



# Hashtag.ai

The Persistent Memory Layer.



# AI Systems Forget. Workflows Cannot.

Inconsistent workflow behavior across interactions.



Fabricated or unreliable references that sound convincing but fail scrutiny.

Total loss of organizational memory and policy context.

# Trust Breaks Down at the Workflow Level.



**Ella** (Corporate Lawyer)

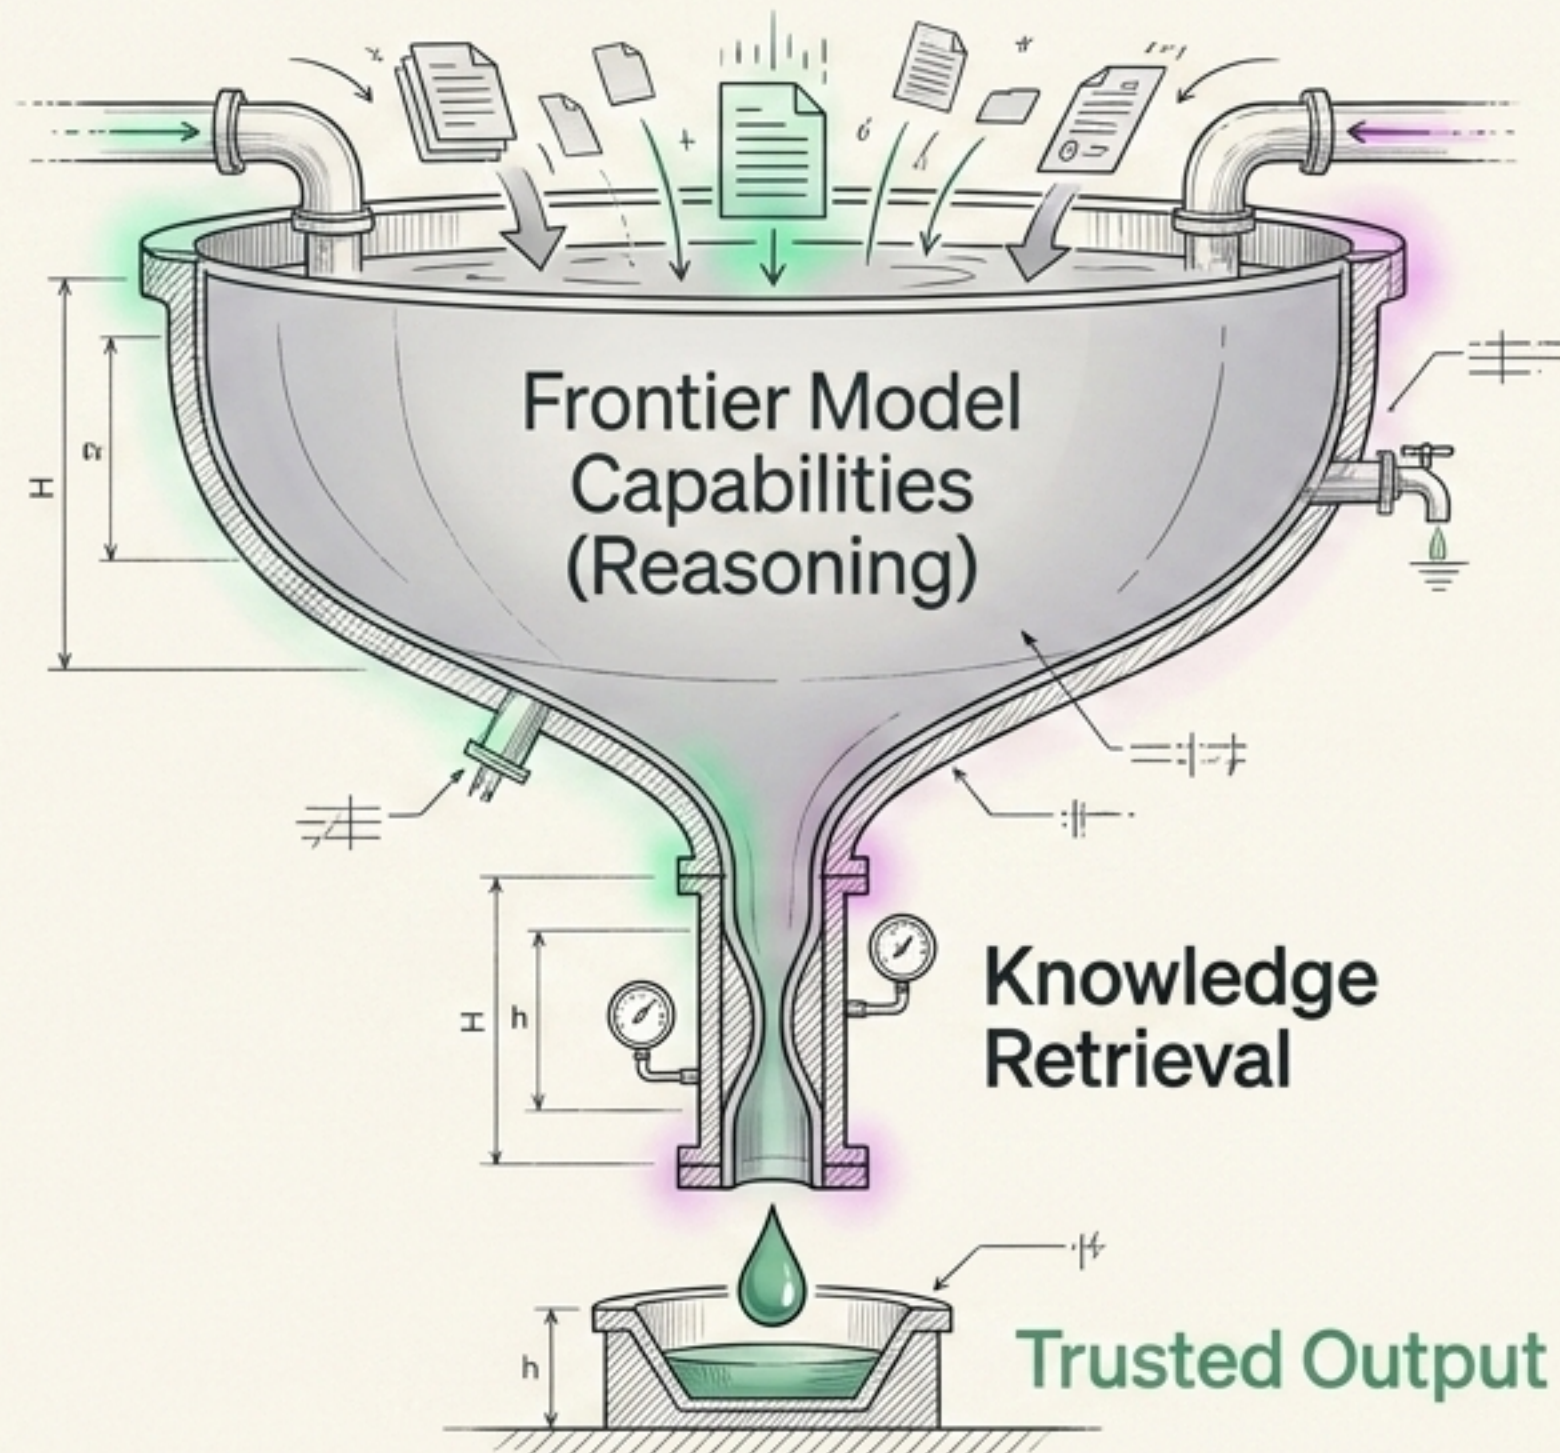


One fabricated citation invalidates an entire process.

Sounding correct is no longer enough for regulated environments.

When part of the workflow is wrong, trust collapses.

# The Bottleneck is Retrieval, Not Reasoning.

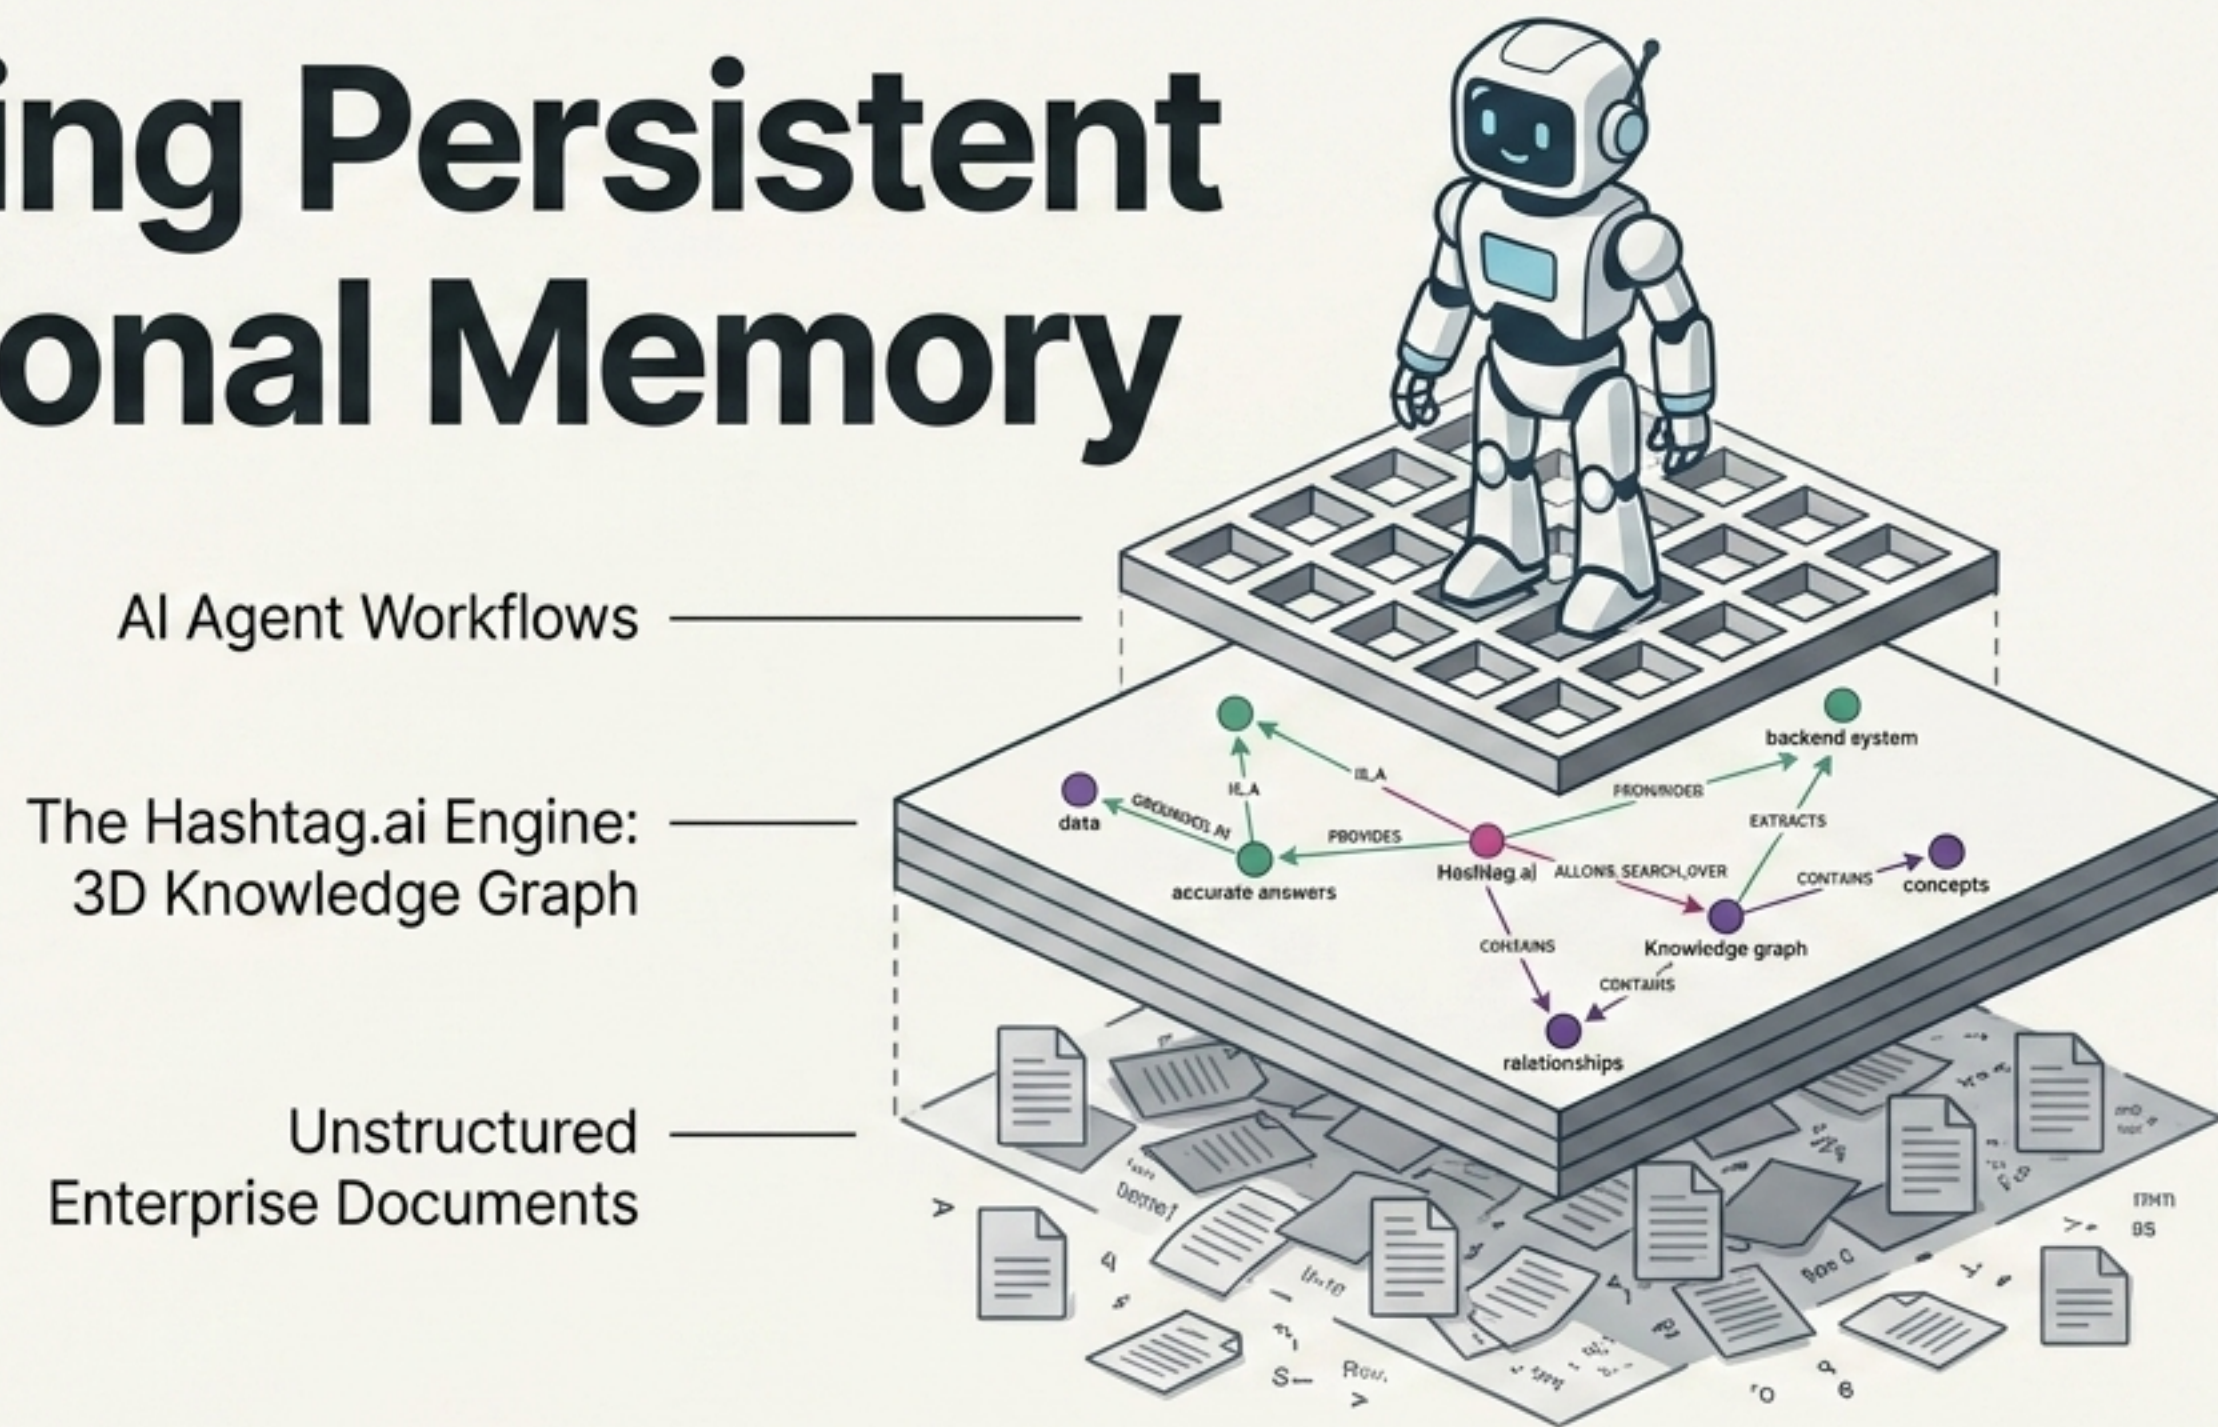


Frontier models are already highly capable at complex reasoning.

The true point of failure is reliable memory across isolated workflows.

We are building the critical infrastructure layer to fix this bottleneck.

# Constructing Persistent Organizational Memory



- Ingest raw data to capture complex workflow relationships.

- Construct structured organizational memory, not just text matches.

- Enable workflows to retrieve grounded context securely.

# Why Existing Approaches Fail to Scale.

|                           | Vertical Apps                        | DIY Infrastructure                         | Hashtag.ai Platform              |
|---------------------------|--------------------------------------|--|----------------------------------|
| Context Continuity        | Limited to small isolated workspaces | Scattered across models                    | Persistent cross-workflow memory |
| Implementation Complexity | Low                                  | Extremely High<br>(Neo4j, Open-source RAG) | Low<br>(Out-of-the-box layer)    |
| Scalability               | Poor                                 | Hard to maintain                           | Highly scalable                  |

# Built for High-Trust, Regulated Workflows

Legal

Finance

Insurance

Compliance

Hashtag.ai Infrastructure Layer

High-trust environments cannot tolerate **hallucinations** or **isolated logic**.

**Unreliable retrieval** carries massive financial and compliance costs.

These sectors **strictly** require **persistent context** and **traceable retrieval**.

# Immediate Commercial Validation.

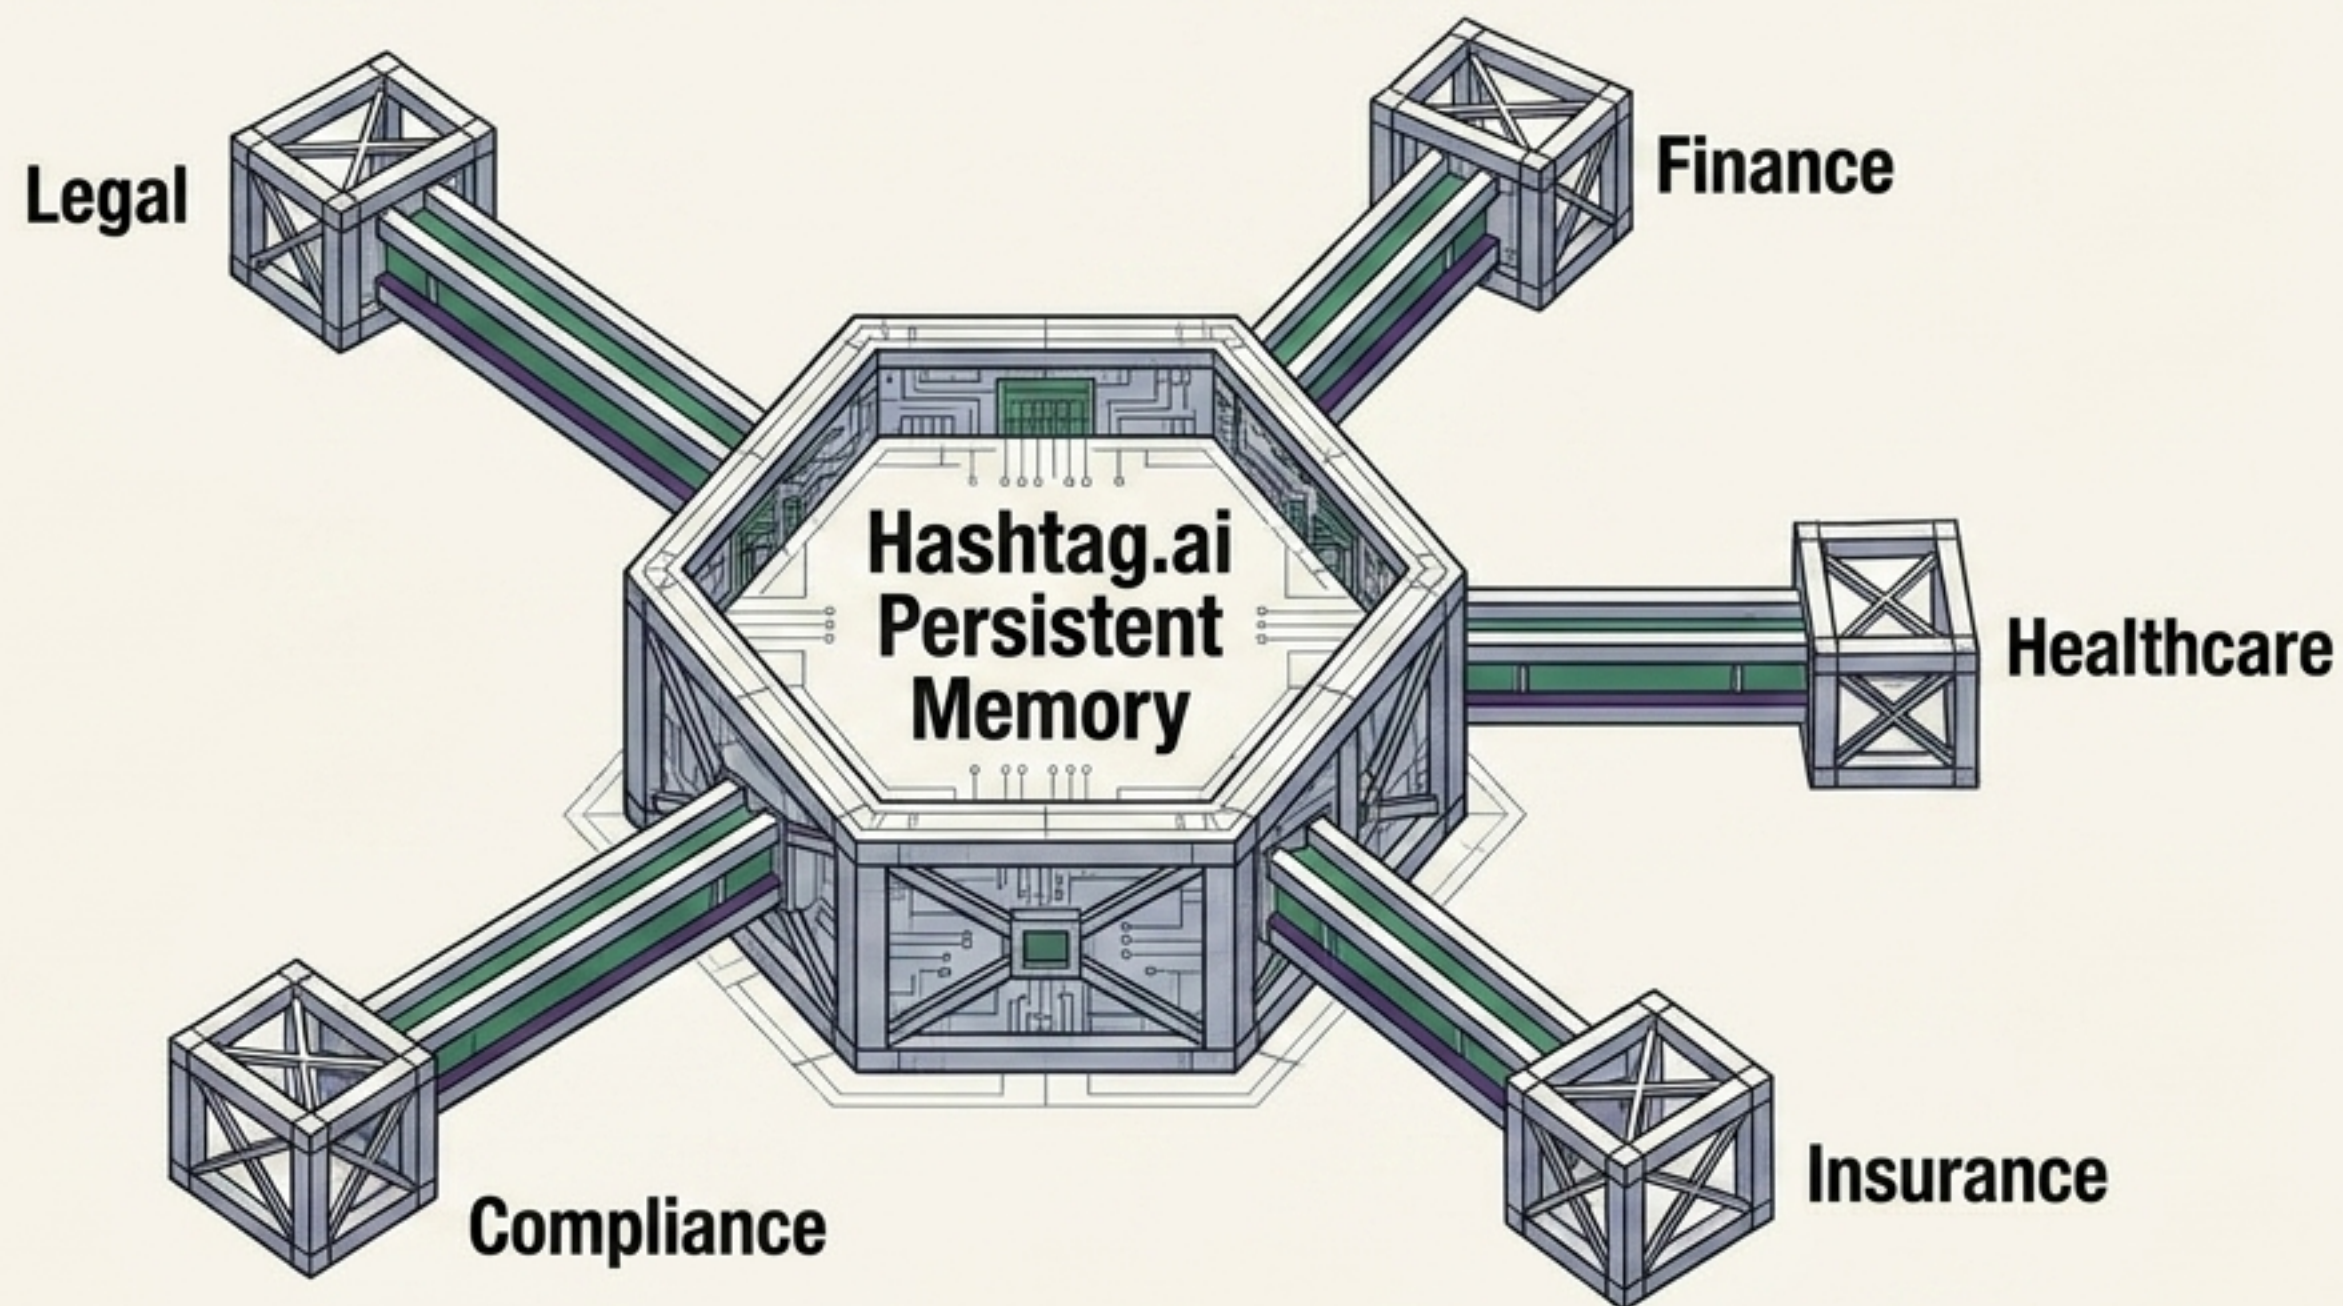


**\$10,000** / month  
pilot agreement secured.

**Week 2**  
of Luminare program.

- Deployed custom memory infrastructure for Australian tax workflows.
- Actively engaged in additional enterprise pricing discussions.

# A Reusable Core Across Industries.



The specific workflow interface changes depending on the industry.

The underlying persistent memory infrastructure remains perfectly reusable.

Creating a highly defensible, horizontal platform across regulated sectors.

# The Architects of Persistent AI.

## Dr. Angie Simmons (CEO)

Led a \$306,000 deep-tech research project funded by the Office of National Intelligence.

## Dr. Anupam Chaudhuri (CTO)

Saved fintech firm Quantaco \$1M/year by architecting and replacing third-party tools with custom LLM systems.

## Dr. Tuan Dung Lai (Advisor)

Swinburne computing lecturer scaling tech education to a massive audience of 448k+ subscribers.



# Give Your AI Workflows a Memory.

Experience persistent organizational memory | Try the platform infrastructure today | Connect with @Angie Simmons on LinkedIn