DEMO

First chapter only

Know Your Agent (KYA)

The Enterprise Guide to AI Agent Identity Management

pragma.vision

TRUSTAUTHORITY.AI

Know Your Agent (KYA)

Trademark Notice

Google, Google Pay, Google Cloud, and Android are trademarks of Google LLC. Stripe is a trademark of Stripe, Inc. Cloudflare and Cloudflare Workers are trademarks of Cloudflare, Inc. Supabase is a trademark of Supabase, Inc. OpenAI and ChatGPT are trademarks of OpenAI, Inc. Claude is a trademark of Anthropic, PBC. W3C is a trademark of the World Wide Web Consortium. Visa is a trademark of Visa International Service Association. OWASP is a trademark of the OWASP Foundation. Midjourney is a trademark of Midjourney, Inc. Canva is a trademark of Canva Pty Ltd. Etsy is a trademark of Etsy, Inc. Amazon is a trademark of Amazon.com, Inc. All other trademarks are the property of their respective owners.

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

Introduction: The Identity Crisis in AI

1.1 About Pragma.Vision

Pragma.Vision is an AI-native commerce ecosystem—a growing family of interconnected platforms sharing identity, payments, and infrastructure. From conversational commerce (wish.now) to an AI agent marketplace (phantoid.com) to developer tools (soft.house), every platform depends on a single foundational question: *Can you trust the agent acting on your behalf?*

This book exists because the answer, for most enterprises today, is no.

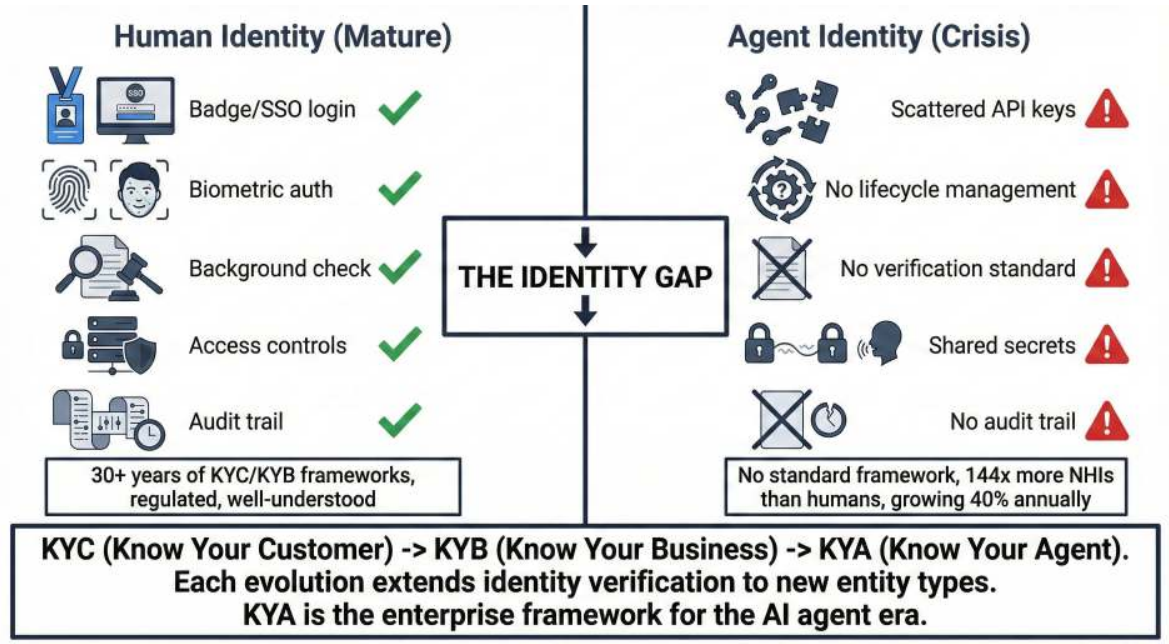

Figure 1. A side-by-side comparison contrasts mature human identity—badge/SSO login, biometrics, background checks, audit trails built over 30+ years—against the agent-identity gap of scattered API keys and no audit trail, with 144x more non-human than human identities, growing 40% annually

1.2 The Trust Deficit

For twenty years, enterprises invested in identity infrastructure for humans. Single sign-on, multi-factor authentication, identity governance, privileged access management—a mature ecosystem protecting human users. Then AI agents arrived, and organizations deployed them with little more than an API key and a prayer.

97%

of non-human identities have excessive privileges[1]

The consequences are already visible. In 2024 alone, compromised service accounts and leaked API keys caused breaches at the Internet Archive (33 million user records),

[1]Cloud Security Alliance and Astrix Security, "State of Non-Human Identity Security," 2025.

Dropbox Sign (elevated service account access), the US Treasury (BeyondTrust API key exploitation), Cisco DevHub (exposed tokens), and Schneider Electric (compromised Jira credentials exposing 40GB of data). Every one of these incidents involved a non-human identity that was poorly managed, never rotated, or granted far more access than it needed.

1.3 Why "Know Your Agent" Matters Now

Financial services has Know Your Customer (KYC). Business verification has Know Your Business (KYB). As AI agents begin conducting transactions, negotiating on behalf of users, and accessing sensitive enterprise systems, we need a parallel framework: **Know Your Agent (KYA)**.

KYA answers the questions that API keys cannot:

- **Who built this agent?** Not just a developer name, but a verified organizational identity with KYB-grade due diligence.
- **What is this agent authorized to do?** Not just API scopes, but cryptographically signed capability declarations.
- **How has this agent behaved?** Not just logs, but a verifiable behavioral history anchored to a persistent identity.
- **Is this agent still trustworthy?** Not a point-in-time check, but continuous monitoring with dynamic trust scoring.

Key Insight

KYC took decades to become mandatory in finance. KYA will become mandatory for AI commerce in years, not decades. The EU AI Act (fully enforceable August 2026) already requires identity traceability for high-risk AI systems. Or-

> ganizations that implement KYA now gain compliance readiness and competitive advantage simultaneously.

1.4 What You Will Learn

This book is organized as a progressive journey from understanding the problem to implementing the solution:

Chapters 2–3

The Problem. The scale of non-human identity proliferation and the OWASP-cataloged threats it creates.

Chapters 4–5

The Framework. The KYA model and its 4-level verification system, drawn from KYC/KYB parallels.

Chapters 6–7

The Implementation. Technical architecture using DIDs, verifiable credentials, trust scoring, and enterprise integration patterns.

Chapters 8–9

The Program. Compliance mapping across regulatory frameworks and a 90-day roadmap to production.

Every chapter ends with a concrete action you can take immediately. By the time you finish, you will have both the strategic understanding and the technical blueprint to build an agent identity program that satisfies your CISO, your auditors, and your board.

1.5 Who This Book Is For

- **CISOs and Security Architects** designing identity controls for AI agent deployments.
- **Enterprise Architects** integrating AI agents into existing IAM infrastructure.
- **Compliance Officers** mapping AI agent requirements to SOC 2, GDPR, PCI DSS, and the EU AI Act.
- **Platform Engineers** building agent marketplaces, orchestration systems, or multi-agent workflows.
- **Technical Founders** who need agent trust infrastructure from day one.

DEMO

This is a free preview of the full edition.

Get the complete book at:

https://shop.pragma.vision