

Kirha

Seed tl;dr (Feb 2026)

Kirha is a usage-based "context layer" for AI agents: a single API that routes queries to premium/private data providers (Insurance, News, Cybersecurity, Adtech, Crypto, Medical, Finance) and returns fresh, provider-verified outputs. Kirha operates a cluster of MCP servers that expose premium data providers API - it does not index any data without a request. With the search market shifting toward AI "generative answer engines"; Gartner forecasts traditional search volume drops 25% by 2026 as chatbots/assistants substitute for classic search.

Fund: Fund I
Expected Close: N/A
Current Series: Seed

Previous Raise: Yes (Pre-seed)
Terms: TBD - Contact Team
Allocation: TBD

Current Problem?

AI has evolved from being a "copilot" to being embedded in workflows. As more workflows become agentic, agents increasingly fail on data access: they need fresh, reliable, auditable context beyond model training data. Kirha frames this as the "context gap" in real-time work.

Why now?

1. Web search is messy for agents: inconsistent, often outdated for niche domains and can be token-inefficient at retrieval time. Kirha's public benchmark argues for a meaningfully higher output quality with fewer tokens vs a web-search baseline.
2. Premium data is fragmented: locked behind subscriptions or enterprise contracts and slow integrations are a poor fit for autonomous agents that need pay-per-call access.

How does Kirha solve this?

Kirha synchronises and aggregates data for agent consumption. It sits in the middle of being a data and search marketplace providing Context-as-a-Service via:

- Unified Search API + SDK: one endpoint to query premium providers.
- Deterministic tool planning + routing: for generalistic agents enables predictability of costs and auditability and the reviewability.
- Micropayments + auditability: integration of x402 to pay-per-resource will enable agent-native payments for data that are auditable on the blockchain.

Who are the users/customers?

Users: AI agents, individuals, businesses.

Monetisation: Usage-based credits + tiered plans (Free + \$20/mo + \$50/mo + Enterprise).

Enterprise: Additional features such as SSO + on-prem deployment.

Developers: API - can bring their own models (further privacy opportunities).

Why will people use Kirha?

Benchmark: Kirha used a LLM-as-Judge to cross-reference scores against human evaluation. It outperformed competitor Exa and web search baseline across 100 domain-specific queries; with 95% fewer tokens injected vs web-search (233k vs 4.6M).

DAU: ~2,000 (last 90d), with roughly 50 new users registering every day - company-reported.

Agents: High reliability, lower token cost and blockchain rails enable pay-per-call.