Landscape and Attack Surfaces

Prompt Injection and Indirect Prompt Injection

Data Poisoning and Model Poisoning

Adversarial Examples and Evasion Attacks

Model Extraction, Inversion, and Membership Inference

Insecure AI Supply Chains and Model Dependencies

Excessive Agency and Tool Abuse in Agentic AI

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Unauthorized Access, API Abuse, and Training Data Risks

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

AI-Generated Misinformation and Emergent Threats

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Chapter 4: AI Security Governance Frameworks and Standards

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

NIST AI Risk Management Framework (AI RMF)

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

NIST Cybersecurity Framework (CSF) and AI Security

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

ISO/IEC 42001: AI Management System

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

ISO/IEC 27001 and AI Security Controls

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

ISO/IEC 23894 and ISO/IEC 38507

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

OWASP LLM Top 10 and Generative AI Security

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

SOC 2 and AI Services

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

EU AI Act and Regulatory Landscape

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Framework Selection and Integration Strategy

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Chapter 5: AI Security Architecture

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

AI Security Architecture Principles

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Governance and Organizational Architecture

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Data Security Architecture for AI

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Model and ML Infrastructure Security

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Application and API Security for AI Systems

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Identity, Access, and Secrets Management

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Network Security, Isolation, and Sandboxing

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

CI/CD, MLOps, and LLMOps Security

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Runtime Protection and Observability

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Third-Party and Cloud AI Services

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Chapter 6: Securing the AI Lifecycle

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Strategy and Requirements Phase Security

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Data Acquisition and Preparation Security

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Model Development and Training Security

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Testing, Validation, and Pre-Deployment Assurance

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Deployment and Integration Security

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Operational Monitoring and Maintenance

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Model Drift, Retraining, and Change Management

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Retirement, Disposal, and Decommissioning

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Chapter 7: AI Governance Programs and Organizational Models

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Establishing an AI Governance Program

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

AI Security Program Design and Structure

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Roles, Responsibilities, and Accountability

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

AI Governance Committees and Decision Bodies

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Policies, Standards, and Operating Procedures

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

AI Classification Schemes and Risk Tiers

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Risk Ownership and Acceptance Processes

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Organizational Models: Centralized, Federated, and Hybrid

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Chapter 8: AI Risk Management and Control Frameworks

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Model Risk Management: Principles and Practice

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

AI Control Frameworks and Control Objectives

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

AI and Model Inventories

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Risk Registers and Risk Assessment Methods

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Security Baselines and Control Matrices

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Assurance Criteria and Evaluation Standards

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Managing Risk Appetite and Risk Tolerance

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Exception Management and Risk Acceptance

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Chapter 9: AI Assurance, Testing, and Validation

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Principles of AI Assurance

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

AI Red-Team Programs and Methodology

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Penetration Testing for AI Systems

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Adversarial Testing and Robustness Evaluation

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Model Evaluation and Validation Frameworks

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Independent Review and Audit Programs

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Bias, Fairness, and Safety Testing

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Assurance Evidence and Documentation

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Chapter 10: AI Security Operations and Continuous Monitoring

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

AI Security Operations Centers and Capabilities

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

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/aisecuritygovernanceandassurance.

Prompt and Response Logging

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Alerting, Detection, and Incident Response

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

AI Incident Classification and Playbooks

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Business Continuity and Resilience for AI Systems

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Third-Party AI Service Monitoring

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Chapter 11: Integrating AI Security with Enterprise Functions

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Integration with Enterprise Cybersecurity

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Integration with Privacy and Data Governance

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Integration with GRC and Risk Management

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Integration with Internal Audit and External Assurance

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Integration with Legal and Regulatory Functions

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Integration with Procurement and Vendor Management

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Integration with Software Development and DevSecOps

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Integration with Business Continuity and IT Operations

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Chapter 12: Advanced Topics - Agentic AI, Cloud AI, and Emerging Risks

Security Challenges of Agentic and Autonomous AI

Cloud AI Services: Security and Shared Responsibility

Foundation Model Security and Governance

RAG Architecture Security

Multimodal AI Security Considerations

Frontier AI and Advanced Model Risks

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Open Source AI Security

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Emerging Threats and Future Directions

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Chapter 13: Measuring Maturity, Compliance, and Continuous Improvement

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

AI Security Maturity Models

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

KPIs, KRIs, and Performance Metrics

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Demonstrating Regulatory Compliance

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Audit Programs and Evidence Requirements

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Traceability and Documentation

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Lessons Learned and Continuous Improvement

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Board Reporting and Executive Communication

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Building an AI Security Culture

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Chapter 14: Conclusion - The Future of AI Security Governance

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Key Takeaways: Principles That Endure

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

The Evolving Threat Landscape

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Technology Trends and Their Security Implications

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

The Role of Standards and Regulation

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

A Call for Integrated, Proactive AI Security

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

Where to Go from Here

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.

References

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/aisecuritygovernanceandassurance.