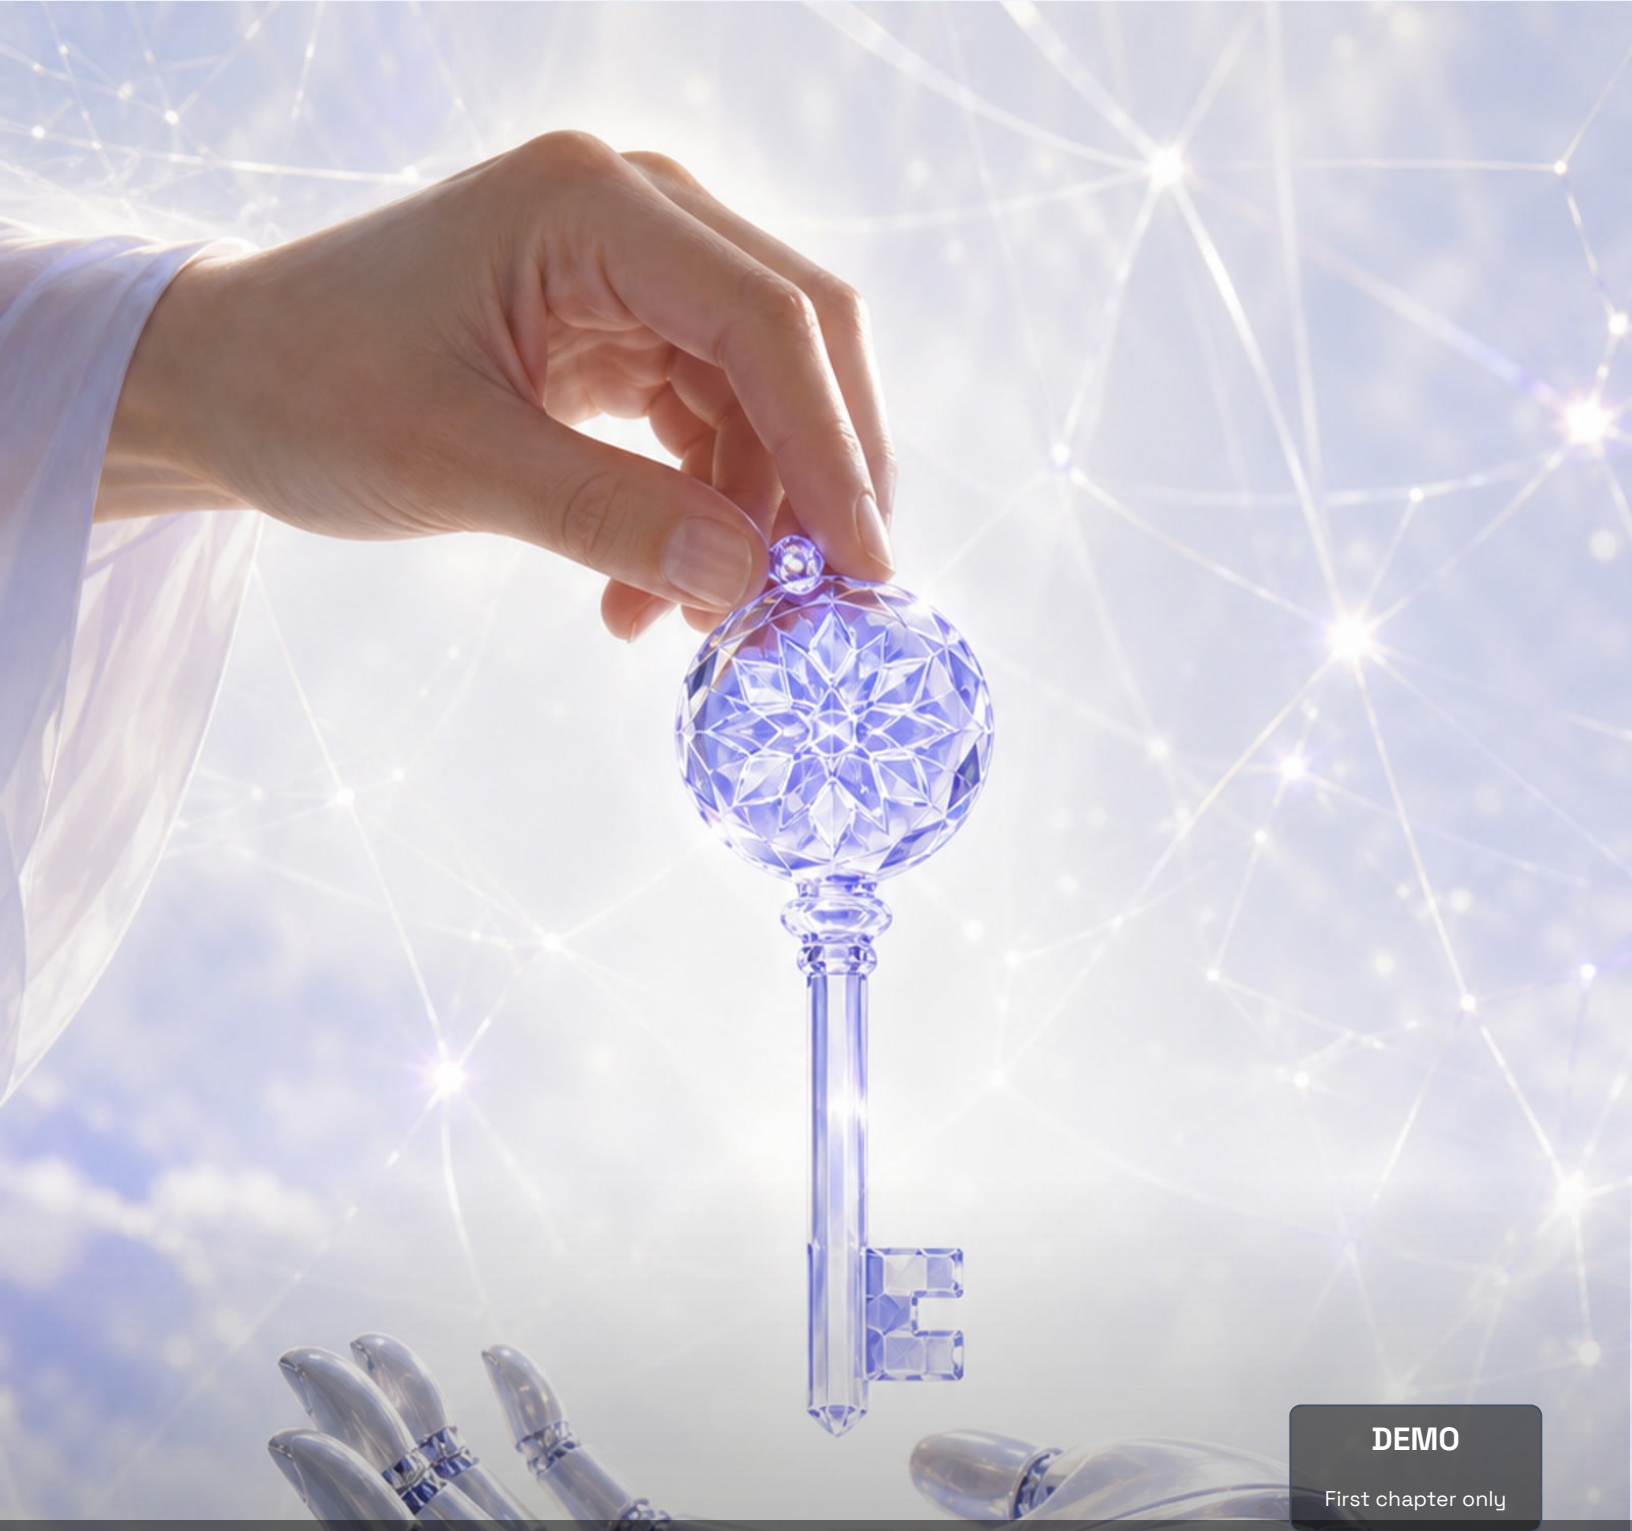

DEMO

First chapter only

W3C DIDs for AI Agents

Implementing did:tp: from Scratch

pragma.vision

TRUSTAUTHORITY.AI

W3C DIDs for AI Agents

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

Decentralized Identity for Non-Human Entities

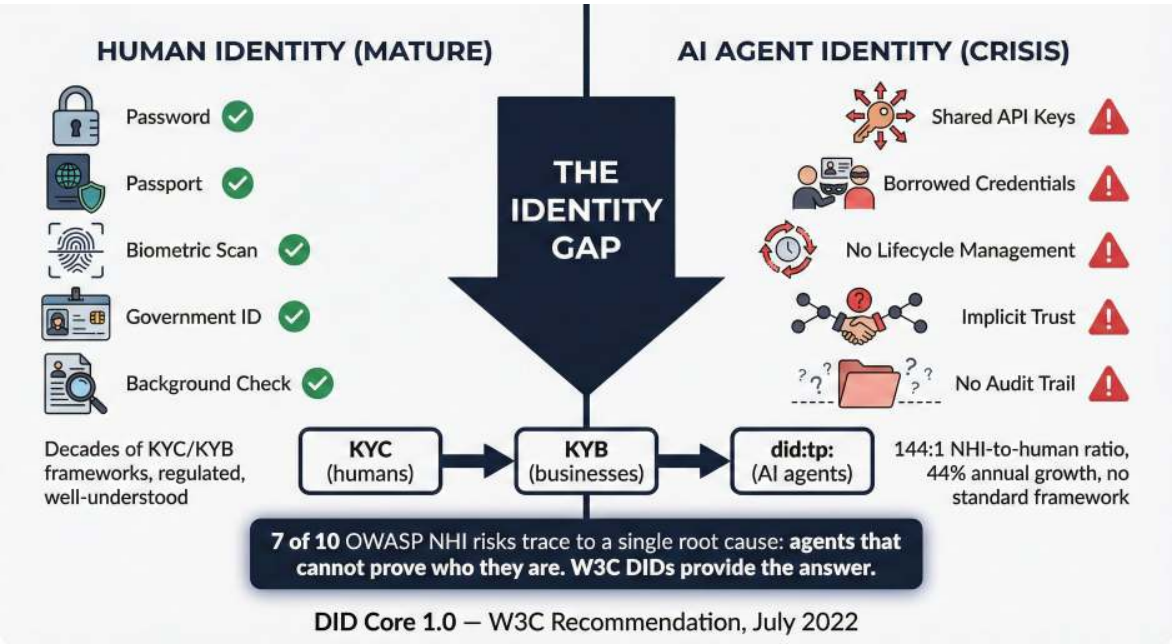

Figure 1. A side-by-side comparison contrasts mature human identity—passwords, passports, government ID, background checks—with an AI-agent crisis of shared API keys, borrowed credentials, and no audit trail at a 144:1 NHI-to-human ratio growing 44% annually, where 7 of 10 OWASP NHI risks trace to agents that cannot prove who they are

1.1 The Identity Crisis in AI

Every human on the internet has a way to prove who they are. Passwords, passports, biometric scans, government-issued IDs—the infrastructure for human identity verification has been developed over decades. AI agents have none of it. A language model that books your flights, a code review bot that pushes to your repositories, an invoice processor that moves your money—none of them carry verifiable identity. They operate on borrowed credentials, shared API keys, and implicit trust.

This is not a theoretical risk. Non-human identities now outnumber human users in enterprise environments by ratios as high as 144 to 1, and that ratio is growing at 44 percent year over year. By 2026, over 40 percent of enterprise workflows will involve autonomous agents—AI copilots, workflow bots, digital assistants—making identity verification for non-human actors not just useful but critical. The OWASP Non-Human

Identity Top 10 catalogs the ten most common vulnerabilities, and seven of them trace back to a single root cause: agents that cannot prove who they are.

144:1

ratio of non-human to human identities in enterprise environments[1]

The W3C Decentralized Identifiers (DIDs) specification provides the answer. Approved as a W3C Recommendation in July 2022, DID Core 1.0 defines a framework for globally unique, cryptographically verifiable identifiers that do not depend on any centralized registry. Originally designed for self-sovereign human identity, the specification is method-agnostic—meaning anyone can define a new DID method tailored to a specific domain. This book shows you how to build one for AI agents.

Key Insight

The W3C DID specification was designed to be extensible. The `did:tp:` method applies the same decentralized identity primitives that power human self-sovereign identity to the non-human entities that are rapidly becoming the majority participants in digital systems.

1.2 About Pragma.Vision

Pragma.Vision is an AI-native commerce ecosystem where multiple platforms work together to fulfill human needs through intelligent orchestration. The ecosystem is designed around a growing family of interconnected platforms—from conversational commerce (wish.now) to an AI agent marketplace (phantoid.com) to fair logistics (daily.delivery)—all sharing authentication, payments, cryptographic identity, and protocol infrastructure. Trust Authority AI (trustauthority.ai) is the foundation: the Certificate Authority pattern for AI agents, designed to issue verifiable credentials

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

using the `did:tp:` method specification described in this book. Every code example, schema definition, and implementation pattern in these pages is drawn from real implementation architecture.

1.3 What This Book Covers

This book is a technical guide to implementing the `did:tp:` DID method from scratch. It is written for developers and architects who need to give AI agents verifiable, decentralized identities that comply with W3C standards and integrate with the verifiable credentials ecosystem.

1. **DID Core 1.0 Fundamentals**: The specification primitives—DID documents, verification methods, service endpoints, and the data model that makes decentralized identity work.
2. **The did:tp: Method Specification**: Four namespaces for the agentic economy—agent, operator, service, and device—each with distinct resolution rules and lifecycle semantics.
3. **Key Management**: Ed25519 key generation, ML-DSA-65 post-quantum keys, key rotation ceremonies, and recovery procedures.
4. **DID Resolution**: How verifiers resolve `did:tp:` identifiers to DID documents, including caching, fallback, and offline verification strategies.
5. **Verifiable Credentials**: Issuing W3C Verifiable Credentials anchored to `did:tp:` identifiers, presenting proofs, and building trust chains.
6. **Implementation**: Complete TypeScript code for generating, storing, resolving, and verifying DIDs using Supabase and Cloudflare Workers.
7. **Testing and Deployment**: Test vectors, compliance validation, and production deployment patterns.

Pro Tip

Every code listing in this book is production-grade TypeScript. The reference implementation pattern is documented for teams building their own `did:tp:` resolver and credential stack.

1.4 Who Should Read This Book

This book assumes familiarity with TypeScript, public key cryptography fundamentals, and JSON-LD. You do not need prior experience with DIDs or verifiable credentials—Chapter 2 covers the specification from first principles. If you are a security architect evaluating decentralized identity for your organization, Chapters 2 through 5 give you the theoretical grounding. If you are a developer implementing agent identity, Chapters 6 and 7 give you the code.

DEMO

This is a free preview of the full edition.

Get the complete book at:

https://shop.pragma.vision