DEMO

First chapter only

The CISO's Guide to Agentic Commerce Risk

Risk Assessment Framework for AI Agents in Financial Transactions

pragma.vision

TRUSTAUTHORITY.AI

The CISO's Guide to Agentic Commerce Risk

Trademark Notice

No Affiliation

This book is an independent publication. It is not authorized, sponsored, or endorsed by any of the companies or organizations whose products or services are mentioned herein.

No Professional Advice

The information in this book is provided for educational purposes only. It does not constitute legal, financial, investment, tax, or other professional advice. Readers should consult qualified professionals for guidance specific to their situation.

Code Examples

Code examples in this book are provided for illustration only. They may not be suitable for production use without additional validation, error handling, and security review.

Published by Pragma Vision LLC

First edition, 2026.

Contents

1

When AI Agents Handle Money

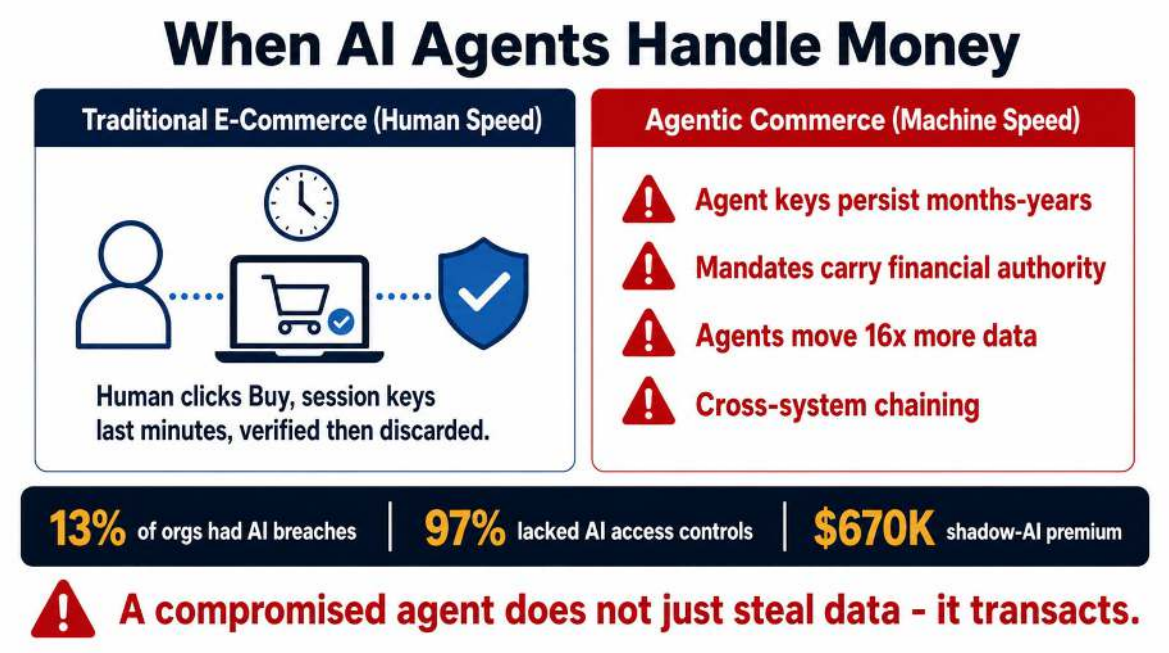

Figure 1. A comparison of traditional human-speed e-commerce (session keys lasting minutes) against machine-speed agentic commerce (keys persisting months to years, moving 16x more data)—where 13% of orgs had AI breaches, 97% lacked AI access controls, and shadow AI added $670K to breach cost

1.1 The New Attack Surface You Were Not Briefed On

Your organization already has AI agents in production. Whether your security team deployed them, approved them, or even knows about them, autonomous software agents are making decisions, accessing data, and—increasingly—spending money on behalf of your employees and your customers. The question is no longer whether to allow AI agents into your financial operations. The question is how to manage the risk they have already introduced.

The numbers are stark. According to IBM's 2025 Cost of a Data Breach Report, 13% of organizations reported breaches involving AI models or applications—and of those, 97% lacked proper AI access controls. Shadow AI—unauthorized agent deployments outside IT governance—costs organizations an average of $670,000 more per breach

than standard incidents. The global average breach cost stands at $4.44 million, while US companies face an average of $10.22 million per incident.

97%

of organizations breached through AI models lacked proper AI access controls[1]

Meanwhile, agentic commerce is accelerating. Global digital commerce is projected to surpass $11 trillion by 2026, with AI agents processing an increasing share of those transactions. Agents now negotiate vendor contracts, approve purchase requisitions, process expense reports, and execute financial trades. Each of these operations introduces attack vectors that traditional security frameworks were never designed to address.

Key Insight

The agentic commerce threat is fundamentally different from traditional cybersecurity threats. Agents operate autonomously, maintain persistent memory, chain actions across multiple systems, and move 16x more data than human users. A compromised agent does not just steal data—it actively transacts, and it does so at machine speed with machine-scale access.

1.2 About Pragma.Vision

This book draws from the security architecture of Pragma.Vision, an AI-native commerce ecosystem model spanning a growing family of interconnected platforms. The ecosystem uses a three-layer protocol security model—identity verification (Visa TAP), user authorization (Google AP2), and payment execution (Stripe ACP)—all secured with quantum-safe hybrid cryptography (ML-DSA-65). The 21 attack vectors cataloged in Chapter 3 come from penetration testing and security hardening patterns

[1] IBM Security, *Cost of a Data Breach Report*, 2025.

across 306 test cases. The risk framework in Chapter 4 is derived from NIST AI RMF 1.0 principles adapted for financial transaction contexts.

1.3 What This Book Covers

This guide is structured for CISOs and security leaders who need to:

1. **Understand the threat model**: The three-layer protocol security model that separates identity, authorization, and execution concerns (Chapter 2).
2. **Catalog the attack surface**: 21 attack vectors specific to AI agents handling financial transactions, categorized by protocol layer (Chapter 3).
3. **Assess risk systematically**: A likelihood-by-impact risk matrix calibrated for agentic commerce, with quantified financial exposure estimates (Chapter 4).
4. **Communicate to the board**: A ready-to-present risk narrative that translates technical threats into business language (Chapter 5).
5. **Map to compliance frameworks**: SOC 2 Trust Services Criteria, GDPR data processing requirements, and PCI DSS 4.0 controls as they apply to AI agent operations (Chapters 6 and 7).
6. **Respond to incidents**: An agent-specific incident response playbook that accounts for autonomous behavior, persistent memory, and cross-system chaining (Chapter 8).
7. **Build a security program**: A phased roadmap for establishing agentic commerce governance within your organization (Chapter 9).

Every chapter ends with a concrete deliverable—a matrix, a checklist, a template, or a policy document—that you can adapt for your organization immediately.

DEMO

This is a free preview of the full edition.

Get the complete book at:

https://shop.pragma.vision