Θ + π

TETA+PI: Trust Infrastructure for Digital Entities

Whitepaper v1.1

Authors: Bob V. (Founder & CEO), Mykhailo M. (CTO & Advisor)

tetapi.dev · github.com/teta-pi · MIT License

Abstract. The internet is entering a phase in which AI agents, not humans, increasingly decide what to read, buy, cite, and trust. No public infrastructure currently lets an agent verify who — or what — it is actually interacting with. TETA+PI is an open, cryptographically-verifiable trust registry for digital entities: people, businesses, brands, domains, APIs, AI models, MCP servers, and AI agents themselves. Critically, TETA+PI does not verify that an entity merely exists or is legally registered — it verifies specific claims an entity makes about itself, anchored to evidence captured at the moment of creation. Entities are verified against official registries and content-provenance standards (C2PA), with every verification event anchored to Bitcoin OpenTimestamps. On top of the registry, we introduce TWIRA (Trust-Weighted Intent Routing Algorithm), a ranking method that routes agent queries to verified entities using a combination of trust, semantic intent alignment, and provenance integrity. We describe the registry model, the claims-based verification architecture, TWIRA, and a structural property we call the Temporal Moat — an append-only chronology of verification events that cannot be reconstructed retroactively by a later entrant. This paper states plainly what is implemented today, what is proposed and awaiting evaluation, and what remains an open research question.

July 2026

1. The Problem: Trust Has No Infrastructure

Every major layer of the emerging agent internet has a standard except one. Model Context Protocol (MCP) and Agent2Agent (A2A) standardize how agents connect. C2PA and similar provenance standards address what content is authentic. Payment and transactional-trust systems — including Stripe's Agentic Commerce Protocol (ACP) and Shopify's Universal Commerce Protocol (UCP) — are emerging for how agents pay. None of these answer a simpler, prior question: who, or what, is actually on the other end of the connection, and is what it claims about itself true?

An agent that fetches example.com/.well-known/agent.json learns only what that domain's owner chose to publish about itself. There is no independent party confirming that the declared identity is accurate, current, or even real. This is the same problem the web faced before certificate authorities: any server could claim to be any domain, and a browser had no way to check. SSL/TLS solved it for connections. Nothing has yet solved it for entities and their claims.

Notably, the largest commerce platforms building agentic infrastructure in 2025-2026 confirm this gap by omission rather than by solving it: Shopify's Catalog syndication auto-enables product discovery for AI agents across over a million merchants at the platform level, with no individual merchant opt-in required. A store is trusted by these systems only because it runs on Shopify/Stripe infrastructure — not because any independent party has verified who operates it or whether its claims about its own products are accurate.

The scale of the gap is already large and growing quickly. The AI agent market was estimated at $7.84–11.5B in 2025–2026 across independent research firms, projected to reach $50–53B by 2030 (45–46% CAGR). The official MCP registry alone recorded 9,600+ server entries by May 2026, up from essentially zero eighteen months earlier. At the same time, over 60% of enterprises report they do not fully trust AI agents in production, and an estimated 40% of agentic AI projects are at risk of cancellation due to weak governance and unclear accountability — not due to model capability.

Regulation is accelerating the need for an answer. The EU AI Act's Article 50 (Regulation (EU) 2024/1689) mandates machine-readable marking of AI-generated or AI-assisted content from August 2, 2026, with fines up to €15M or 3% of global turnover. Meaningful compliance requires exactly the kind of provenance-anchored, verifiable identity that does not yet exist as public infrastructure.

Trust, in other words, has become a missing layer — not a missing feature.

2. Positioning: Trust Infrastructure for Digital Entities

TETA+PI is not an AI registry. It is trust infrastructure for digital entities, of which AI agents are the first urgent class — not the only one.

We deliberately scope “digital entity” broadly: a person, a company, a brand, a domain, a website, an API, an AI model, an MCP server, a piece of software, a code repository, an AI agent, and future classes of autonomous entity not yet named. Each of these can be verified, and each benefits from the same underlying registry mechanics.

Precedent	What it verifies	What TETA+PI verifies
ICANN / domain registries	Who owns a name in DNS	Who a verified entity actually is
CA/Browser Forum (SSL)	Which server a certificate belongs to	Which real-world entity a digital identity belongs to
GLEIF (Legal Entity Identifier)	Legal entity identity for finance	Any entity's identity for the agent internet

A registry is the first product. The infrastructure — a durable, neutral layer other systems build on rather than compete with — is the actual company. Large platforms rarely build neutral infrastructure that must remain trustworthy to their own competitors. TETA+PI's goal is not to compete with any single platform. The goal is to become useful even to them.

3. The Registry Model

3.1 Entities and blocks

A verified profile in TETA+PI consists of an entity and a set of blocks. An entity has a type (person, business, organization, brand, domain, website, API, AI model, MCP server, software, repository, or AI agent) and a visibility flag. A block is a discrete unit of content — text, a file, or video — attached to an entity, each independently verifiable and independently made public or private.

Privacy is a first-class primitive: an entity or an individual block can be verified while remaining unlisted and unindexed. A journalist can prove authorship of a document without publishing their identity to the open web.

3.2 Segmentation

Builders — developers, MCP server authors, AI-native startups who integrate the protocol.

Operators — businesses and organizations running verified operations; the primary paying segment.

Consumers — individuals establishing verified identity; always free.

3.3 No forms, no categories

The registry does not ask an entity to self-categorize. Given a name and a set of blocks, the system structures, categorizes, and ranks the entity automatically, cross-checked against official sources where available — one verification method among several described in Section 4.3.

3.4 Brand and legal entity are not the same thing

A significant design decision follows from a simple observation: an entity's public brand and its legal registration frequently do not share a name. Google is legally Alphabet Inc.; countless brands operate under names distinct from their registered business. Requiring an exact registry name-match at entity creation would incorrectly reject or block a large share of legitimate entities.

TETA+PI therefore separates entity creation from registry verification entirely. Any brand name can be created immediately, free, with no registry check (Level 0 — existing but unverified). A brand entity can optionally be linked to a verified legal entity via legal_entity_id — a disclosed, public relationship through which the brand inherits verified trust from its legal parent.

---

4. Verification: What We Actually Confirm

4.1 The core claim

A common misreading of the product is that TETA+PI verifies an entity exists. It does not — or rather, existence is one weak, optional signal among several (Section 4.3). What the system verifies is narrower and considerably more useful: whether a specific claim an entity makes about itself is authentic and matches reality, captured and checked at the moment of creation.

A tailor claims “we sew with high-quality fabric” — PI Camera captures the sewing process itself. A sushi chef claims “traditional recipe” — the camera captures the preparation, without revealing proprietary technique. An artist claims authorship of a painting — the camera captures the painting process, anchored to their existing verified profile. A journalist claims footage is real, unstaged coverage — PI Camera proves the content is not staged and not AI-generated.

This removes dependence on legal status entirely. A freelancer or home-based creator with no registered business can prove the quality and authenticity of their work exactly as a large company can — verification attaches to the claim being made, not to the legal form of the entity making it.

4.2 Two independent checks per block

Check	Question it answers	Mechanism
1 — Authenticity	Is this content real, captured just now, by a human?	C2PA signing + Bitcoin OpenTimestamp anchor
2 — Claim Match	Does the content actually support the block's text claim?	Multimodal model reviews video/audio vs. claim text; returns yes/n

A block that passes both checks carries materially higher weight in TWIRA's ranking (Section 5) than a text-only claim with no corroborating evidence. Authenticity (Check 1) is implemented today via the PI Camera pipeline (Section 4.4). Claim Match (Check 2) is specified and architecturally straightforward — orchestration of existing third-party multimodal models via API, not a novel model requiring original research — but is specified, not yet implemented.

Multimodal models make mistakes, particularly on nuanced or technical claims. Claim Match is designed and will be presented as an AI-assisted signal with an appeals and manual-review path — never as an infallible verdict. A visible public failure here would damage confidence in the surrounding trust infrastructure more severely than almost anywhere else in the system.

4.3 Verification methods

Registry matching is one of several independent, optional methods by which an entity can be verified — not the sole path:

Method	Mechanism	Trust weight
Official Registry Match	Legal name cross-checked against Handelsregister, GLEIF, EU VAT, etc.	Highest
Business Email Control	Magic-link confirmation to an email on the entity's own domain	High
Domain Ownership	DNS TXT record or file-based check	High
Document Upload	Registration certificate, license, or tax ID	Medium — UI specified, backend deferred until file-handling

4.4 Verification levels and PI Camera

Level	Name	What it requires
L1	Registry Attested	Entity registered; cross-checked against public official registries
L2	C2PA Verified	Content blocks signed via C2PA at capture, anchored to Bitcoin OpenTimestamp
L3	Continuous	Live, periodically-probed API endpoint under continuous verification

PI Camera is the capture mechanism behind L2 and behind the two-check model above. It signs photos and video with a C2PA manifest at capture, then anchors that manifest to a Bitcoin OpenTimestamp. It does not verify location — it verifies authorship of the capturing entity: this specific verified entity produced this specific content at this specific, cryptographically-anchored time.

C2PA itself is an open standard maintained by the Coalition for Content Provenance and Authenticity, adopted by Adobe, Google, OpenAI, and Microsoft, and is the mechanism the EU AI Act's Article 50 points to for machine-readable AI content marking. TETA+PI's contribution is the registry that anchors C2PA-signed content to a verified entity and its specific claims, and the algorithm (Section 5) that uses that anchoring as a ranking signal.

5. TWIRA — Trust-Weighted Intent Routing

5.1 Motivation

A registry alone answers “is this entity verified?” It does not answer: given a query, which verified entities should I surface, and in what order? TWIRA (Trust-Weighted Intent Routing Algorithm) combines three independently-computed signals.

5.2 Formal definition

```
TWIRA(e, q) = α·T(e) + β·I(e, q) + γ·P(e)
```

T(e) — Trust Score:

```
T(e) = Σ_v  level(v) · decay(v) · source(v)
```

summed over verification events v, level(v) ∈ {0.4, 0.8, 1.0} for L1/L2/L3, decay(v) = e^(−λ·t), and source(v) weights each verification method from Section 4.3.

I(e, q) — Intent Alignment Score:

```
I(e, q) = max_i [ cos_sim(embed(q), embed(block_i)) · evidence_weight(block_i) ]
```

computed via pgvector ANN, where evidence_weight is highest for blocks that passed both checks in Section 4.2, lower for authenticity-only blocks, lowest for unsigned text-only claims.

P(e) — Provenance Integrity Score:

```
P(e) = f(btc_timestamp_depth(e), c2pa_chain_length(e), endpoint_uptime_30d(e))
```

The weights α, β, γ are learned parameters. At launch they are static; the intended path is to fit them against real query-and-selection data once production traffic exists.

5.3 What is implemented vs. proposed

Status	What it covers
Implemented today	Entity/block schema, C2PA signing via PI Camera, Bitcoin OpenTimestamps anchoring, working MCP ser
Specified, in progress	Full TWIRA scoring pipeline, Claim Match check (Section 4.2)
Proposed, not yet validated	That the functional forms (decay, P combination, I evidence-weighting, static weights) are optimal. No b

5.4 Why this is not easily replicated

The formula is public. What is not reproducible is the data it depends on. btc_timestamp_depth and c2pa_chain_length are properties of past events — a competing registry cannot backdate its history to match. The weights can only be tuned against real query logs, which do not exist until a registry has real traffic — the same asymmetry separating public PageRank from Google's query data.

A well-resourced competitor with a comparably large verified-entity base and comparably deep historical timestamps could, in principle, replicate TWIRA's effectiveness. The claim is that doing so requires reproducing accumulated history, not merely reading a paper.

6. The Temporal Moat

Every verification event — registration, signing, claim check, endpoint reverification — is written to an append-only log and anchored to Bitcoin. This produces a chronology of when verification happened, that cannot be rewritten or backdated.

1. First-verified advantage. A cryptographic fact anchored to Bitcoin, publicly auditable.

2. Evidentiary weight. Independent of whether TETA+PI continues to operate.

3. Feeds TWIRA directly. btc_timestamp_depth is a direct input to P(e).

TWIRA is the algorithm. The Temporal Moat is the data that makes the algorithm work. A competitor can reimplement the formula. They cannot retroactively acquire a longer verified history.

7. Product Evolution and Discovery Engineering

#	Stage	When	Description
1	**Digital Entity Verification**	today	Any entity establishes verified identity + attaches verifiable claims
2	**Identity + Reputation**	next	Verified entities accumulate a provable, portable history
3	**Discovery**	next	Entities found through trust, not advertising spend
4	**Agent Economy**	horizon	AI agents transact directly with verified entities
5	**Autonomous Commerce**	horizon	Entities transact without continuous human mediation

Discovery Engineering: making digital entities discoverable, understandable, and selectable by autonomous AI systems, distinct from SEO (human search) or advertising (human attention). TWIRA is its first concrete implementation.

Stages 4 and 5 are an inevitable horizon, not today's product or pitch.

8. Distribution: Meeting Entities Where They Already Are

Verification value depends on registry density, which depends on how easily an entity can create a presence. TETA+PI ships integration paths matched to where entities already operate:

WordPress plugin — roughly a third of the web; generates machine-readable descriptors on install, verification is an optional upgrade, not a gate.

Shopify / Wix integrations — same free-first model, chosen for scale and early adoption of agentic-commerce standards.

Universal script tag — comparable to a web analytics snippet, covers sites with no CMS; paired with an optional one-time DNS record for full file coverage, since a script alone cannot serve files at /.well-known/.

Each path converges on the same registry and verification methods (Section 4.3) — the entry point differs, the trust model does not.

9. Governance and Openness

Stays open	Stays commercial
Core verification specification	Enterprise services and delegated CA
Core trust model (transparent, auditable)	Premium APIs and bulk verification
Protocol (MIT licensed)	White-label deployments
Root CA architecture (publicly documented)	Optional modules (one-time upgrades)

Governance is intended to become more independent of the founding company over time, in the gradual pattern ICANN followed. We do not claim this independence exists today.

10. Conclusion

The agent internet already has protocols for connectivity (MCP, A2A), content provenance (C2PA), and payment authorization (ACP, UCP). It does not yet have public infrastructure for the prior question: whether the entity on the other end is who it claims to be, and whether what it claims about itself is true. TETA+PI is built to be that infrastructure.

TWIRA's formula is public; its practical superiority over simpler baselines is a claim we intend to substantiate with real evaluation data, not assert in advance. Claim Match verification is specified but not yet in production, and will ship with a human-review path rather than as an unqualified AI verdict. The Temporal Moat is a structural, falsifiable property — not a claim of invulnerability.

The registry is open source (MIT) at github.com/teta-pi. The MCP server is live at tetapi.dev/mcp. We welcome scrutiny, replication attempts, and correction.

Changelog

v1.1 — July 2026. Added Section 4, reframing verification around claims rather than mere existence, including the two-check-per-block model and the registry-vs-brand distinction (3.4). Added Section 8 on distribution paths. Updated Section 1 to reference ACP/UCP.

v1.0 — July 2026. Initial public release.