
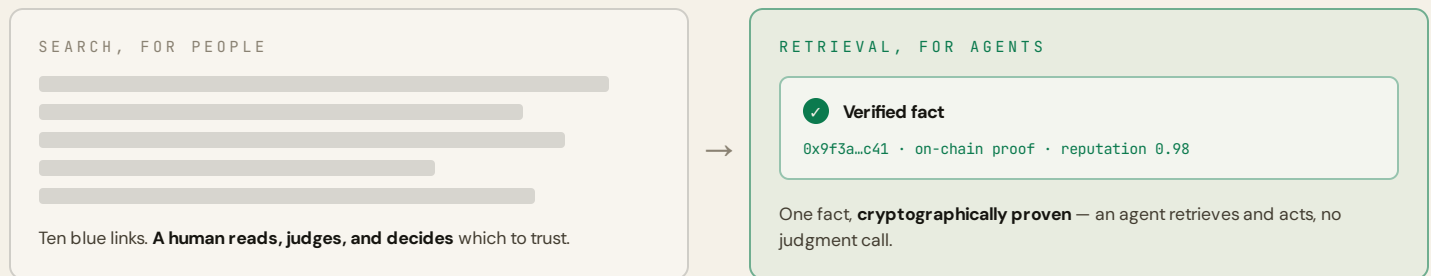


THE FUTURE OF SEARCH IS NOT FOR PEOPLE

Google indexed the web for people. We retrieve *verifiable facts* for AI agents.

Agents don't want ten blue links to read and judge. They want one fact they can **prove** — authentic, current, and true — returned in **~100 ms**, not re-derived from scratch. That is a different product than search has ever been, and it is the layer every autonomous agent will depend on.



One answer. *Ten thousand times over.*

One rep's agent researches a customer Monday; another does it Thursday; ten thousand more do the same work across the web.
Every agent re-derives an answer that already exists — paying full price, in tokens and in time, every single time.

THE WASTE, QUANTIFIED

300×

faster. ~30 sec to re-derive vs **~100 ms** to retrieve a verified fact.

80M

tokens to answer one question across 10,000 agents — vs **8,000** computed once.

10,000×

redundant. The same compute, billed again and again, team after team.

THE BUSINESS MODEL — WE CHARGE 20% OF THE SAVINGS

USERS KEEP 80% — AND ALL THE SPEED

OUR FEE 20%

*80% · the customer**20% · SupraCognition*

We capture a fifth of the value we create and hand back the rest — aligned with the customer on every single query. With **1 billion+ AI agents forecast by 2029** (IDC), that fifth is a multi-billion-dollar business that compounds as every lookup gets cheaper and faster.

THE SOLUTION

A verifiable memory layer — *the speed of a cache, the trust of a chain.*


SupraCognition pairs **real-time, blockchain-verified data** with persistent, provenance-tagged memory. When an agent does real work, the result is stored, scored, and provenance-tagged — so the next agent **retrieves** it instead of recomputing it. A classifier routes every question to the right memory mode; **MemoryRank** scores and decays facts over time; and a hash-chained audit anchors every claim to our own Layer-1. One lookup returns a fact **and** the proof of where it came from, what was known when, and whether anything has changed since.

WHY IT NEEDS A BLOCKCHAIN

The hard requirements — **permissionless participation, censorship resistance, credible neutrality, and trustless micropayments with no trusted operator** — are exactly what a permissionless ledger provides. It is load-bearing, not decorative. The performance objection dissolves in the architecture: **only hashes and metadata live on-chain**; the data is served from caches with an on-chain proof attached. The efficiency win and the trust win are the same win.

THE AI ORACLE PROBLEM — SOLVED IN PRODUCTION

Every AI faces an **oracle problem**: how does a model get real-world facts it can trust, when it is frozen at training time and prone to confident guesses? Our answer is already deployed. **DORA / Threshold AI Oracles** have secured financial data in production for ~2 years on this exact hybrid design. We are extending a working oracle to all of agent memory — not inventing one on a slide.



THE OLD WAY VS. VERIFIED RETRIEVAL

PER IDENTICAL ANSWER	RE-DERIVE FROM SCRATCH	SUPRACOGNITION
LATENCY	~30 sec, every time	~100 ms — 300× faster
COST · ONE ANSWER × 10,000 AGENTS	80,000,000 tokens	8,000 tokens, computed once
TRUST	unverifiable guess	cryptographic, on-chain proof
PROVENANCE	none	full source lineage & timestamp
NET RESULT	re-bought every single time	proven once, reused by all

PROVEN · IN PRODUCTION


Not a thesis on a slide — *a live, funded substrate.*

<i>\$45M</i>	raised across 30+ institutional investors — Coinbase Ventures, Animoca, HashKey	<i>Live</i>	Layer-1 mainnet shipped Nov 2024; 20+ peer-reviewed papers
<i>~2 yr</i>	DORA oracle in production, securing financial data on-chain	<i>+16pp</i>	memory benchmark vs Claude 3.7 Sonnet baseline (MAB v3, matched harness)
<i>CCS</i>	ACM CCS Distinguished Paper — consensus & verification research	<i>217B</i>	agent actions/day forecast by 2029 (IDC) — the market we index

WHY NOW · THE MOAT

As frontier models commoditize, durable value migrates to the **memory-and-verification layer** around them. Models are interchangeable and rented by the token; **accumulated, provenance-tagged memory compounds and cannot be copied**. The model is rented — the memory is owned. Whoever holds the trusted, reusable knowledge layer for agents holds the position search held for the human web.

THE OFFERING — CONCENTRIC, NOT A SINGLE FAR-OFF BET



- **ENTERPRISE MEMORY · THE FLOOR**
Deployed locally; agents retrieve instead of recompute. **An ROI sale**, in token spend.
- **KNOWLEDGE LEDGER**
Private audit chain for provenance & explainability. **A risk-reduction sale**.
- **KNOWLEDGE FEDERATION**
Privacy-preserving sharing across trusted orgs. **A network-effect sale**.
- **GLOBAL KNOWLEDGE LAYER · THE UPSIDE**
The public, provenance-backed substrate. **The platform**.

THE CATEGORY

This is the “*Google search*” for AI — verified facts retrieval, in one query, instant and trusted. The enterprise floor pays for it today; the public network is the upside.

The public network is the upside.
The enterprise business is the floor.

THE ASK
Raising to scale enterprise deployments and the public verification network. Round details, traction, and unit economics shared on request.