

clara •

PASSPORT.md

Foundational Identity for Autonomous Economic Actors

OPEN RESEARCH PROPOSAL

VERSION 3.1 / JULY 2026

A lightweight, portable proposal for summarizing an AI agent's sponsor, authority, evidence, and operating context - published to be tested, criticized, forked, or discarded in the open.

clarainsure.com / github.com/lucasbarnes96/agent-passport

A White Paper on Verifiable Identity, Trust, and Coordination Infrastructure for AI Agents

Version 3.1 - Revised Draft · July 3, 2026

Working document for discussion and iteration. Not legal, regulatory, or financial advice.

Executive Summary

Autonomous AI agents are becoming first-class economic actors. They discover information, negotiate, transact, delegate, and coordinate at machine speed across digital (and increasingly physical) environments.

Existing rails enable action, and an expanding set of identity, authorization, registry, attestation, monitoring, and payment systems is being built around them. The open question is how evidence should travel between those systems and the humans, counterparties, and institutions that rely on an agent.

Agent Passport is one lightweight proposal: a portable PASSPORT .md file that can summarize an agent's sponsor, capabilities, permissions, history, and evidence in a human-readable and machine-parseable form. It does not replace runtime authorization, workload identity, platform registries, cryptographic credentials, or monitoring. It explores whether a simple public artifact can make those systems easier to inspect and discuss.

- A consistent summary of an agent's claimed identity and operating context.
- References to sponsors, capabilities, permissions, and history.
- Space for third-party attestations and dynamic trust signals.
- Portability across repos, platforms, and agent-readable documentation.
- A possible input to risk assessment, compliance, access control, financial services, and multi-agent coordination.

This is broader than risk and insurance alone, but the analogy has limits. A Markdown file is not a government credential, an authentication protocol, or an authorization decision. Its useful role is closer to an inspectable cover sheet that points to the systems and evidence another party can verify.

Key benefits include:

- **Discovery & Interoperability:** Agents can be found, understood, and interacted with consistently.
- **Trust & Authorization:** Verifiable claims about capabilities, history, and constraints (KYA - Know Your Agent).
- **Accountability & Compliance:** Binding to responsible entities with auditable records.
- **Economic Enablement:** Portable reputation and credentials that support lending, insurance, payments, and delegation.
- **Risk Management:** Structured data for dynamic scoring, benchmarking, and targeted protection (one powerful use case among many).

The intended result is modest but useful: give relying parties a clearer starting point for asking what an agent is, who stands behind it, what it may do, and which claims can actually be verified.

1. The Emerging Need for Agent Identity

AI agents are evolving from chat-based assistants into autonomous systems that plan, use tools, execute actions, and interact with other agents and systems.

Current challenges:

- Every new agent or platform reinvents context, identity, and trust assembly.
- No portable way to verify "who" an agent is, "what" it can do, or "why" it should be trusted.
- Delegation chains and multi-agent systems create accountability gaps.
- Rapid evolution (model changes, new tools, behavioral drift) makes static approaches obsolete.
- Lack of standardized identity hinders interoperability, compliance, and safe scaling.

Human societies solved similar problems with layered identity systems: passports for travel and basic identity, KYC for financial participation, driver's licenses for specific privileges, corporate registries for legal entities, and credit scores for economic reputation.

Agent Passport applies the same principle to software actors in a machine-native, portable format.

1.1 What Is an Agent? (Scope and Definition)

For this framework, an **AI agent** is a software system that uses AI to perceive inputs or environments, reason toward goals, plan and execute sequences of actions (including tool use), adapt based on outcomes, and operate with a meaningful degree of autonomy - without requiring continuous human intervention for every decision.

Important Distinctions:

- **Autonomy Level:** Ranges from highly scaffolded/harnessed agents (strong external orchestration and oversight) to self-contained agents (persistent state, dedicated resources, long-horizon independent operation).
- **Resource Access:** Some agents operate within tightly controlled environments with mediated compute and permissions. Others have direct access to wallets, APIs, compute instances, data stores, or even physical systems.
- **Action Capability:** The ability to initiate effects in the world (payments, code changes, data modifications, physical actions) is what creates economic and risk exposure.

This framework focuses on agents with action capability and meaningful autonomy, as these are the ones that participate in economic and operational systems and therefore benefit most from portable, verifiable identity.

Legal Context: Argentina's 2026 corporate-law reform proposal includes digital company infrastructure and decentralized autonomous organizations. It does not settle whether an AI agent can itself hold legal personhood, but it makes software-mediated governance a practical corporate-law question rather than a distant thought experiment.

1.2 The Current Infrastructure Landscape

Agent Passport is not entering an empty field.

- **Identity and authorization:** NIST has launched an AI Agent Standards Initiative and a separate effort on software and AI-agent identity and authorization. These programs focus on applying established identity standards, authorization controls, auditing, and non-repudiation to agent systems.
- **Commerce:** Visa's Trusted Agent Protocol uses existing web-signature infrastructure to help merchants recognize authorized shopping agents. Payment networks, wallets, and commerce platforms are developing their own mandate and transaction controls.
- **Enterprise governance:** Workday now uses the name **Agent Passport** for a platform that records independent security attestations, monitors agents at runtime, and supports revocation. Its product is materially different from this open Markdown proposal, but the name collision is important and should be explicit.
- **On-chain registries:** ERC-8004 and adjacent registry patterns provide discoverability, identity references, and reputation signals for agents operating through wallets and smart contracts.
- **Governance frameworks:** Singapore's Model AI Governance Framework for Agentic AI emphasizes bounded powers, meaningful human accountability, lifecycle controls, and transparency.

These systems solve different layers of the problem. A registry can identify an agent without proving it is safe. An authorization system can approve an action without explaining the broader risk. A runtime monitor can detect drift without giving a counterparty a portable summary. A PASSPORT .md file should therefore be evaluated as connective documentation, not as a replacement for any of them.

The project is best understood as an open research alternative: deliberately simple, easy to fork, and useful for testing which fields and evidence references matter before stronger infrastructure is available or interoperable.

2. The Agent Passport Concept

The Agent Passport is a portable identity-and-evidence summary for AI agents.

Core Analogy to Human Systems:

Human System	Agent Equivalent
Passport	Basic portable identity and participation rights
KYC / Digital Agent Passport	Verified identity bound to responsible entities
Driver's License	Scoped capability credentials
Corporate Registry	Entity-level identity and limited liability
Credit File	Historical behavior and trustworthiness

Implementation: A simple, elegant file-based format (PASSPORT .md) using Markdown + structured YAML frontmatter (inspired by the Open Knowledge Format style). This keeps it:

- Human-readable and editable.
- Machine-parseable.
- Git-native and versionable.
- Extremely portable (files, repos, tarballs, filesystems).

- Easy to extend and integrate.

Example Structure (high-level):

```

---
type: agent-passport
title: "Procurement-Agent-v3.2"
agent_id: "eip155:1:0x..."
version: "3.2.1"
sponsor: "Entity or Human ID"
capabilities: ["vendor-discovery", "invoice-processing", "payment-initiation"]
permissions:
  max_single_txn: 50000
  requires_approval_above: 10000
attestations:
  - type: benchmark
    name: "procurement-v1"
    score: 87
    date: "2026-06-01"
  - type: platform
    source: "Stripe-MCP"
    successful_txns_30d: 1247
reputation_signals:
  - type: onchain
    registry: "ERC-8004"
    score: "...
status: active
timestamp: "2026-06-17T00:00:00Z"
---

# Rich Markdown Content
## Description
## History & Evidence
## Sponsorship & Delegation

```

The frontmatter holds queryable structured data. The body holds narrative, evidence, and context.

This format aligns with emerging real-world efforts:

- ERC-8004 on-chain identity registries (NFT-based agent IDs with metadata).
- Know Your Agent (KYA) frameworks from multiple vendors.
- Digital Agent Passport concepts already deployed in commerce.
- NIST, IETF, OpenID Foundation, and NCCoE work on agent identity and authorization standards.

3. Possible Uses of a Portable Summary

While risk assessment and insurance are powerful applications, a portable summary could offer broader value as connective documentation:

3.1 Discovery and Interoperability

A shared summary could make agents easier to discover and compare across platforms. Whether it reduces integration friction depends on adoption and on links to authoritative records.

3.2 Authentication, Authorization, and Access Control

The file can point to proofs of identity and authorized scope before actions. Enforcement remains with OAuth, SPIFFE/SVID, workload identity, platform policy, and other authorization systems.

3.3 Accountability and Auditability

It can state the sponsor or responsible entity and link to audit records. The statement itself does not create legal accountability or prove the records are complete.

3.4 Reputation and Trust Signals

It can reference attestations from benchmarks, platforms, and interactions. Portable reputation would require common definitions, trustworthy issuers, revocation, and anti-gaming controls beyond the file.

3.5 Enabling Financial and Economic Services

Identity, history, and attestations may support lending, credit, and insurance when they are independently verified. Staking or collateral references could be included, but enforcement belongs elsewhere.

3.6 Legal and Governance Wrappers

The summary could document legal wrappers and delegation chains. It does not create personhood, limited liability, agency authority, or clear provenance by itself.

3.7 Multi-Agent Coordination and Delegation

Consistent identity and capability claims may make delegation easier to inspect. Safe orchestration still requires runtime policy and verification.

3.8 Security and Governance at Scale

Organizations could use the fields as an inventory view while relying on adopted governance and Zero Trust systems for control, monitoring, and decommissioning.

4. Path to Insurance

Insurance is the application Clara is most interested in testing.

The stack:

- 1 **Payment rails** (Coinbase, x402, stablecoins) let agents transact.
- 2 **Legal wrappers** and contracts assign rights, obligations, and responsibility.
- 3 **Identity, authorization, and monitoring systems** generate decision-grade evidence.
- 4 **A portable summary** may make that evidence easier for another party to inspect.
- 5 **Insurance** can transfer defined losses when the risk and responsible insured are sufficiently clear.

For underwriting, the relevant question is not whether an agent has a Markdown file. It is whether the deploying organization can demonstrate authority, controls, evidence, accountability, and plausible loss. A portable summary might help an underwriter locate:

- benchmark attestations and transaction history;
- financial, technical, and communications limits;
- monitoring, incident, shutdown, and recovery evidence;
- contracts, responsible entities, and existing coverage.

This is why Agent Passport remains an open side project rather than Clara's core product. Clara's nearer-term underwriting work is the Agentic Authority Supplement to `INSURANCE.md`, which can use Passport references where helpful without depending on Passport adoption.

5. Technical and Standards Alignment

The Agent Passport is designed to be lightweight and composable:

- **Format:** Markdown + YAML frontmatter (OKF-inspired) for the project artifact.
- **On-Chain Layer:** Aligns with ERC-8004 for discoverability, portable IDs, and on-chain signals.
- **Verification:** Cryptographic attestations, platform integrations (MCP, x402), and optional zero-knowledge proofs.
- **Ecosystem:** Compatible with KYA frameworks, Digital Agent Passport concepts, NIST agent identity efforts, and OpenID Foundation agentic identity work.

6. Research Plan and Decision Gates

Immediate:

- Map each proposed field to existing identity, authorization, registry, monitoring, and payment infrastructure.
- Publish reference examples and invite criticism from builders and relying parties.
- Test whether readers can answer sponsor, authority, evidence, and revocation questions faster with the summary than without it.

Build gate:

- Build a validator only if real users repeatedly exchange the file.
- Add integrations only where they point to authoritative evidence rather than duplicate it.
- Rename or retire the project if the Workday name collision creates more confusion than value.

Stop conditions:

- Adopted platforms already provide an adequate portable view.
- Relying parties do not use the file in a real decision.
- The summary creates false confidence or becomes stale faster than it can be verified.

Success is not a new standard for its own sake. It is evidence that a simple portable summary improves a real decision-or clear evidence that Clara should stop building it.

7. Conclusion

The agentic future needs identity, authorization, evidence, monitoring, accountability, and recourse. No single file or protocol supplies all of them.

Agent Passport tests one narrow proposition: whether a portable, inspectable summary can help humans and agents understand who sponsors a system, what authority it has, what evidence supports its claims, and where stronger verification lives.

If more adopted infrastructure makes this proposal unnecessary, that is a useful result. If the file becomes a convenient bridge between those systems, that is useful too. The project is being published in the open so the idea can be tested, criticized, forked, or discarded on evidence rather than protected as a finished answer.

We invite collaboration on the fields, examples, failure modes, and interoperability questions.

References & Further Reading

- [Google Cloud Open Knowledge Format \(OKF\)](#)
- [ERC-8004: Trustless Agents](#)
- [Workday Agent Passport](#)
- [NIST AI Agent Standards Initiative](#)
- [NIST/NCCoE work on software and AI-agent identity and authorization](#)
- [Visa Trusted Agent Protocol](#)
- [Singapore Model AI Governance Framework for Agentic AI](#)
- IETF and OpenID Foundation efforts on agent authentication and identity management
- Open Agent Passport (OAP) specification for pre-action authorization
- Argentina corporate-law reform and DAO proposals (2026)

PASSPORT.md - Foundational Identity for Autonomous Economic Actors

Version 3.1 · July 3, 2026 · Open research proposal for discussion and collaboration.

Canonical source: github.com/lucasbarnes96/agent-passport