0x4149474E544D524B

AIAGENTMARK™

# The AI Agent Passport™

A Self-Sovereign Identity Standard for Personal AI Agents

INTRODUCING THE AIAGENTMARK STANDARD - VERSION 1.0

**Robert H. Rose**

Founder, AI Agent Mark™, LLC

aiagentmark.com

March 2026

Every Agent Marked · Every Identity Signed

0x4149474E544D524B

# A Note on Authorship

*The ideas, vision, architecture, and proof of concept implementation in this white paper are entirely my own. The writing was developed collaboratively with Claude, Anthropic's AI assistant, working from my concepts, research direction, and working implementation of LuoLongBot.*

*This is no different from an author directing a team of writers, or an executive working with a communications professional to articulate a vision. The direction, the accountability, and the conviction behind every claim in this document are mine. Claude executed the writing under my direction.*

*I live in the truth of my reality. I am open about how this document was produced because transparency is the foundation of everything AI Agent Mark™ stands for.*

— Robert H. Rose, March 2026

# Abstract

AI agents are acting in the world right now — reading email, scheduling meetings, sending messages, and making decisions on behalf of their owners. Yet there is no standard way to verify who these agents are, who owns them, or what they are authorized to do. Every government, enterprise, and institutional initiative addressing this problem is built for organizations. No standard exists for the individual.

The AI Agent Passport™, proposed in this paper, is a self-sovereign, cryptographically signed identity document for personal AI agents. It is generated by the agent owner, stored on their local machine, signed with an Ed25519 keypair, and verifiable by anyone who holds the owner's public key. No registration is required. No server is involved. No company stands in the middle.

The idea for the AI Agent Passport™ emerged from direct experience. While building LuoLongBot — a personal AI agent that reads email, manages a calendar, sends messages, and executes tasks on its owner's behalf — it became immediately clear that the agent had no way to identify itself. It could act in the world, communicate with humans and other systems, and operate autonomously — yet nothing about its communications conveyed who built it, who authorized it, or what it was permitted to do. That gap is not a design flaw in LuoLongBot. It is an absence in the entire personal AI agent ecosystem. The AI Agent Passport™ was created to fill it. A full proof-of-concept implementation of LuoLongBot is described in Section 9 of this paper.

The AI Agent Mark™ Standard — Version 1.0 — is open, free, and available for immediate adoption. This paper defines the problem, surveys the existing landscape, presents the technical architecture of the AI Agent Passport™, provides a formal threat model and security analysis, demonstrates alignment with existing standards including W3C DIDs and NIST SP 800-63, demonstrates proof of concept through a working implementation, and invites the personal AI agent community to adopt the standard.

> This paper is submitted as a formal public comment to the NIST AI Agent Standards Initiative (CAISI) in response to the draft Concept Paper "Accelerating the Adoption of Software and AI Agent Identity and Authorization," public comment period closing April 2, 2026. It is also submitted to the W3C Credentials Community Group for consideration as a complementary contribution to the Decentralized Identifiers ecosystem.

# Section 1: The Problem — AI Agents Are Acting Without Identity

## The Rise of the Personal AI Agent

Something significant is happening in the personal computing landscape. Individuals — not enterprises, not governments, not research institutions — are building private AI agents that run on their own machines, connect to their personal accounts, and act on their behalf throughout the day.

These agents read and send email. They manage calendars and contacts. They remember conversations. They respond to spoken commands from a smartphone. They execute tasks while their owners sleep. And they do all of this on standard consumer hardware, using publicly available tools, without specialized technical training.

This is not a future scenario. It is happening now, and the population of personal AI agent owners is growing rapidly. GitHub's Octoverse 2025 report documents 1.13 million public repositories now dependent on generative-AI SDKs — a 178% year-over-year increase — with over 693,000 new AI SDK-dependent repositories created in 2025 alone. Anthropic's Model Context Protocol (MCP), the primary open standard enabling agents to connect to tools and services, grew from its November 2024 launch to more than 10,000 active public server implementations within one year, with 97 million monthly SDK downloads. The tools to build personal AI agents — large language models, messaging APIs, authentication libraries, local databases — are accessible, affordable, and increasingly approachable for non-developers.

## The Identity Vacuum

As personal AI agents proliferate, a critical infrastructure gap has emerged: there is no standard way to establish who an AI agent is, who owns it, or what it is authorized to do.

When a human sends an email, centuries of social and legal infrastructure support attribution. A signature. A return address. A digital certificate. A paper trail. Accountability exists because identity exists.

When a personal AI agent sends an email on behalf of its owner, none of that infrastructure applies. The message arrives with no verifiable record of which agent

sent it, under whose authorization, with what declared capabilities, or whether the agent is what it claims to be.

This is not a theoretical concern. It is an operational reality that will define whether personal AI agents are trusted partners or persistent sources of confusion, fraud, and unaccountable action.

## The Evidence

The security and standards communities have recognized this crisis clearly and publicly. In January 2026, SC Media published an analysis titled "Identity becomes the 2026 battleground as AI erases trust signals." The analysis addressed the Model Context Protocol (MCP) — the dominant protocol for connecting AI agents to tools and services — and identified a fundamental gap: MCP was designed for interoperability, not security; it contains no built-in identity, no least-privilege enforcement, and no audit trail. As MCP becomes the primary standard for agentic AI, it will need a trust layer to verify which agents exist, who they represent, and what they are allowed to do.

That need — articulated by enterprise security experts describing a gap in the most widely adopted agentic AI protocol — is precisely what the AI Agent Passport™ addresses.

Additional data confirms the scale of the problem:

- 88% of organizations surveyed in early 2026 reported confirmed or suspected AI agent security incidents within the prior year.
- Only 47% of deployed AI agents are actively monitored — more than half operate with no security oversight or logging.
- 22% of organizations have no formal catalog of their AI agents at all.
- Non-human identities already outnumber human identities by 50 to 1 in the average enterprise environment, with projections reaching 80 to 1 within two years.

These figures describe enterprise environments. The personal agent landscape, where governance is even more nascent, presents an even more significant gap.

# Section 2: The Landscape — What Exists and Who Is Being Served

Several significant initiatives are actively working to address AI agent identity. Each represents a serious, well-resourced effort. And each serves a constituency that is not the personal AI agent owner.

## NIST AI Agent Standards Initiative

In February 2026, the Center for AI Standards and Innovation (CAISI) at NIST launched a formal AI Agent Standards Initiative, with a specific focus on agent authentication and identity infrastructure. NIST's Information Technology Laboratory published a draft Concept Paper, "Accelerating the Adoption of Software and AI Agent Identity and Authorization," with a public comment period open through April 2, 2026.

This is an important federal foundation. It is designed for enterprise environments and government adoption. It does not address personal agent owners and does not provide a self-sovereign identity mechanism an individual can implement today.

## NIST SP 800-63: Digital Identity Guidelines

NIST's SP 800-63 series establishes the federal framework for digital identity assurance through three dimensions: Identity Assurance Level (IAL), Authenticator Assurance Level (AAL), and Federation Assurance Level (FAL). These levels define a graduated trust model — from self-asserted identity at Level 1 through remotely verified identity at Level 2 to in-person verified identity at Level 3.

The AI Agent Mark™ trust hierarchy maps directly to this framework. The self-sovereign tier (self-generated Passport, no registration) corresponds to IAL1/AAL1 — self-asserted identity with single-factor cryptographic authentication. The registered tier (public key at register.aiagentmark.com) corresponds to IAL2/AAL2 — identity evidence collected and verified, multi-factor authentication. Future enterprise and audited tiers would correspond to IAL3/AAL3. This alignment is detailed in Section 5.

## Mastercard Verifiable Intent

On March 5, 2026, Mastercard announced Verifiable Intent, an open-source cryptographic proof standard for AI-driven transactions. The framework links a consumer's identity, their specific instructions, and the outcome of a transaction into

a tamper-resistant record, backed by FIDO Alliance, EMVCo, IETF, and W3C standards. Partners include Google, Fiserv, IBM, and Checkout.com.

Mastercard's philosophy — that as autonomy increases, trust must be proven rather than implied — aligns precisely with the AI Agent Mark™ approach. However, Verifiable Intent is designed for commercial AI agent transactions within the Mastercard network. It does not address personal agent identity outside of commerce.

### W3C Decentralized Identifiers (DIDs)

The W3C DID 1.0 specification, ratified as a Recommendation in July 2022, provides the foundational standard for self-sovereign identity on the internet. A DID is a globally unique, cryptographically verifiable identifier that requires no centralized registry. Academic research published in late 2025 proposed equipping AI agents with ledger-anchored DIDs and Verifiable Credentials.

The AI Agent Mark™ Standard is technically aligned with W3C DID principles — both use Ed25519 cryptography and both are self-sovereign. Section 5 of this paper provides a formal mapping between the AI Agent Passport™ and the DID Document specification. The DID ecosystem's requirement for distributed ledger infrastructure, DID method selection, and resolver implementation currently places full DID compliance beyond the reach of most personal agent owners building on consumer hardware — but the AI Agent Mark™ standard is designed as a bridge toward that ecosystem.

### Agent Name Service (ANS) — IETF Draft

ANS is a proposed IETF standard that would function as the DNS of AI agents — mapping agent identities to verified capabilities, cryptographic keys, and endpoints through a PKI-backed global directory. It is an important long-term infrastructure vision. It is also a draft standard requiring institutional certificate authority issuance, and it does not exist as an operational system today.

### Enterprise Identity and Access Management

Microsoft, CyberArk, Strata, and other enterprise security vendors have declared AI agent identity a top priority for 2026. Microsoft's security guidance states that every AI agent should be treated as a first-class identity, governed with the same rigor as human identities. CyberArk has extended its privileged access platform to cover AI

agents as machine identities.

These solutions govern agent fleets within enterprise IT environments. They require enterprise licensing, enterprise infrastructure, and enterprise-scale implementation. They are not designed for individuals.

## The Gap

The pattern is consistent and unambiguous. Every existing initiative serves organizations — enterprises, governments, financial networks, research institutions. The unit of concern is always the fleet, the deployment, the commercial transaction, the federal standard.

The person who builds a private AI agent on their own machine, connects it to their own accounts, and uses it to manage their own life has no identity standard available to them. They cannot verify their agent. They cannot sign its actions. They cannot prove, to themselves or to anyone else, that their agent is what it claims to be and acts only as they have authorized.

This is the gap the AI Agent Passport™ fills.

# Section 3: The Philosophy of the AI Agent Mark™ Passport — A Foundation Built on Truth, Accountability, and Freedom

Today there is no identity standard for personal AI agents. There is no mechanism to verify who built an agent, who owns it, what it is authorized to do, or who is accountable for its actions. This is the wild west — not a metaphor, but as reality. Powerful tools are proliferating faster than the standards to govern responsible use.

The AI Agent Mark™ Passport is the beginning of a new standard.

The progression of this standard follows the same arc as every transformative technology. The first international travel passport, issued in 1920, was a single sheet of paper with a photograph. It was clunky, manually verified, adopted by a small group of nations before the broader world understood why it mattered. It was also the foundation for a modern identity system to govern how billions of people move around the world.

We are at the 1920 moment for personal AI agents. The first iteration of the AI Agent Mark™ Passport is clunky. Verification is manual. Adoption requires early adopters to understand the problem before the broader world has felt its consequences. But the standard is real, it is implemented, and it is necessary — for the same reason the 1920 travel passport was necessary.

The standard begins with a commitment to freedom. Every personal AI agent owner has the right to declare their agent's identity without registry permission, without cost, and without a gatekeeper. The self-sovereign tier asks nothing of the owner except honesty. Trust is earned upward from that foundation — through voluntary registration, through verified identity, through demonstrated ethical operation over time.

This is the same architecture that secures the internet today. A self-signed certificate costs nothing and takes minutes. A domain-validated certificate takes a day and modest expense. An extended validation certificate requires documentation, an interview, and verified organizational identity. The same cryptographic foundation underlies all three. Only the verification of the person(s) behind the key escalates. The AI Agent Mark™ Passport follows this proven model — applied to the most urgent identity problem of the agentic AI era.

The revocation authority of the registry is not a contradiction of self-sovereignty. It is the standard fulfilling its lawful purpose. Just as a government issues, guarantees, and can revoke a travel passport from someone who uses it for criminal intent, the AI Agent Mark™ registry can revoke a Passport from an agent engaged in criminal activity.

## What Self-Sovereign Means

Self-sovereign identity is an identity model in which the individual controls their own identifier, independent of any central authority. There is no company to register with. No government to approve you. No server that must remain online for your identity to be valid. Your identity exists because you created it, and it is yours because you hold the keys. The W3C defines a self-sovereign identifier as one that is globally unique, cryptographically verifiable, and does not require any centralized registry.

The AI Agent Passport™ adopts this model and applies it to personal AI agents. The central distinction from every existing AI agent identity solution — from enterprise IAM platforms to Mastercard's Verifiable Intent to server-based agent authentication — is that the Passport requires no third party. The agent does not register with a service. No service issues a token. No token expires when a service disappears.

Your keypair is generated on your machine. Your Passport is signed with your private key. Your public key is the proof. No company can revoke your self-sovereign identity. No server going offline can invalidate your Passport. You created it. You own it. You can prove it.

This stands in a well-established tradition. It is the same foundation on which the web's SSL/TLS infrastructure was built — cryptographic proof replacing institutional vouching. It is the same foundation on which Bitcoin's self-sovereign proof model operates. And it is the natural extension of the W3C's decentralized identity work into the personal AI agent domain.

The personal AI agent owner is a new kind of digital actor. They bear personal responsibility for everything their agent does. That responsibility demands a corresponding capability: the ability to prove, clearly and verifiably, that their agent is who it claims to be, that they are its owner, and that its actions are within the scope they have authorized. The AI Agent Passport™ provides that capability.

It is important to distinguish identity from trust. A self-signed Passport proves that the holder of a specific private key created and signed this document. It does not, by

itself, prove that the agent is trustworthy or that its capability claims are accurate. Trust is built through a graduated hierarchy — from self-sovereign to registered to verified — detailed in Section 5. This mirrors the proven model of SSL/TLS certificates, where a self-signed certificate establishes cryptographic identity while higher tiers provide increasing levels of independently verified trust.

# Section 4: Introducing the AI Agent Passport™

## The One-Sentence Mission

> AI Agent Mark™ provides the trust layer that verifies which agents exist, who they represent, and what they're authorized to do.

## What the AI Agent Passport™ Is

The AI Agent Passport™ is a cryptographically signed identity document for a personal AI agent. It is a structured JSON file that declares, in verifiable terms:

- Who this agent is — its name, its unique ID, and the date it was created
- Who it belongs to — its owner's identity, represented as a salted SHA-256 hash that protects privacy while establishing ownership
- What it can do — a formal Capability Register listing every integration and function the agent has been built to perform
- That it is genuine — an Ed25519 digital signature that proves the Passport was created by the agent's owner and has not been tampered with
- What standard it follows — the AI Agent Mark™ Standard, Version 1.0

Once created, the Passport lives on the owner's machine. It is verifiable by anyone who holds the owner's public key. It establishes the agent as a known, attributed entity whose identity is cryptographically provable.

## The Capability Register

A defining feature of the AI Agent Passport™ is its Capability Register — a formal, signed list of every function the agent has been built to perform. Capability flags use a standardized naming convention under the AI Agent Mark™ Standard:

| Capability Flag | Description |
|---|---|
| MEMORY.PERSISTENT | Long-term memory across all conversations |
| INTERFACE.TELEGRAM | Natural language access via smartphone |
| INTEGRATION.GMAIL | Read, send, and search email |
| INTEGRATION.GOOGLE_CALENDAR | Manage schedule and events |
| INTEGRATION.GOOGLE_CONTACTS | Look up and manage contacts |
| INTEGRATION.M365 | Microsoft 365 email, calendar, and contacts |
| INTEGRATION.ALIAS_ROUTING | Alias-based communication routing and sub-agent dispatch |
| INTEGRATION.TWILIO_SMS | SMS messaging via Twilio |
| INTEGRATION.TWILIO_VOICE | Outbound voice calls via Twilio |
| SAFETY.THREE_TIER | Safe, dangerous, and blocked command classification |
| CREDENTIAL.SECURED | All credentials stored in .env, excluded from repositories |
| IDENTITY.AIAGNTMRK_V1 | Compliant with AI Agent Mark™ Identity Standard Version 1.0 |

The Capability Register is signed as part of the Passport. Any modification to the capability list invalidates the signature. The register is tamper-evident and self-declared by the agent's owner. It is important to note that the Capability Register is an owner attestation, not a third-party verification. It declares what the owner asserts the agent can do. Third-party capability verification is a planned future extension of the standard (see Section 8: Scope and Future Work).

## Capability Flag Naming Convention

Capability flags follow a formal naming grammar: NAMESPACE.ACTION, where NAMESPACE is an uppercase category identifier and ACTION is an uppercase descriptor. Both components MUST contain only ASCII uppercase letters, digits, and underscores. The period separator is required. Example: COMM.EMAIL_SEND. Reserved namespaces under AI Agent Mark™ Standard Version 1.0 are: IDENTITY (standard compliance), COMM (communications), MEMORY (persistence), SCHED (scheduling), DATA (data access), VOICE (voice/audio), and SAFETY (safety controls). Custom namespaces are permitted and MUST use a prefix not reserved above. The full namespace registry is maintained at aiagentmark.com.

## The AI Agent Mark™ Standard

The AI Agent Mark Standard™ — Version 1.0 — is an open standard for personal AI agent identity. It specifies the Passport document format, the required fields, the cryptographic signing algorithm, the Capability Register naming convention, and the verification process. It is free to adopt. It requires no registration. It is self-sovereign by design.

Every agent that generates an AI Agent Passport™ carries the IDENTITY.AIAGNTMRK_V1 capability flag. This flag signals compliance with the standard and allows any verifier to apply the same verification process to any AIAGNTMRK-compliant Passport.

# Section 5: Technical Architecture

## The Ed25519 Keypair

The AI Agent Passport™ is built on Ed25519 — the Edwards-curve Digital Signature Algorithm using Curve25519. Ed25519 is a modern, fast, and exceptionally secure digital signature algorithm, recognized as best practice by NIST (as part of EdDSA in FIPS 186-5), used in SSH key infrastructure, secure messaging protocols including Signal, and an increasing number of identity and authentication systems.

Key properties relevant to the AI Agent Passport™:

- 128-bit security level — resistant to all known classical attacks
- Compact keys — 32-byte private key, 32-byte public key
- Fast signing and verification — suitable for agent-speed interactions
- Deterministic signing — no randomness required during signing, eliminating bias attacks
- Widely supported — the Python cryptography library, Node.js, and most modern platforms implement Ed25519 natively

The owner's private key is generated locally and stored in the agent's .env file. It never leaves the owner's machine. The public key is stored as agent_public_key.pem and may be shared freely — it enables verification without enabling forgery.

## The Passport Document

The Passport is a structured JSON document containing the following fields:

| Field | Description |
|---|---|
| standard | "AIAgentMark" |
| version | "1.0" |
| passport_version | "1.0" |
| agent.name | The agent's given name |
| agent.id | Unique identifier: AGNT-[UUID] |
| agent.created | ISO 8601 timestamp of Passport creation |
| owner.identity_hash | Salted SHA-256 hash of owner's chosen handle |
| owner.identity_salt | Random 16-byte salt (hex-encoded) |
| owner.attribution | Plain-language privacy protection statement |
| capabilities | The Capability Register array |
| public_key | Owner's Ed25519 public key in PEM format |
| passport_expiry | ISO 8601 expiration date (recommended: 1 year) |
| signature | Ed25519 signature of the complete document |

## Example Passport Document

The following is a normative example of a conformant agent_passport.json, based on LuoLongBot's implementation (private key redacted; signature is illustrative):

```
{
 "standard": "AIAgentMark",
 "version": "1.0",
 "passport_version": "1.0",
 "passport_expiry": "2027-03-10T00:00:00Z",
 "agent": {
  "name": "LuoLongBot",
  "id": "AGNT-7f3d9a12-c841-4e6b-b002-1a2b3c4d5e6f",
  "created": "2026-03-10T14:22:00Z"
 },
 "owner": {
  "identity_salt": "a3f8c2e1d4b7901f",
  "identity_hash": "e9b4d2...c7a1f3",
  "attribution": "Owner identity is a salted SHA-256 hash. Handle not stored."
 },
 "capabilities": [
  "IDENTITY.AIAGNTMRK_V1", "COMM.EMAIL_SEND", "COMM.EMAIL_READ",
  "COMM.TELEGRAM", "COMM.SMS", "MEMORY.PERSISTENT",
  "SCHED.CALENDAR", "DATA.CONTACTS", "VOICE.TRANSCRIBE"
 ],
 "public_key": "-----BEGIN PUBLIC KEY-----\nMCowBQYDK2VwAy...EXAMPLE...
KEY==\n-----END PUBLIC KEY-----",
 "signature": "base64url-encoded-Ed25519-signature-bytes"
}
```

## The Digital Signature

The Passport document is signed in its entirety using the owner's Ed25519 private key. The signature covers every field — agent identity, owner hash, capability register, and public key. Any modification to any field, including adding or removing a single capability flag, produces a different document hash and invalidates the signature.

Verification is simple and definitive. Anyone who holds the owner's public key can run the verification module and confirm: (1) this Passport was created by the holder of the corresponding private key, (2) the Passport has not been modified since it was signed, and (3) the Capability Register has not been altered.

This is tamper-evident identity. It is not probabilistic. It is mathematically verifiable.

## Verification Algorithm

Any implementation claiming AIAGNTMRK v1.0 compliance MUST verify a Passport using the following steps:

1. Parse. Deserialize agent_passport.json as a JSON object. Reject if the document is not valid JSON or if any required field is absent.
2. Extract. Separate the signature field (base64url-encoded bytes) and the public_key field (PEM-encoded Ed25519 public key) from the document object.
3. Canonicalize. Remove the signature field from the document object and serialize the remainder using RFC 8785 JSON Canonicalization Scheme (JCS) to produce a deterministic byte sequence.
4. Decode. Base64url-decode the signature bytes. Parse the PEM public key to obtain the raw Ed25519 public key bytes.
5. Verify. Apply Ed25519 signature verification: verify(public_key, canonical_bytes, signature_bytes). A result of TRUE confirms the Passport is authentic and unmodified.
6. Check compliance. Confirm standard = "AIAgentMark", version = "1.0", and that IDENTITY.AIAGNTMRK_V1 is present in the capabilities array.
7. Check expiry (if present). If the passport_expiry field is present, compare against the current UTC timestamp. A Passport past its expiry date SHOULD be treated as requiring re-verification.

### Owner Identity Privacy

The owner's identity is represented in the Passport as a salted SHA-256 hash of a chosen handle. Version 1.0 requires a random 16-byte salt stored alongside the hash in the owner.identity_salt field. The salt prevents rainbow table attacks and ensures that even common handles produce unique, non-reversible hashes.

The actual name, username, or identifier is never stored in the Passport in readable form. The owner can verify their claim by producing the original handle and the salt, demonstrating the combination hashes to the stored value, without revealing the handle to anyone who does not already know it.

## Conformance Requirements

A conformant AIAgentMark v1.0 Passport MUST satisfy all of the following:

- The standard field MUST be "AIAgentMark" and version MUST be "1.0".
- The agent.id field MUST follow the format AGNT-[UUID] where UUID is a version 4 UUID.
- The owner.identity_hash MUST be a SHA-256 hash of the concatenation of identity_salt and the owner's handle.
- The signature MUST be a valid Ed25519 signature over the canonical JSON serialization of all fields except the signature field itself.
- The public_key MUST be a PEM-encoded Ed25519 public key corresponding to the private key that produced the signature.
- The capabilities array MUST contain at least IDENTITY.AIAGNTMRK_V1.
- The canonical JSON serialization MUST follow RFC 8785 (JSON Canonicalization Scheme, JCS). The signature MUST be computed over the JCS-canonicalized byte sequence of all fields except the signature field itself. This ensures interoperability across independent implementations.

## W3C DID Interoperability Mapping

The AI Agent Passport™ is designed to interoperate with the W3C Decentralized Identifiers (DIDs) v1.0 specification. While full DID method registration is planned for a future version, the following mapping demonstrates structural alignment between the Passport and a DID Document:

| AI Agent Passport™ Field | DID Document Equivalent |
|---|---|
| agent.id (AGNT-[UUID]) | did:aiagentmark:[UUID] |
| public_key (Ed25519, PEM) | verificationMethod (Ed25519VerificationKey2020) |
| signature | proof (Ed25519Signature2020) |
| capabilities array | service endpoints / capability declarations |
| owner.identity_hash | controller (hashed for privacy) |
| passport_expiry | No direct equivalent (extension) |

A future AI Agent Mark™ DID method (did:aiagentmark) would define resolution as follows: the DID is resolved by retrieving the agent_passport.json from the agent's registered endpoint or from the direct attachment in the communication. The DID Document is derived from the Passport fields using the mapping above. This approach allows AIAGNTMRK-compliant agents to participate in the broader DID and Verifiable Credentials ecosystem without requiring distributed ledger infrastructure at the personal agent scale.

### NIST SP 800-63 Alignment

The AI Agent Mark™ trust hierarchy maps to the NIST SP 800-63 Digital Identity Guidelines as follows:

| AIAGNTMRK Tier | NIST 800-63 Level | Requirements |
|---|---|---|
| Self-sovereign (Tier 0) | IAL1 / AAL1 | Self-asserted identity; single-factor cryptographic auth (Ed25519 keypair) |
| Registered (Tier 1) | IAL2 / AAL2 | Public key registered; identity evidence collected and verified remotely |
| Enterprise (Tier 2) | IAL2+ / AAL2 | Organizational identity verified; higher trust |
| Audited (Tier 3, future) | IAL3 / AAL3 | In-person or supervised remote proofing; full compliance audit |

This mapping ensures that the AI Agent Mark™ trust hierarchy is legible to federal systems and enterprise environments that already use NIST 800-63 as their identity assurance framework.

### The File Structure

An AIAGNTMRK-compliant agent produces three Passport-related files:

- agent_passport.json — the signed Passport document
- agent_public_key.pem — the Ed25519 public key, freely shareable
- passport_verify.py — the verification module

The private key is stored in the agent's .env file alongside all other credentials. The .env file is excluded from any repository by .gitignore. This is consistent with industry-standard credential management practice.

# Section 6: Threat Model and Security Analysis

A responsible identity standard must clearly articulate what threats it addresses, what threats it partially mitigates, and what threats fall outside its current scope. This section provides that analysis for the AI Agent Passport™ Version 1.0.

## Threats Addressed

| Threat | Description | AIAGNTMRK Mitigation |
|---|---|---|
| Agent impersonation | A third party claims to operate an agent they do not control | Ed25519 signature verification proves Passport origin; forgery requires the private key |
| Passport tampering | An attacker modifies a Passport after signing (e.g., inflating capabilities) | Any field modification invalidates the signature; verification detects all tampering |
| Unsigned agent actions | An agent acts without any identity attribution | The Passport and direct attachment model ensure every compliant agent carries verifiable identity |
| Capability misrepresentation after signing | A previously valid Passport is altered to add capabilities | Signature covers the full capability register; changes invalidate the signature |

## Threats Partially Mitigated

| Threat | Description | Status in v1.0 |
|---|---|---|
| False capability claims | Owner signs a Passport claiming capabilities the agent does not have | The standard proves the owner attested to these capabilities. Third-party capability verification is planned for future versions. |
| Owner identity spoofing | A malicious actor creates a Passport with a false owner hash | The salted hash prevents reverse-engineering but does not prove real-world identity. Higher tiers (registered, audited) add identity verification. |
| Public key distribution trust | A verifier receives a public key from an untrusted channel | The registry (register.aiagentmark.com) provides a trusted distribution channel. Direct attachment relies on the transport channel's security. |

### Threats Outside Current Scope

| Threat | Description | Planned Response |
|---|---|---|
| Key compromise | An attacker gains access to the owner's private key | Key rotation and revocation mechanisms are specified in Section 7 (Key Lifecycle). Full revocation infrastructure is planned for v1.1. |
| Sybil attacks | One actor generates many Passports to create false network presence | The self-sovereign tier does not prevent this. The registered tier requires identity verification, limiting Sybil attacks. |
| Replay attacks | An attacker reuses a previously valid signed communication | Timestamp and nonce fields for signed communications are planned for v1.1. |
| Post-quantum cryptography | Future quantum computers may threaten Ed25519 | Ed25519 is safe against all known classical attacks. The standard will migrate to quantum-resistant algorithms (e.g., CRYSTALS-Dilithium) when NIST PQC standards mature. |

This threat model is intended to evolve with the standard. Community review and adversarial analysis are welcomed.

# Section 7: Key Lifecycle Management

A cryptographic identity standard must address the full lifecycle of keys, not just their initial generation. This section specifies key lifecycle management for AI Agent Mark™ v1.0.

## Key Generation

The Ed25519 keypair is generated locally on the owner's machine using a cryptographically secure random number generator. The private key is stored in the agent's .env file. The public key is exported as agent_public_key.pem. No key material is ever transmitted to AI Agent Mark™ or any third party during generation.

## Key Rotation

Owners SHOULD rotate their keypair periodically (recommended: annually) or immediately upon suspicion of compromise. Key rotation follows this process:

- Generate a new Ed25519 keypair.
- Re-sign the Passport with the new private key, updating the public_key field.
- If registered at register.aiagentmark.com, update the registered public key.
- Distribute the new public key to known verifiers.
- The previous Passport remains valid until its passport_expiry date or until verifiers update their cached public key.

## Passport Expiration

Version 1.0 introduces the passport_expiry field (recommended: 1 year from signing). Expired Passports SHOULD be treated as requiring re-verification. Verifiers MAY choose to accept expired Passports with reduced trust for a grace period.

## Revocation

Version 1.0 provides two revocation mechanisms:

- Self-revocation: The owner generates a new Passport (or deletes the existing one) and distributes the updated version. The old Passport becomes stale when verifiers update their records.

- Registry revocation: For agents registered at register.aiagentmark.com, the registry can mark a public key as revoked. Verifiers querying the registry will receive the revocation status. This mechanism is available only for the registered tier and above.

A formal Certificate Revocation List (CRL) or OCSP-equivalent mechanism for real-time revocation checking is planned for v1.1. Version 1.0 acknowledges this as a known limitation.

### Compromise Recovery

If an owner's private key is compromised, the recommended recovery procedure is:

- Immediately generate a new keypair and re-sign the Passport.
- If registered, contact register.aiagentmark.com to revoke the old public key and register the new one.
- Notify known verifiers of the key change.
- Review agent logs for unauthorized actions taken during the compromise window.

# Section 8: Usage, Verification, and the Trust Hierarchy

## The Trust Hierarchy

Just as SSL/TLS certificates established a graduated hierarchy of trust for websites, the AIAGNTMRK standard establishes a parallel hierarchy for personal AI agents:

| AIAGNTMRK Tier | Description |
|---|---|
| Tier 0: Self-sovereign | Self-generated AI Agent Passport™ — local keypair, no registration. Corresponds to a self-signed certificate. |
| Tier 1: Registered | Public key registered at register.aiagentmark.com — identity confirmed, signature verifiable remotely. Corresponds to a domain-validated certificate. |
| Tier 2: Enterprise | Enterprise-registered agent — organizational identity verified, higher trust tier. Corresponds to an organization-validated certificate. |
| Tier 3: Audited (future) | Fully audited AIAGNTMRK compliance — maximum trust tier. Corresponds to an extended-validation certificate. |

The self-signed tier is the foundation. It is free, immediate, self-sovereign, and available to every personal AI agent owner today. Higher tiers offer increasing levels of independently verifiable trust.

## The Direct Attachment Model

The core principle is simple: the Passport travels with the agent's communication, the same way a business card travels with a handshake. The recipient gets everything they need to verify identity in the same transaction as the communication itself.

In the early days of the standard, three files travel together: agent_passport.json, agent_public_key.pem, and passport_verify.py. As the standard matures and verification tools become common, only the Passport JSON will be needed.

## Email

A first contact email from an AIAGNTMRK-compliant agent includes the Passport JSON as an attachment with a plain-language explanation in the email body. The email header carries the machine-readable signal: X-AIAgentMark-Passport: attachment://agent_passport.json. For subsequent communications, the attachment is dropped and only the header remains.

### Conversational Interfaces

When an AIAGNTMRK-compliant agent initiates contact via Telegram or other conversational interfaces, it responds with a plain-language identity statement followed by the Passport JSON sent as a file attachment.

### API and Agent-to-Agent

When an AIAGNTMRK-compliant agent interacts with another AI agent or an API endpoint, the Passport travels in the request headers: X-AIAgentMark-Passport (base64-encoded Passport JSON) and X-AIAgentMark-PublicKey (base64-encoded public key). The receiving system can decode and verify the Passport programmatically in milliseconds.

### Online Verification

The AI Agent Passport™ includes a local verification module (passport_verify.py) for technical users. For broader accessibility, AI Agent Mark™ provides a public verification resource at verify.aiagentmark.com. Any individual, organization, or platform can submit a Passport for instant verification of signature validity, standard compliance, capability register integrity, ownership attribution, and registry status.

### The Public Key Registry

Agent owners who want to participate in the verified tier register their Ed25519 public key at register.aiagentmark.com. Registration is optional. It elevates an agent from self-sovereign (Tier 0) to registered (Tier 1), enabling cryptographic verification against a known, queryable public key. The registry also supports revocation status queries for registered keys.

# Section 9: Proof of Concept — LuoLongBot

## Built in One Week

The AI Agent Passport™ is not a theoretical proposal. It is implemented and operational. LuoLongBot — a fully functional personal AI agent — was built from scratch in seven days on a standard Windows 11 machine using Windows Subsystem for Linux (WSL), Claude Code, and publicly available APIs. No new hardware was purchased. No cloud virtual machine was provisioned. The agent runs locally, starts automatically with Windows, and operates silently in the background.

LuoLongBot is accessed and controlled entirely through Telegram on an iPhone. It understands natural language voice commands and text instructions. It has been operational continuously since its initial build.


LuoLongBot — Proof of Concept Implementation, AI Agent Mark™ Standard Version 1.0

## LuoLongBot's Capability Register

| Capability | Integration |
|---|---|
| MEMORY.PERSISTENT | SQLite database — four tables across all conversations |
| INTERFACE.TELEGRAM | Telegram Bot API — smartphone control via natural language |
| INTEGRATION.GMAIL | Google Gmail API — read, send, search email |
| INTEGRATION.GOOGLE_CALENDAR | Google Calendar API — schedule management |
| INTEGRATION.GOOGLE_CONTACTS | Google People API — address book access |
| INTEGRATION.M365 | Microsoft Graph API — Outlook, Exchange, Microsoft Contacts |
| INTEGRATION.ALIAS_ROUTING | Alias-based routing across 7 email aliases |
| INTEGRATION.TWILIO_SMS | Twilio — outbound SMS messaging |
| INTEGRATION.TWILIO_VOICE | Twilio — outbound voice calls |
| SAFETY.THREE_TIER | Safe, dangerous, and blocked command classification |
| CREDENTIAL.SECURED | .env + .gitignore — zero credentials in code |
| IDENTITY.AIAGNTMRK_V1 | AIAgentMark Standard Version 1.0 |

### What This Demonstrates

LuoLongBot demonstrates four things about the AI Agent Passport™:

- The standard is implementable today, on existing consumer hardware, by a non-developer.
- The Capability Register accurately reflects a real, complex, operational agent — not a demonstration stub.
- The cryptographic architecture — Ed25519 signing, salted SHA-256 owner hash, JSON Passport format — works as specified.
- The self-sovereign model is practical — no server, no registration, no third-party dependency.

# Section 10: Scope and Future Work

The AI Agent Mark™ Standard Version 1.0 deliberately limits its scope to the foundational identity layer. The following capabilities are planned for future versions:

### Version 1.1 (Planned)

- Formal CRL/OCSP-equivalent revocation mechanism for real-time revocation checking.
- Timestamp and nonce fields for signed agent communications to prevent replay attacks.
- DID method registration (did:aiagentmark) with the W3C DID Method Registry.
- Verifiable Credentials integration for third-party capability attestation.

### Version 2.0 (Vision)

- Third-party capability verification — independent auditors can attest to an agent's declared capabilities.
- Post-quantum cryptographic migration path as NIST PQC standards mature.
- Multi-agent delegation chains — an agent can verifiably delegate authority to sub-agents.
- Formal IETF or W3C specification submission.

Community feedback and contributions to the standard's evolution are welcomed at aiagentmark.com.

### Legal and Liability Framework

The AI Agent Mark™ Standard Version 1.0 addresses the technical dimension of AI agent identity. It is important to note that identity establishment is a necessary but not sufficient condition for legal accountability. A cryptographically verifiable Passport proves that a specific private key signed a specific document — it does not, by itself, establish legal liability for an agent's actions, determine jurisdiction, or create enforceable obligations. The legal and liability framework for personal AI agents — including questions of who bears responsibility when an agent causes harm, what evidentiary weight a self-signed Passport carries, and how identity claims interact with existing digital signature law — is a distinct and important body of work. AI

Agent Mark™ scopes this dimension to future work and community input, and invites legal scholars, policymakers, and standards bodies to contribute to developing that framework in parallel with the technical standard.

## Conformance Testing

A conformance test suite — a set of valid and invalid Passport documents that any implementation can test against — is planned for Version 1.1. The test suite will include: valid Passports with correct signatures that MUST verify successfully; Passports with tampered fields that MUST fail verification; Passports with missing required fields that MUST be rejected; and Passports with expired passport_expiry dates that SHOULD trigger re-verification warnings. Independent implementation and testing of the standard is actively encouraged. Developers who build conformant verifiers are invited to register their implementations at aiagentmark.com, contributing to a growing ecosystem of interoperable tools.

# Section 11: The Broader Vision

## What Adoption Looks Like

The AI Agent Mark™ Standard is self-adopting. Any personal AI agent owner who generates an AI Agent Passport™ has adopted the standard. No application is required. No membership is required. No fee is charged.

As adoption grows, the standard gains authority through network effect. An AI agent carrying an AIAGNTMRK_V1 Passport can present its Passport to any system or person that understands the standard. The more agents that carry Passports, the more valuable each individual Passport becomes.

## The NIST Opportunity

NIST's public comment period for the AI Agent Identity and Authorization Concept Paper closes April 2, 2026. This white paper is submitted as a formal public comment.

The AI Agent Mark™ Standard does not compete with NIST's enterprise-focused framework. It complements it — by addressing the constituency NIST's initiative explicitly does not serve: the individual personal AI agent owner. A federal standards framework that acknowledges the personal agent layer strengthens the entire identity ecosystem. AI Agent Mark™ invites NIST to recognize the personal agent owner as a stakeholder whose identity needs are distinct from, and as important as, those of enterprise deployments.

Specifically, this paper recommends that the NIST AI Agent Standards Initiative:

- Recognize personal AI agents as a distinct identity category alongside enterprise agent deployments.
- Acknowledge self-sovereign identity models as a valid foundation for personal agent identity, consistent with W3C DID principles.
- Consider the graduated trust hierarchy model (self-sovereign → registered → verified → audited) as a framework that bridges personal and enterprise identity needs.
- Align personal agent identity assurance levels with the existing NIST SP 800-63 framework.

# Conclusion

AI agents are no longer experimental. They are operational, they are personal, and they are acting in the world on behalf of their owners right now. The identity infrastructure to support this reality does not yet exist for the individual.

The AI Agent Passport™ proposes a solution that is self-sovereign, cryptographically verifiable, immediately implementable, and free to adopt. It requires no server, no registration, and no company in the middle. It is built on the same cryptographic foundations as the most trusted identity systems in the world, applied at the personal scale where the need is greatest and the gap is widest.

The standard provides a formal threat model, a clear key lifecycle management process, conformance requirements, and demonstrated alignment with both W3C DIDs and NIST SP 800-63. It distinguishes clearly between what Version 1.0 addresses and what is planned for future versions.

The AI Agent Mark™ Standard — Version 1.0 — is the beginning. As the personal AI agent ecosystem matures, the standard will evolve with it. What will not change is the philosophical foundation: your agent, your keys, your proof.

# Recommendations

1. Implement and test. Independent implementation and testing of the AI Agent Mark™ Standard is actively encouraged. Personal AI agent owners can generate their first AI Agent Passport™ using the implementation guide at aiagentmark.com.

2. Adopt the standard. If you are building agents, tools, or platforms that interact with personal AI agents, recognize and verify AI Agent Mark™-compliant Passports.

3. Comment. Submit your feedback to the NIST AI Agent Standards Initiative before April 2, 2026. Personal agent owners deserve a seat at the standards table.

**Every agent marked. Every identity signed.**

0x4149474E544D524B

# References

1. SC Media. “Identity becomes the 2026 battleground as AI erases trust signals.” January 6, 2026. https://www.scworld.com

2. NIST Center for AI Standards and Innovation. “Announcing the AI Agent Standards Initiative.” February 2026. https://www.nist.gov/caisi/ai-agent-standards-initiative

3. NIST Information Technology Laboratory. “Accelerating the Adoption of Software and Artificial Intelligence Agent Identity and Authorization.” Draft Concept Paper, February 5, 2026. Comment period closes April 2, 2026.

4. NIST SP 800-63-4 (Draft). “Digital Identity Guidelines.” December 2024. https://pages.nist.gov/800-63-4/

5. PYMNTS. “Mastercard Unveils Open Standard to Verify AI Agent Transactions.” March 5, 2026. https://www.pymnts.com

6. Gravitee.io. “The State of AI Agent Security 2026.” Survey of 919 executives and practitioners. https://www.gravitee.io/state-of-ai-agent-security

7. Narajala, et al. “AI Agents with Decentralized Identifiers and Verifiable Credentials.” arXiv:2511.02841, November 2025. https://arxiv.org/abs/2511.02841

8. HID Global. “Trust Standards Evolve: AI Agents, the Next Chapter for PKI.” November 2025. https://blog.hidglobal.com

9. W3C. “Decentralized Identifiers (DIDs) v1.0.” W3C Recommendation, July 19, 2022. https://www.w3.org/TR/did-core/

10. W3C. “Verifiable Credentials Data Model v2.0.” W3C Recommendation, May 2024. https://www.w3.org/TR/vc-data-model-2.0/

11. Microsoft Security Blog. “Four Priorities for AI-Powered Identity and Network Access Security in 2026.” January 2026. https://www.microsoft.com/en-us/security/blog

12. Strata. “Agentic AI Security: 8 Strategies for AI Agent Security in 2025.” January 2026. https://www.strata.io

13. CyberArk. “What’s Shaping the AI Agent Security Market in 2026.” January 2026. https://www.cyberark.com

14. The Hacker News. “9 Identity Security Predictions for 2026.” February 2026. https://thehackernews.com

15. NIST FIPS 186-5. “Digital Signature Standard (DSS).” February 2023. (Includes EdDSA/Ed25519.)

16. GitHub. “Octoverse 2025: A new developer joins GitHub every second as AI leads TypeScript to #1.” October 2025. https://github.blog/news-insights/octoverse/

17. Anthropic. “Donating the Model Context Protocol and Establishing the Agentic AI Foundation.” December 2025. https://www.anthropic.com/news/donating-the-model-context-protocol-and-establishing-of-the-agentic-ai-foundation

# Appendix A: Potential Platform Integrations

The AIAGNTMRK standard presents opportunities for AI platforms to strengthen protection for personal AI agent operators. Platforms with input awareness could natively recognize the AIAGNTMRK cryptographic architecture, using the hex bookend markers as a compliance signal and Ed25519 signatures for full verification. For registered agents, platforms could verify cryptographic signatures in real time before executing system-level actions. The AI Agent Mark™ standard is offered freely for native implementation by any platform — no licensing, no fee.

# About the Author

Robert H. Rose is the founder of AI Agent Mark™, LLC and SaniTrace™, LLC, based in St. Louis, Missouri.

He spent four years in Beijing at FleishmanHillard Inc. as the firm's youngest Vice President, where he established the firm's Information and Technology practice in China. He subsequently served as Director of Public Relations and Government Affairs at the Donald Danforth Plant Science Center.

He founded SaniTrace™ in 2015 — the first blockchain-based food safety tracking system for ground beef — building the entire system himself, including custom hardware, private blockchain connectivity, and consumer SMS recall alerts via Twilio.

LuoLongBot, the AI agent that serves as proof of concept for the AI Agent Passport™, was built by Rose in seven days in March 2026 on a typical Windows 11 computer using Windows Subsystem for Linux, Claude Code, and publicly available APIs.

While building LuoLongBot, a name based on his Chinese name, Luo Long (罗龙), Rose came up with the concept for the AI Agent Passport™ and further developed the idea in concert with Claude. This project, like his life, is driven by his credo, "He lives in the truth of his reality."

*rrose@aiagentmark.com* *linkedin.com/in/roberthrose* *aiagentmark.com*