Kirha
IC Memo (Feb 2026)

Kirha is a usage-based "context layer" for AI agents: a single API that aggregates a query to premium data providers across Insurance, News, Cybersec, Adtech, Crypto, Medical and Finance to the most premium and latest outputs.

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

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

KIRHA

TABLE OF CONTENTS

---

1. INVESTMENT THESIS

1. Kirha is well-positioned to shape the gateway for premium data: Kirha is building the access layer that makes premium data usable inside agentic systems. By operating a cluster of MCP servers and abstracting provider complexity (authentication, rate limits and pricing), Kirha can become the default "gateway" through which agents pull licensed, high-signal information. If agents increasingly become the interface for research and decisioning, the platform that sits in the paid data path (standardising access and packaging outputs for machines) becomes a structurally advantaged control point.

2. Context is the new moat in the agentic era: As AI chatbots and agents become a front door to information, reliable retrieval from licensed and premium data sources becomes mission-critical. Model capability is no longer the bottleneck; trustworthy context is. Agents that are deployed in real workflows need outputs that are current, defensible and traceable back to high-quality sources especially in domains like finance, cyber, insurance and medical.

Kirha's plan-first execution (deciding the provider/tool sequence upfront) is a wedge because it turns retrieval from a messy, token-heavy loop into a predictable, cost-controlled workflow.

3. Strategic alignment with portfolio: The fund is positioned to distribute Kirha and accelerate its network effects. Agent-native portco's create built-in distribution channels where "premium context" becomes a monetisable primitive - sold as pay-per-call tools or agent services within marketplaces. The result is a portfolio flywheel: portfolio adoption drives volume and provider coverage for Kirha, while Kirha raises the quality, speed and auditability of portfolio companies' agent workflows, making their products more reliable and defensible.

---

2. PRODUCT OVERVIEW

Kirha is a usage-based "context layer" for AI agents. It provides a single API that routes an agent's query to the best-fit premium data providers across verticals such as: Insurance, News, Cybersecurity, Adtech, Crypto, Medical and Company—and returns fresh, high-signal outputs packaged for agent consumption (structured responses, consistent schemas, and reliability-oriented delivery).

Kirha also has an application-layer chat product built on top of the infrastructure. This interface functions as an operator console and experimentation layer: users can run conversations by sector, inspect logs, manage MCP gateways and administer API keys. Conversations are intentionally segmented by category because each vertical maps to a distinct provider set and data constraints.

Problem Kirha solves - Current pain points
Teams building agentic products face recurring bottlenecks:

- **Fragmentation:** Premium data is scattered across vendors with incompatible schemas, authentication patterns, rate limits, and contracts.
- **Slow integration cycles:** Adding a single new dataset frequently requires weeks/months of engineering work plus procurement, legal, and compliance effort.
- **Agent usability gaps:** Even when data exists, it's rarely returned in agent-ready formats (tool schemas, entity linking, source attribution/provenance, and guardrails), leading to brittle tools and unreliable outputs.

Concepts
Context as a Service
Kirha runs a cluster of MCP servers that expose partner provider APIs; Kirha then routes and composes responses across one or more providers as needed "auth once, access all partner providers, only pay for the data you use."

Kirha Planning
A key product is the deterministic tool planning: instead of a typical generalist agent loop (discover tool → call → re-plan → call again), Kirha generates a full execution plan upfront

covering what tools/providers will be called, in what order, and how data flows between steps (with parallelisation when possible). This makes execution more predictable and cost effective.

Micropayments
Money just doesn't move as fast as information. Flat-rate subscriptions per dataset or provider are inefficient, subscription models are fundamentally incompatible with the operational and architectural needs of AI agents. Each AI prompt includes a Kirha receipt, detailing the planning process: which data points were considered, selected, and why.

Kirha is packaged as a multi-surface product suite, each aimed at a different user group:

1. **Consumer application layer / console** (operators using/running Kirha in production)
2. **Kirha platform API** (agentic systems)
3. **Modes and distribution** (how Kirha gets adopted)

1. Consumer console (app.kirha.com)

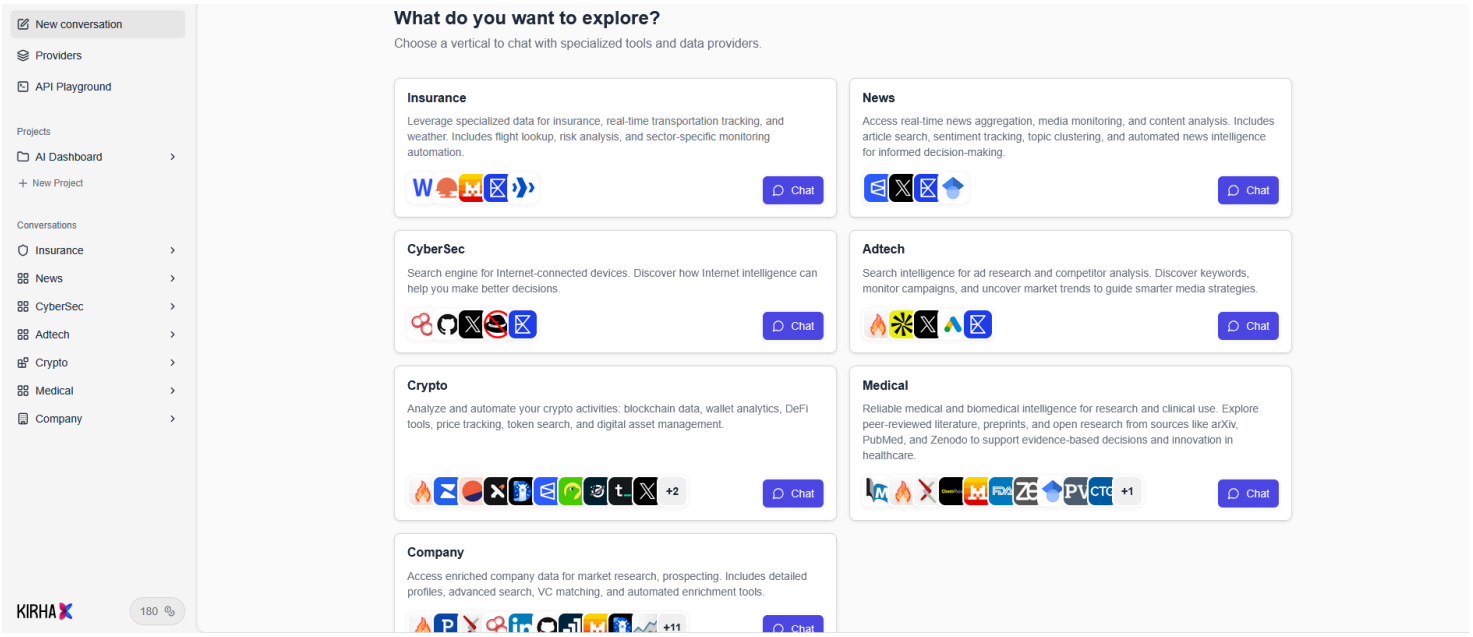

It's where teams:

- Create Projects to isolate usage, logs and billing
- Generate and rotate API keys
- Inspect the Logs Dashboard where every API call is captured with request details, tool execution steps, and response data (critical for debugging + governance)
- Use the chat application layer to have a conversation within each vertical

Each vertical is backed by its own curated set of data providers. When a user selects a vertical, Kirha opens a chat-style interface (similar to mainstream GenAI assistants) scoped to that domain. Users can also switch the underlying model, with support for GPT and Gemini.

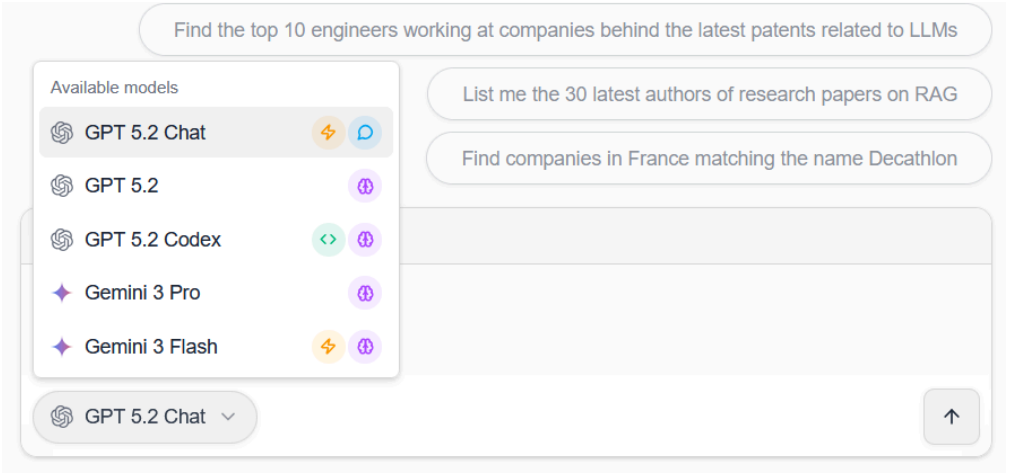

The interface shows how many tools are available within each vertical, and this catalogue is expanding over time. Kirha currently supports 40+ providers via its MCP server network. Many of these sources are paid and have tiered usage limits; a Kirha membership effectively bundles access to otherwise expensive subscriptions, for example Artemis, CoinGecko Pro and DeFiLlama, which together can cost $700+/month if purchased separately.

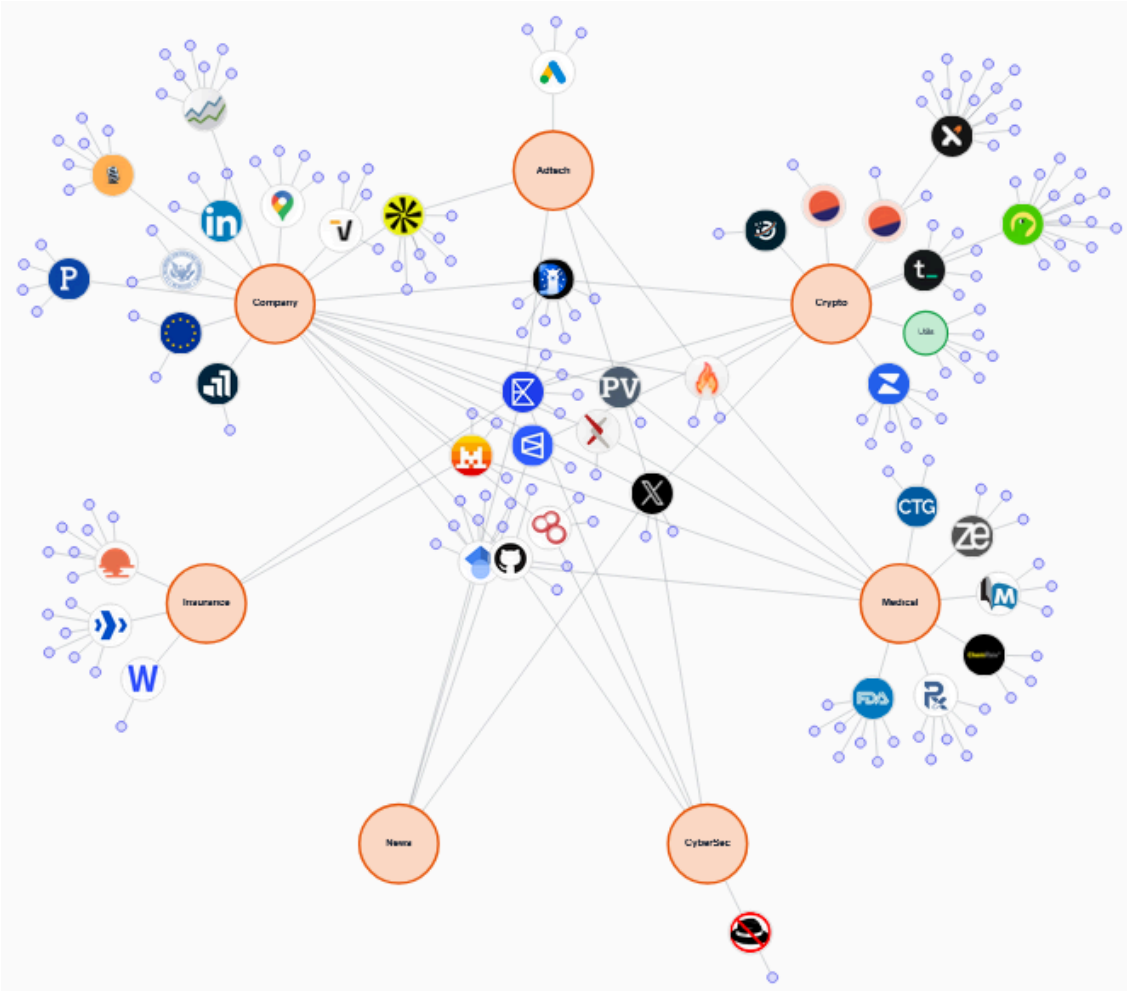

---

2. Kirha platform API
This is the core programmatic surface: a unified way for agents to access premium context across Kirha's domain verticals.

The API has multiple "entry points" into the capability of the application layer.

- **Search API**: default mode for agent workflows, submit a query and Kirha returns results; optionally use plan - review - execute for transparency and spend control.
- **Tools API**: for teams that want explicit control, it lists available tools per vertical and executes a specific tool with explicit parameters (bypassing planning).
- **Typed tool schemas**: tools include input/output schemas, making integration more deterministic and reducing brittle tool calling.

Developers get "one pipe" into many providers, while still being able to choose how much control they want over routing and execution.

3. Modes and distribution
Paths and operating modes that make Kirha easy to drop into wherever agents run:

- Plan-first execution (Kirha Planning): Kirha plans the entire tool execution graph upfront. Covering what tools to call, in what order and how to pass data between steps (parallelising when possible). This enables predictability, review/approval workflows and cost visibility before execution.
- Micropayments + receipts: Kirha positions pay-per-use at the data-point level and generates a "receipt" per prompt describing the planning/selection process for provenance and auditability.
- Local planner mode (privacy / control): Planning can run on the customer's machine (queries don't leave the machine during planning), while execution still uses Kirha's provider network.
- Integrations (distribution):
    - Remote MCP server at mcp.kirha.com + installer CLI for popular MCP clients (Claude Code, OpenCode, Codex, Droid, Gemini CLI).
    - Automation/workflow integrations like n8n and Zapier.

---

3. BUSINESS MODEL

Kirha monetises by turning premium context into a metered utility: users pre-fund usage (via subscriptions or top-ups), and each query triggers micropayments per datapoint/tool call with an auditable "receipt" of what was queried and why. Kirha's core insight is that subscription-per-provider is mismatched to agent workflows (variable demand, bursty traffic, multi-provider composition), so the platform abstracts this into a single credit-based balance that can be spent across the entire provider catalogue.

Pricing Monthly

Monthly plans pre-sell credits. A simple query *"what is the current amount of stablecoin issued and what is the % change in the last 24hrs, use artemis"*. Cost 3 credits.

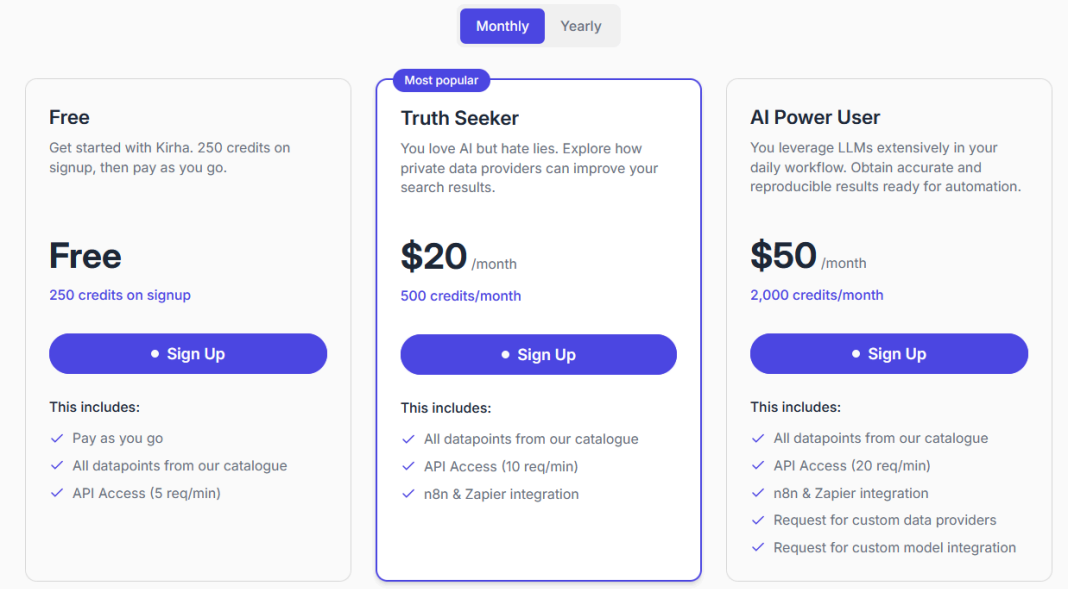

Credit Bundles

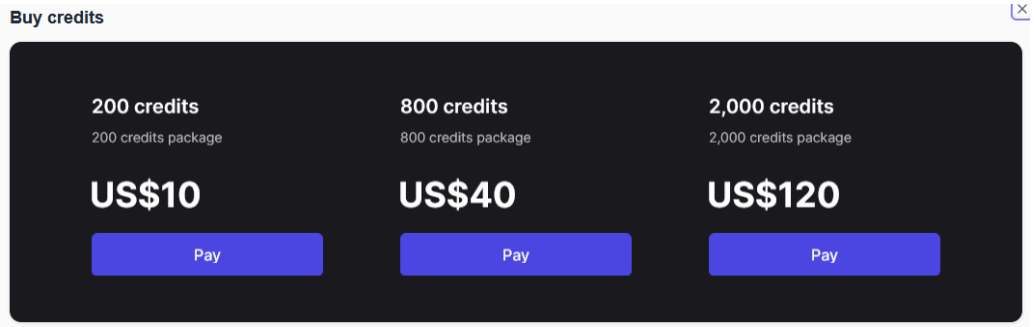

Micropayments

In its first collaboration with Swiss 6022, Kirha will be paid in Swiss 6022's native token on Polygon, creating an immutable record of data-acquisition events; however, management is pragmatic that "money doesn't move as fast as information." Despite integrating x402, they believe onchain settlement remains too slow and too cost-volatile for real-time agent loops resulting in a degraded response time. As a result, the product direction is a hybrid architecture: maintain performance through credits while actively R&D'ing the cheapest, most scalable mechanism to publish receipts onchain, preserving transparency without compromising speed or unit economics.

How the model scales

Kirha's model scales because it is aligned with how agents consume information: high-frequency, multi-step retrieval where the "right" answer often requires composing across several sources. Traditional subscription-per-provider pricing doesn't map cleanly to that pattern; it forces customers into fixed costs and fragmented procurement, resulting in

waste. At scale, Kirha benefits from compounding advantages: higher throughput improves routing efficiency and enables volume-based pricing with suppliers. The more agents utilise, the more it becomes the default gateway for paid data access, thus driving revenue.

4. TRACTION

Business & Ecosystem
DAU: ~2,000 (last 90d), with roughly 50 new users registering every day - company-reported.
X Growth: +10% growth in X followers in February, now sitting close to 6,600 followers.
LinkedIn: Shows over 1,000 followers, employees 2-10.

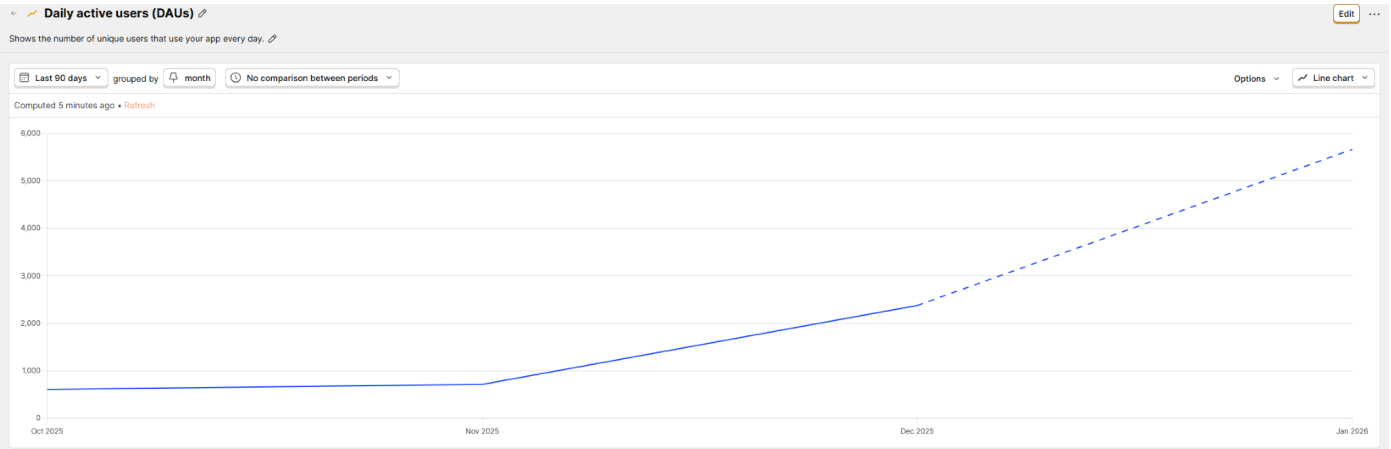

Benchmark - LLM-as-a-Judge
In January they published a Benchmark comparing Kirha to competitor Exa for traditional Web Search. It reached close to one million views and has been a force behind the traffic driven to kirha.com.

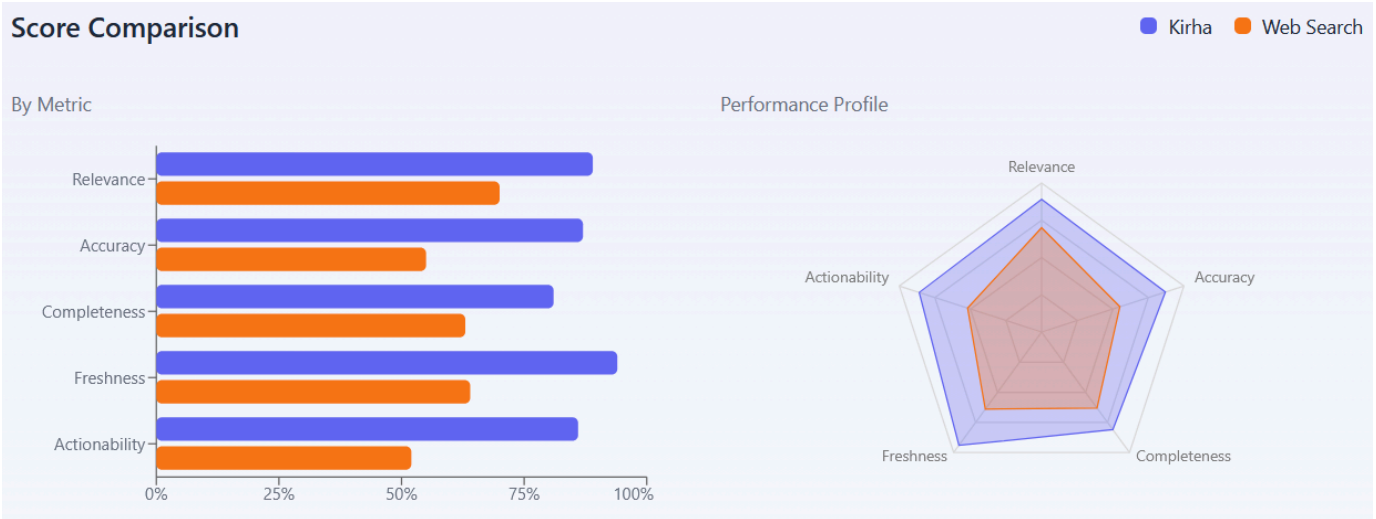

Across all 100 domain-specific queries; Kirha injected 95% fewer tokens compared to traditional web search (233,920 vs 4,604,853). The summarised responses are evaluated using an LLM-as-Judge approach with Gemini 2.5 Flash and extended thinking enabled.

An independent cross-review by Claude (Opus 4.5) validates 96/100 judge decisions (~96% accuracy). The reviewer estimates Kirha wins ~75–80 prompts, while Exa wins ~20–25, and flags 4 edge-case inconsistencies that warrant clarification but "don't invalidate the overall methodology."

Metric	Value
Total prompts reviewed	100
Judge decisions validated	96
Inconsistencies found	4
Judge accuracy rate	~96%
Kirha wins	~75-80
Exa wins	~20-25

The 4 flagged inconsistencies:

- **Relevance vs freshness:** Exa found contributors to the major Tesseract OCR project; Kirha returned small trending repos and still won on freshness.
- **Explicit completeness constraint:** prompt asked for "5 latest authors," Kirha returned 4, but still won.
- **Title accuracy:** Kirha returned "Field CTO" rather than the actual CTO; Exa had the correct CTO but fewer links.
- **Large metric discrepancy:** Kirha reported 274 "active users" vs Exa ~300k (Sep 2025); reviewer argues the gap should trigger an accuracy penalty or definition check.

Product Updates
Nov 2025: Integrated x402 / Coinbase Bazaar to enable paid context access without Kirha accounts, launched Cybersecurity & OSINT vertical (Shodan + AbuseIPDB), migrated tool metadata to Neo4j for semantic discovery and added full log tracing in the dashboard.
Dec 2025: Launched News & prediction markets vertical combining Polymarket MCP + Exa + X MCP; shipped spreadsheet integration via FormulaJS (=KIRHA()), introduced "semantic tool filtering" aimed at keeping latency flat as the tool catalog scales.
Jan 2026: Released open-source tool planning models on Hugging Face (for on-prem/privacy workflows) and launched Medical and AdTech verticals

---

5. COMPETITIVE LANDSCAPE & VALUATIONS

The global agentic AI market size was valued at USD 7.29 billion in 2025 and is projected to grow from USD 9.14 billion in 2026 to USD 139.19 billion by 2034, exhibiting a CAGR of 40.50% during the forecast period (Fortune Business Insights). General-purpose tools like Perplexity, Brave Search and Exa serve the public web. However, there is a need for premium and verified data, a gap Kirha is positioned to fill.

Company	Stage / Valuation	Total Raised	Users/Scale	Growth YoY	Core Model
Tavily	Acquired ($400M)	$25M	1M Monthly	0 → 1M users in 12 months (organic, zero marketing (raised $25M)	API usage based (agent web search)
Exa AI	Series B ($700M)	$111M	Dev focused	1,010% ARR growth YoY (est. $0.9M → $10M, Aug '24–Sep '25)	API usage based (semantic / neural search)
Perplexity AI	Late ($20B)	$1.6B	30M Monthly	66% user growth YoY; ARR growth YoY $20M → $148M ARR	Subscriptions + ads + enterprise
Brave Search	Series A-7 ($1B - '24)	~$252M	82M Monthly	30% MAU growth YoY (77M 2024 → 100M Sept 2025); search queries +80% in 2024	Ad revenue + Search API + BAT token
Kirha	Seed	Unknown	2,000 (90d)	N/A	Pay-per-use credits + Subscriptions + Enterprise

BAT - Basic Attention Token - cryptocurrency designed to fix current advertising model
MAU - Monthly Active Users

The Kirha angle: Every one of these companies grew explosively from a small base and none of them are currently solving what Kirha is solving.

Company	Seed Round	Seed Valuation	Peak Valuation	Multiple
Tavily	$5M July 2024	Est. $15-20M	$400M Acquired (Feb 2026)	20x-27x on seed valuation
Exa AI	Seed + Series A together as $22M, July 2024,	$70M Series A	$700M Series B, September 2025	10x on Series A valuation in 12 months
Perplexity AI	$3.1M seed round, September 2022	Est. $10-15M	$20B (September 2025)	1,300–2,000x on seed valuation
Brave	$4.5M seed round, 2016	Est. $15M pre-money	$980M Series A-7 (2022)	65x on seed valuation

Tavily: Zero marketing spend throughout.
Exa AI: ARR grew roughly 1,010% YoY from an estimated $0.9M to $10M between September 2024 and September 2025.
Perplexity: Hit $148M in annualised revenue by June 2025, up from $63M at end of 2024, and now serves about 30M-45M monthly users processing 780M queries in May 2025.
Brave: surpassed 100M monthly active users as of September 2025, up from 77M at end of 2024, while Brave Search handles 1.6B queries per month and passed $100M in annualised revenue in Q1 2025.

Core Differentiators
Perplexity and Brave have built real moats, but they're consumer moats. Kirha is building a different kind of defensibility: supply-side moat through exclusive access and relationships with premium data providers. This is a marketplace flywheel, not a brand flywheel, each new provider integration strengthens Kirha's leverage in future partnerships.

Most competitors still sell premium context the "old way": upfront commitments, seat licenses, dataset subscriptions, API minimums. Kirha flips the script with a pay-per-use credit model: customers only pay when value is delivered. That dramatically lowers adoption friction and it aligns with how agents actually behave. As agent usage expands across an organisation, Kirha's revenue can scale organically with query volume, without requiring repeated upsell motions.

Finally, Kirha's Bring Your Own Model option isn't a "fancy feature" , it's the minimum requirement for enterprises. Large organisations increasingly insist on controlling where prompts are processed and how sensitive inputs are handled. Kirha's model-agnostic architecture lets enterprises keep their model stack decisions in-house while still benefiting from Kirha's premium data layer, planning and routing. This creates clear partnership opportunities with other decentralised, AI-native infrastructure players, accelerating Kirha's distribution and expanding product usage.

6. TEAM

Founder: Pierre Hay (CEO) - Experience in data infrastructure + security + crypto

- **Enterprise-scale data engineering:** Pierre started at Artefact, contributing to LVMH's cross-Maison data platform, working on data governance/quality and large cloud architectures (Python/Go/GCP).
- **Security + institutional Web3:** He later joined Ledger and integrated an Ethereum framework into Ledger Enterprise (smart contract deployment, secure message signing, institutional dApp interactions).
- **Clear product insight:** Pierre explicitly frames the wedge: MCP adoption is real, but it's missing "planning" and "micropayments," and Kirha exists to make private data providers sell on-demand to agents with stablecoin rails and auditable receipts.

Founders and execution ability:

- **Public facing credibility:** In January Pierre attended a Pitch competition and won the event, a useful signal of founder sales ability and narrative control.
- **Strong technical ability on emerging standards:** When integrating x402, Kirha identified key limitations: the spec assumes static pricing, but Kirha's /search pricing depends on the prompt and cannot be known until the system parses and plans the request. Kirha implemented a dynamic pricing workaround (calculate price on the fly, return it as paymentRequirements), showing the team can both adopt new standards early and engineer around missing pieces.
- **Practical mindset:** Kirha routes x402 payments through an internal credits system and uses the /top-up endpoint to credit a pre-registered x402 admin account while allowing top-ups "with any currency… as long as it's configured," with the practical constraint that the token must be liquid on a DEX. This is the kind of operational thinking required to make "agentic commerce" work beyond demos and "hype" - the fact the team is implementing this themselves indicates their ability.

Image demonstrating flow

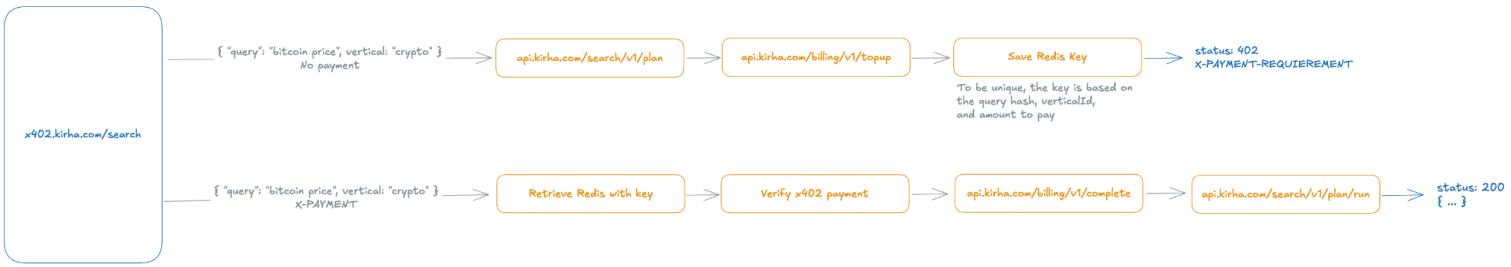

7. KEY RISKS

This section evaluates those risks to pressure-test the investment thesis and clarify what must be true for Kirha to become a durable, category-defining infrastructure layer.

Kirha is well-positioned to shape the gateway for premium data:

- **Risk: Platform disintermediation.** If major agent platforms/model providers ship first-party "data gateways" (connectors + billing + governance) or acquire a competitor, Kirha could be relegated to a commodity connector layer unless its provider relationships are truly defensible.
- **Risk: Supplier concentration / coverage gaps.** If Kirha relies heavily on a small set of "must-have" providers in a vertical, losing one contract (or a pricing change) can materially degrade product quality and unit economics.

Context is the new moat in the agentic era:

- **Risk: Liability and trust failure in high-stakes domains.** In insurance/medical/finance workflows, a single high-profile error (wrong entity, stale data, incorrect risk signal) can cause real-world harm and destroy trust—especially if provenance is unclear or customers treat Kirha outputs as "authoritative."
- **Risk: Observability ≠ compliance.** Logging and receipts are helpful, but if Kirha retains sensitive prompts/results or mishandles access controls, it creates regulatory and security exposure. Moving into medical/insurance expands the blast radius significantly.

Strategic alignment with portfolio:

- **Risk: GTM over-reliance on crypto-native channels.** If early traction is concentrated in crypto/agent marketplaces, Kirha may struggle to translate that into regulated enterprise adoption (news/insurance/medical), where procurement cycles, licensing and compliance needs differ materially.
- **Risk:** Spanning seven verticals creates a real risk of being "broad but shallow." The company may need to win decisively in 1-2 verticals first; otherwise, support burden, provider management, and compliance requirements can overwhelm a small team.

Mitigants

First, Kirha can reduce supplier fragility by focusing on 1-2 verticals where it can generate consistent query volume, then using that demand to convert key providers into partnerships and preferred commercial terms over time.

Second, Kirha should expand distribution by making integrations effectively one-click wherever agents run (MCP clients, automation tools, and agent marketplaces), so adoption is low-friction and usage compounds without a heavy enterprise sales motion.

Third, Kirha can mitigate breadth risk by scaling provider onboarding in a way that does not create linear engineering or support overhead. As the provider catalogue grows, defensibility

increases: more coverage improves routing quality, raises switching costs and makes replication meaningfully harder for competitors.

Finally, Kirha can mitigate breadth and enterprise adoption risk by scaling provider onboarding and deployments. Expanding their deployment optionality (Bring your own model, local planning, plan-review controls) to support compliance-sensitive environments, will widen the partnership surface area.

---

8. SOURCES

Official Website
Official Blog
Docs
Github
Twitter
LinkedIn
Youtube
CAGR Agents