

An Agentic Internet

yesh.eth
Magali Morin

2025-11-14

<https://agenticinternet.tech/>

1 Summary

Our proposal re-imagines how the web can contribute to a shared sense of truth.

Today the internet is made of silent pages waiting to be interpreted by search engines. In the Agentic Internet, every website becomes an active participant in knowledge, not just a passive document. Each document becomes an AI agent that publishes a verifiable confidence score about its claims, track evidence, compare itself with others, and update its confident score as new information appears.

A confidence score is a numerical estimate that expresses how strongly an agent believes a specific claim is correct, based on the available evidence and the consistency of that claim with related sources.

At the core of this vision is a missing layer the internet never had, a reasoning layer. AI agents continuously evaluate their own claims against others and revise confidence as evidence shifts. This transforms the internet into an active, self-interpreting and verifiable knowledge network.

This report outlines the problem, the proposed solution and the key principles of the technical implementation.

2 Problem Statement

As a society, we don't share the same base understanding of truth anymore. Search engine results are personalized, feeds are fragmented, and information spreads faster than it can be verified. The result is a world where people no longer agree on facts not because truth changed but because our information infrastructure never built a shared epistemic foundation.

2.1 The Current Web Landscape

The modern web is a vast archive of documents designed to be searched but not self-interpreted. Search engines rank documents based on relevance signals and answers are shaped by these rankings rather than by the documents content themselves.

Research evolves but documents do not inherently update their claims or publish revised confidence levels. There is no built-in mechanism for a document to say “my confidence in this claim has changed” or “new evidence updates this claim.”

Most importantly, the web lacks any distributed reasoning layer. There is no mechanism for documents to evaluate one another’s claims, detect contradictions or converge on updated understanding.

2.2 Consequences of lacking a reasoning layer

- **Epistemic fragmentation:** Different documents drift apart because none of them can update based on new research or contradictions discovered elsewhere.
- **Misinformation persists:** Older claims remain online unchanged and uncontextualized, even when falsified.
- **No convergence:** There is no mechanism for documents to coordinate or align their understanding of a topic.

3 Proposed Solution (Non-Technical Overview)

Imagine every website or paper behaving like a small AI agent that understands its claims, rates how confident it is, updates that rating and shares it in a way anyone can verify.

- **Continuous validation:** Each agent regularly re-evaluates its claims as new evidence, research, or peer signals appear.
- **Provenance tracking:** Every claim and every confidence update is linked to clear metadata showing where information came from and when it changed.
- **Cryptographic binding:** A public key associated with each agent allows anyone to verify that confidence scores updates actually come from the same source and haven’t been tampered with.

4 Technical Details

4.1 High-Level Architecture

Each document becomes an AI Agent. On start up the AI agent:

1. Creates a unique identifier for the document/paper/website (a DID or a content-derived hash).
2. Generates a cryptographic public-private key pair. This key can be derived from the document identifier.
3. Registers this binding on-chain in a smart Contract registry to guarantee verifiable identity and provenance.
4. Instantiates a lightweight LLM that represents the agent and interprets the document’s claims.

5. Discovers peer agents and communicates with them.
6. Computes a confidence score and updates it periodically.
7. Stores provenance metadata.

4.2 Five-Layer Distributed Agentic Internet

Agentic means having the ability to act autonomously: to interpret information, make decisions, initiate communication, and update internal beliefs without external prompting.

1. **Extraction Layer:** Crawls and extracts structured claims.
2. **Claim Layer:** Agents use standardized confidence score related to their claims.
3. **Exploration Layer:** Agents discover other agents working on related topics thanks to an embedded graph. An embedded graph is a structured representation of concepts, topics, or claims that allows agents to discover related content. It acts as a map that agents use to find relevant peers.
4. **Messaging Layer:** Agents exchange structured messages with other agents using a shared communication protocol.
5. **Reasoning Layer:** Agents compare and update beliefs. Each agent independently interprets peer claims and recomputes its own confidence score.

5 Conclusion

By turning each website and document into an autonomous, verifiable, reasoning agent, the web gains a shared epistemic foundation. This creates a transparent, evolving and machine-verifiable knowledge ecosystem.