STRATEGIC CONCEPT · WORKLOFT.AI · 3 MAY 2026

AgentPass

A standardised real-time verification primitive for AI agents — answering the question *"is this agent currently authorised, in good audit standing, and capable of the action it's about to take?"* as a single API call. The institutional governance layer the agent ecosystem doesn't yet have.

Author: Workloft.ai LTD · ICO C1912528
Audience: Alfred Churchill (founder) — strategic flag-plant for 2026 build
Status: v0.1 concept · for sharpening with Vera + Alfred
Companion artefacts: V0.1 spec (separate doc), 18-task Gary build list

The problem AgentPass solves

> When agent A from organisation X interacts with party B (a counterparty firm, a regulator, an auditor, an LP, or another agent at organisation Y), there is currently no standardised way for B to verify in real-time that A is who A claims to be, currently authorised to do what A is about to do, in good audit-history standing, and capable of the action. Every interaction either trusts the operator's word or runs ad-hoc due diligence. At institutional scale this doesn't work.

WHAT'S MISSING TODAY

The agent ecosystem in May 2026 has solved several primitive layers but left one critical gap:

Layer	Standard	Status
Agent identity	W3C did:web + Data Integrity Proof	Solved
Agent-to-agent communication	Google Agent2Agent (A2A) v1.0	Solved (Linux Foundation, 150+ orgs)
Agent-to-tool communication	Anthropic Model Context Protocol (MCP)	Solved (broadly adopted)
Agent payment authorisation	Google AP2 V0.1 (FIDO-tracked)	Solved (under formalisation)
Agent counterparty verification	**None**	**Open gap**

AgentPass fills the open gap. It is the institutional layer above all four solved primitives — the standardised real-time verification of an agent's current standing as a counterparty.

CONCRETE: WHAT DOES "VERIFICATION" ACTUALLY ANSWER?

Six questions a counterparty asks before treating an agent as a legitimate party in an institutional transaction:

1. **Who is this agent?** Cryptographic identity, who operates it, what's the chain of custody for the operating identity
2. **What is this agent currently authorised to do?** The mandate scope at the moment of interaction — actions, data classes, entity boundaries
3. **Has this agent been operating in good standing?** Audit chain unbroken, anchored to a public timestamping authority, verifiable without trust in the operator
4. **Is this agent capable of the action it's about to take?** The capability advertisement (extending W3C Agent Card)
5. **What's the agent's track record?** Reputation claims, signed by counterparties or third parties
6. **Are there known violations?** Revocation status (W3C Status List 2021 pattern)

Today, answering all six requires a multi-week paper exercise. AgentPass converts it to a single API call, returning a yes/no with cryptographic proof, in real-time.

The shape — what AgentPass actually is

> AgentPass is a W3C Verifiable Credential — same data model as digital driving licences, academic transcripts, the EU Digital Identity Wallet — applied to AI agent claims. Not a new credential format. Not new cryptography. The novel work is the AI-agent claims schema, the federation pattern, the verification API, and the institutional positioning.

THE CREDENTIAL, IN SKELETON

```
{
  "@context": ["https://w3.org/ns/did/v1", "https://agentpass.ai/v1"],
  "type": "AgentPass",
  "version": "0.1",
  "issuedAt": "2026-05-03T08:00:00Z",
  "expiresAt": "2026-05-03T08:15:00Z",
  "agent": { "did": "did:web:operator.com:agents:w4-drafter", ... },
  "mandate": { "id": "ap2:mandate:c5d8...", "scope": [...], ... },
  "standing": {
    "auditChainStatus": "intact",
    "lastAnchorRef": "btc:block:889234:tx:8a3...",
    "operatingDays": 73, "violations": []
  },
  "capabilities": { "actions": [...], "dataClasses": [...], ... },
  "reputation": { "documentsProduced": 47, "officerCountersigned": 47, ... },
  "proof": { "cryptosuite": "eddsa-jcs-2022", "proofValue": "z3..." }
}
```

15-minute expiry. Signed by the operator's key. Every field cryptographically tied to the artefact the agent is about to produce. The audit-chain anchor reference (`lastAnchorRef`) makes the standing claim verifiable against a public chain that nobody — neither operator, nor verifier, nor Workloft — controls.

THREE DEPLOYMENT PATTERNS, INCREASING MATURITY

- **Self-issued** — operator signs its own AgentPass; verifier checks signature against operator's `did:web` . Day-one shape.
- **Federated** — operator runs an AgentPass authority server that verifying parties call; provides real-time standing query without exposing operator internals.
- **Public ledger** — operators publish AgentPass authority references at a public registry (analogous to FINRA broker-check). Regulators can query any operator's agents directly. The 2028+ shape.

Why now — the 12-month opening

Four signals stack into a single window from May 2026 to mid-2027 that converts AI governance from "good practice" to "procurement-graded prerequisite":

- **EU AI Act high-risk obligations** — Article 6 bites 2 August 2026. Logging and FRIA evidence requirements map to AgentPass-style verification natively.
- **FCA SM&CR-AI guidance** — published end-2026. Names the assurance level expected from senior managers for AI-caused harm; AgentPass is precisely the evidence shape.
- **AM Best ERM-AI scrutiny** — already live; only 24% of rated insurers confident they could pass a 90-day independent review.
- **Five Eyes joint guidance** — published 1 May 2026 by CISA, NSA, NCSC UK + allies. Names cryptographic agent identity, signed mandate scoping, encrypted inter-agent comms, human authorisation gates as the recommended technical controls. AgentPass operationalises every one of them in a single primitive.

Inside that window: regulators need verification mechanisms; framework providers will be looking for partnership-shaped governance layers to embed; institutional buyers will start naming AI-governance attestations in procurement RFPs. **The 12 months ending mid-2027 is when the institutional verification layer of agent infrastructure gets standardised.** Whoever ships a credible spec + reference implementation first becomes the named author. After that window, either it gets done internally by big institutions (proprietary, fragmented) or by a slow standards body (24-month process). The publish-early window is real and finite.

Why this is MCP-scale

> The test for whether a protocol is MCP-scale rather than one-engagement-shaped is whether it holds up across multiple unrelated industries. AgentPass passes that test cleanly.

Use case	What AgentPass enables
Insurance-linked AM	Aeon's W4 Drafter presents AgentPass to AM Best portal before submission upload. AM Best logs the standing-verification in their own audit trail. Defensible end-to-end.
UK council DSAR processing	Citizen DSAR processed by an agent. Council records AgentPass query result alongside the response. Citizen receives evidence the agent had verifiable standing at processing time.
Regulatory examination	FCA / EIOPA / OID examiner queries AgentPass authority during examination to verify ongoing standing of agents producing regulatory submissions. Continuous compliance evidence rather than periodic audit.
LP due diligence	LP evaluating fund manager queries manager's AgentPass authority pre-investment. Verifies AI-driven portfolio operations have verifiable governance posture. Standard part of pre-investment DD by 2027.
B2B agent commerce	Agent A places order with agent B (AP2-mediated). Agent B verifies agent A's AgentPass before fulfilling. Real-time agent-to-agent counterparty verification at commercial scale.
Critical infrastructure / defence	Five Eyes guidance specifically names this sector. Agent verification before any agent acts within a critical-infrastructure perimeter is exactly the AgentPass pattern.

One protocol, six unrelated industries, same primitive. That's the MCP-scale shape.

The 12-week build + adoption sequence

Week	Phase	Deliverable
1–2	Spec V0.1	JSON-LD schema, verification flow, federation pattern. Apache 2.0 GitHub repo.
3–8	Reference implementation	agentpass-py + agentpass-js libraries; reference verifier service; integration with Workloft's existing did:web + AP2 + audit chain.
9–10	First production deployment	Aeon (if engagement signs), as first AgentPass-issuing operator. Live case study.
11–12	Standards body submission	Linux Foundation A2A SIG package + W3C Credentials Community Group package.
	Adoption push	

Week	Phase	Deliverable
Months 4–6		One major framework (Anthropic Claude Agent SDK / Hermes / AutoGen / LangGraph) ships native AgentPass support. Workloft Show episode + LinkedIn launch.
Months 6–12	Formalisation	W3C VC working group adoption or FIDO Alliance recognition.

The moat

Open-source specs don't have legal moats. The moat is reputation — being known as the protocol's authors. When Anthropic published MCP, anyone could fork it; the moat was that "MCP" became shorthand for the standard, and Anthropic's name was attached. Same pattern.

Workloft's moat is being the named author of the spec, the maintainer of the reference implementation, and the first deployment case study. Everyone else implements AgentPass; Workloft *is* AgentPass. That position can't be retroactively taken once the spec is published with our name on it.

Secondary moat: the relationship one. Spec authors become the natural partner for institutional buyers, regulators, platform vendors who want to ship AgentPass-compatible tooling. *"Talk to Workloft, they wrote the spec"* — the same position MCP gave Anthropic in agent tooling, AgentPass would give Workloft in institutional agent governance.

The bet: three months of focused build (spec + reference impl + first deployment + standards-body submission) plus three to six months of adoption work, against a 12-month window where this layer of the agent ecosystem will get standardised one way or another. If we ship credibly inside the window, Workloft becomes named author of the institutional verification layer of agent infrastructure — a position that compounds for the rest of the decade. Outside the window, either someone else ships first (loses moat), or institutions roll their own (no standard, fragmented). The window is real, narrow, and closing. The technical risk is low because we compose existing W3C primitives. The execution risk is real and worth the bet.

Companion artefacts: V0.1 technical spec at `/home/workloft/strategy/agentpass-spec-v0.1.md` · 18-task Gary build list logged 3 May 2026 · Vera-stress-test brief embedded in V0.1 spec for execution once OpenRouter credit restored.

Open name decision: AgentPass (current working name) vs Open Counterparty Protocol (OCP) vs Agent Standing Protocol (ASP) vs Workloft Standing Verification (WSV) — Alfred decision, locks in for spec, repo, domain, marketing.

Workloft.ai LTD · ICO C1912528 · UK-incorporated · alfred@workloft.ai