

M C P S E R V E R

N O C O D E

C L O U D H O S T E D

FullEnrich is your B2B data engine.

Find verified emails and mobile numbers across 25+ data vendors.

FullEnrich connects your AI agent to a massive waterfall of 25+ data providers. Instead of guessing, your agent can now find the most probable work email, personal email, and mobile number for any prospect. You only pay credits when the data is actually found and verified, making it a highly efficient way to build and clean contact lists.

A+ Quality Score 96.67/100

b2b-data

contact-enrichment

email-finder

phone-finder

waterfall-enrichment

sales-intelligence



The connectivity layer between AI and the world's software.



Vinkius sits between AI and every application. All communication passes through Vinkius Cloud via the Model Context Protocol (MCP) — with governance, observability, and security at every layer.

Your AI Connections Run Through Vinkius Cloud

The world's largest
managed MCP catalog

Vinkius is the connectivity layer where AI connects to the software your business already runs. We handle the hosting, the security, the credentials, the uptime — you get agents that actually do things.

We operate the world's largest managed MCP catalog. Major SaaS platforms, CRMs, databases, and cloud providers — running, monitored, production-ready. This MCP server is hosted and maintained by the Vinkius Cloud for AI Agents.

The agent doesn't manage credentials, doesn't manage uptime, doesn't manage security. Vinkius does.

— Architecture principle

Four Pillars of the Vinkius Runtime

01 — Security by design

Credentials stay encrypted at rest via AES-256. The AI agent never touches raw keys — they're injected into a sandboxed V8 isolate at runtime. Actions are logged, and connections have an emergency kill switch.

03 — Deterministic observability

Eight immutable metrics per endpoint: request volume, p95 latency, error rate, active connections, cost attribution. A live payload feed logs every tool call with mutation detection.

02 — Built on MCP Fusion

This MCP server was built with **MCP Fusion**, the open-source framework (Apache 2.0) that powers the entire Vinkius catalog. Schema-as-firewall strips undeclared fields, compiled PII redaction runs at zero overhead, and cryptographic lockfiles produce git-diffable audit trails.

04 — Autonomous operations

Servers are deployed, monitored, and patched autonomously. New capabilities and security patches ship weekly. Zero-downtime deployments ensure continuous availability across all managed MCP servers.

AES-256

Encryption at rest

Ed25519

PKI vault signatures

24h TTL

Ephemeral session keys

V8 Isolate

Sandboxed execution

One Token. Instant Access.

Every MCP server on Vinkius is accessed through a **Connection Token**. Tokens are generated in the cloud dashboard and produce a unique MCP endpoint URL. Paste this URL into any MCP-compatible client — no SDK required.

A single token can serve **multiple AI clients simultaneously**, or you can issue separate tokens per client for granular access control. Each token tracks its own request count, last activity timestamp, and can be individually enabled or revoked.

MCP ENDPOINT`https://edge.vinkius.com/{token}/mcp`

Claude



Cursor



VS Code



Windsurf



Grok



Gemini

Security Is the Architecture

Security in Vinkius is not a feature — it's the foundation of the runtime. The gateway enforces multiple independent protection layers between AI agents and third-party APIs.

01 — Ed25519 PKI Vault

Every workspace has an Ed25519 Master Key. Session keys are generated ephemerally (24h TTL) and signed by the Master Key. Credentials never leave the vault boundary.

02 — V8 Isolate Sandboxing

Tool code runs inside isolated-vm V8 isolates with 64 MB memory caps and per-request timeouts. No filesystem access, no network access except through the SSRF-guarded fetch bridge.

03 — SSRF Guard

All outbound HTTP requests are DNS-resolved and validated before execution. Private IP ranges (10.x, 172.16-31.x, 192.168.x, AWS metadata 169.254.x) are blocked at the network layer.

05 — Cryptographic Audit Trail

Every request is signed into a SHA-256 hash chain with Ed25519 signatures. Events form a tamper-proof, SIEM-exportable forensic record.

04 — DLP & PII Redaction

A ResponseGuard pipeline intercepts every tool response. Configurable redaction patterns strip sensitive fields (emails, SSNs, card numbers) before data reaches the AI agent.

06 — Honeypot Trap System

Phantom credentials are injected into isolated environments. If a honeypot is used outside Vinkius infrastructure, the server is quarantined instantly.

Emergency Kill Switch

EU AI Act Art. 14(1)
Compliant

The kill switch is an **emergency halt** mechanism — not a simple toggle. When triggered, it executes three actions atomically:

01 — Server deactivated

The MCP server is immediately taken offline across the entire cluster.

02 — All tokens revoked

Every connection token is invalidated. Total lockout — reconnection blocked until new tokens are issued.

03 — WebSocket connections killed

Active connections terminated via Redis pubsub broadcast. Propagates to every runtime node in the cluster.

Full Visibility. Zero Guesswork.

The Vinkius cloud dashboard includes a full MCP Governance suite — real-time analytics and security controls for production AI operations.

Control Plane

KPI dashboard with request volume, latency, success rate, token consumption, and AI-generated operational briefings.

FinOps

Cost tracking per tool, payload compression savings, budget optimization signals, and consumption trends.

Firewall & DLP

PII redaction activity, sensitive data protection counters, and security event timeline.

Agent Activity

Which AI clients are connecting, how often, and what they're doing — real-time session tracking.

Tool Health

Slowest and most error-prone tools, with actionable root-cause insights and performance baselines.

Incident Log

Error trends, failure rates, status-code breakdowns, and forensic audit trail access.

Get started at cloud.vinkius.com — connect your AI agent in under 60 seconds.

FullEnrich MCP

8 tools available

Cloud-hosted on Vinkius

You can now give your AI agent the ability to hunt down contact info like a professional growth team. FullEnrich works by running a waterfall across dozens of different data vendors to find the best possible match for a person or company. If you provide a LinkedIn URL, your agent will see significantly higher hit rates for both emails and phone numbers. You can use it to build massive prospect lists from scratch using structured filters, or you can use it to clean up your existing CRM by resolving email addresses back to real people. Because the enrichment process is asynchronous, your agent handles the heavy lifting of submitting batches and checking for results while you move on to other tasks. It's a direct way to turn raw company names or LinkedIn profiles into actionable contact data without manual searching.

Core Capabilities

01 — Waterfall Enrichment

Your agent runs a single request across 25+ vendors to find the best contact data.

02 — Reverse Lookup

Your agent turns an email address into a full professional profile.

03 — Professional Search

Your agent filters through 800 million people to find specific targets.

04 — Company Intelligence

Your agent pulls firmographic data for any company in the database.

05 — Credit Management

Your agent monitors your balance to prevent overspending on large batches.

One Click on Vinkius — From Prompt to Execution

Available at vinkius.com/en/ai-agent-connect/fullenrich — connect your AI agent in three steps.

- 01 Subscribe to the FullEnrich MCP on Vinkius.
- 02 Get your API key from the FullEnrich dashboard.
- 03 Connect the MCP to your preferred AI client like Claude or Cursor.
- 04 Tell your agent to search for people or companies.
- 05 Instruct the agent to run the enrichment waterfall on your results.

Connecting FullEnrich to your AI client takes just a few steps to start automating your outbound data.

Built For

This MCP is built for teams that need to turn raw company or person data into verified contact lists.

GTM and RevOps Engineers

They use the agent to automate prospect list building and CRM updates.

Sales Development Reps

They rely on the agent to find deliverable emails and mobile numbers.

Data Teams

They use the agent to programmatically normalize and verify contact records.

What Changes When You Connect

- 01 Pay only when the data is actually found and verified.
- 02 Higher hit rates for emails and phones when using LinkedIn URLs.

03 Access to 25+ different data vendors through a single request.

04 Asynchronous batch processing for large contact lists.

Real-World Applications

Automated Prospecting

Your agent finds specific job titles in certain industries and immediately enriches them with contact info.

Lead Generation

Search the database of 50 million companies to find new targets that fit your ideal customer profile.

CRM Cleaning

You provide a list of old email addresses to resolve the people and companies behind them.

Contact Verification

Run existing lists through the waterfall to ensure you only use deliverable emails and mobile numbers.

8 Tools Available

#	TOOL	DESCRIPTION
01	lookup_company	This tool fetches a single company's firmographic profile. Use a domain or professional URL for the most accurate match.
02	search_companies	This tool searches a database of 50 million companies using structured filters. It is a synchronous operation that does not consume credits.
03	start_enrichment	This tool initiates a bulk waterfall enrichment across 25+ vendors. It returns an enrichment ID so your agent can track the progress of the batch.
04	start_reverse_lookup	This tool resolves people and companies from specific email addresses. It works asynchronously and returns an ID to track the results.
05	get_credits	This tool checks your current workspace credit balance. Use it to ensure your agent doesn't attempt to run a batch you can't afford.
06	lookup_person	This tool retrieves a full professional profile for a single individual. Providing a professional network URL ensures the best results.
07	get_enrichment_results	This tool fetches the status and final data for a bulk enrichment job. You should poll this periodically once the job is submitted.
08	search_people	This tool queries a database of 800 million professionals using specific filters. This is a synchronous tool that uses no credits.

See It in Action

Real prompts you can use once this MCP is connected to your AI agent through Vinkius Cloud.

- U** Find 20 heads of sales operations at software companies with 200-1000 employees in the US, then enrich them with work emails.



I searched the people database for Heads of Sales Operations at US-based software companies with 200 to 1000 employees. I submitted the first 20 contacts for enrichment. The batch is processing, and 18 of the 20 work emails are already verified.

- U** Who is behind john.snow@acme.com and what's their current company?



I ran a reverse lookup for john.snow@acme.com. The profile belongs to John Snow, who is the Head of Sales Operations at Example Inc, a company with 250 employees in San Francisco.

- U** Check my credit balance before we start the next batch.



Your current workspace credit balance has been retrieved. You have enough credits to proceed with the planned enrichment.

Frequently Asked Questions

01 How much does enrichment cost?

You are billed on hits. You only pay credits when verified data is actually found for a contact.

02 How do I get better results for emails and phone numbers?

Using a LinkedIn URL instead of just a name significantly improves hit rates for both work emails and mobile numbers.

03 Is the enrichment process instant?

No, enrichment is asynchronous. Your agent submits the batch, gets an ID, and then polls for the results once they are ready.

04 Can I search for companies without using credits?

Yes, the search_companies tool is synchronous and does not consume any credits.

05 How many people are in the database?

The database contains 800 million professionals and 50 million companies.

Go Live in 60 Seconds

Get your connection token from cloud.vinkius.com, then paste the endpoint URL into any MCP-compatible client.

YOUR MCP ENDPOINT

`https://edge.vinkius.com/[TOKEN]/mcp`

CLIENT	WHERE TO CONFIGURE
 Claude AI	Profile → Customize → Connectors → "+" → Add custom connector → Paste endpoint
 Cursor	Settings → Features → MCP Servers → "+" Add New MCP Server → Type: SSE → Paste endpoint
 VS Code	Ctrl/Cmd+Shift+P → "MCP: Add Server" → add <code>"fullenrich": { "url": "..." }</code>
 Windsurf	MCP Settings → <code>mcp_settings.json</code> → Add endpoint URL
 ChatGPT	Settings → Tools & plugins → Add MCP server → Paste endpoint
 Gemini	Extensions → Add MCP Server → Paste endpoint URL

ASK AN AI ABOUT THIS

Let your preferred AI explain this MCP server

-  Ask ChatGPT 
-  Ask Claude 
-  Ask Perplexity 
-  Ask Gemini 
-  Ask Grok 

READY TO CONNECT

FullEnrich is live on Vinkius Cloud.

Get your connection token, paste it into your AI agent, and start building. No SDK. No deployment. Just results.

Start at cloud.vinkius.com →

vinkius.com · support@vinkius.com

INDEPENDENT PLATFORM DISCLAIMER

Vinkius is an independent platform and is not affiliated with, endorsed by, sponsored by, verified by, or otherwise authorized by FullEnrich. All third-party trademarks, logos, and brand names are the property of their respective owners. Their use in this document is strictly for informational purposes to identify service compatibility and interoperability.

DOCUMENT INFORMATION

Generated	September 2026
MCP Server	FullEnrich MCP
Server ID	01a09f71-e03d-70ba-9f84-e1a3e685a720
Platform	Vinkius Cloud for AI Agents
Endpoint	<code>https://edge.vinkius.com/{token}/mcp</code>

LICENSE & USAGE

This document is generated automatically by the Vinkius PDF Engine. Content reflects the MCP server configuration at the time of generation and may change as updates are deployed. For the most current information, visit vinkius.com/en/ai-agent-connect/fullenrich.