MCP SERVER | NO CODE | CLOUD HOSTED

Mastodon Intelligence MCP

Track the real pulse of decentralized conversations.

Mastodon Intelligence tracks real-time activity and trends across any decentralized social network instance. Monitor what topics are spiking, which external articles gain traction, and the general community pulse using authenticated API access.

A+ Quality Score 100/100

fediverse | mastodon-trends | social-listening | osint | public-api

MAINTAINED BY VINKIUS CLOUD FOR AI AGENTS · JUNE 2026
vinkius.com/mcp/mastodon-intelligence

SCAN TO VIEW

The infrastructure that powers AI agents in the real world.

Vinkius connects AI to the world's software through secure, enterprise-grade infrastructure — enabling real-world execution at scale, built on the Model Context Protocol (MCP).

Your AI Connections Run Through Vinkius Cloud

The world's largest managed MCP catalog

Vinkius is the cloud infrastructure where AI agents connect 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.

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.

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.

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	Ed25519	24h TTL	V8 Isolate
Encryption at rest	PKI vault signatures	Ephemeral session keys	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.

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.

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.

06 — Honeytoken Trap System

Phantom credentials are injected into isolated environments. If a honeytoken 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.

Agent Activity

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

FinOps

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

Tool Health

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

Firewall & DLP

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

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.

Mastodon Intelligence MCP

6 tools available
Cloud-hosted on Vinkius

Tracking conversations across large social platforms used to feel like a guessing game—you only see the highlights reel, filtered by algorithms designed for engagement. This MCP changes that. It lets your AI client pull raw intelligence directly from Mastodon instances, giving you an unfiltered view of community sentiment.

You can discover what topics are genuinely gaining steam using `get_trending_hashtags`, or monitor the live feed to see exactly how people talk about a subject in real time via `get_public_timeline`. Need to know what content is actually influencing the conversation? You can run `get_trending_links` to surface the exact external articles and news stories that the community is sharing right now. This gives you true editorial relevance data, which is crucial for market research or content planning. Since Vinkius hosts this MCP, your agent connects once from any compatible client and gets access to all these specialized tools. It's about getting accurate signals—the ones driven by genuine user activity, not platform whims.

Core Capabilities

01 — Analyze network scale

Get statistics and a description of the specific Mastodon instance you are targeting for analysis.

02 — Monitor topic-specific chatter

Track public posts mentioning a precise hashtag, allowing you to focus on niche industry conversations.

03 — Watch the real-time feed

Retrieve the latest public posts from the open decentralized network as they happen.

04 — Identify viral external content

Determine which outside articles and news sources are currently being shared most frequently across the Mastodon community.

05 — Find high-impact posts

Surface the posts that have received the highest amount of boosts and favorites, indicating organic community interest.

06 — Discover usage trends

Identify currently trending hashtags based purely on how many times they are used by users.

One Click on Vinkius — From Prompt to Execution

Available at **vinkius.com/mcp/mastodon-intelligence** — connect your AI agent in three steps.

01 First, you subscribe to this MCP in Vinkius and provide the specific Mastodon instance URL and your access token.

02 Next, you ask your AI client a question—for example, 'What are people saying about #AI today?' or 'Show me trending news links.'

03 Finally, your agent executes the necessary tool calls (like `get_posts_by_hashtag` or `get_trending_links`) and returns structured data directly to you.

The bottom line is that you get raw, actionable social listening data without having to build and manage complex API integrations yourself.

Built For

Anyone who relies on public opinion or community consensus needs this. If your job depends on knowing what's genuinely spiking in culture, technology, or finance—and you hate platform algorithms—you need this MCP.

Market Researcher

You run `get_posts_by_hashtag` to track how sentiment around a new competitor is evolving over weeks without algorithmic bias.

Content Strategist

You use `get_trending_links` and `get_trending_hashtags` to pinpoint specific, timely topics that your next week of content should cover.

Developer Advocate

You monitor the public feed using `get_public_timeline` to keep a real-time pulse on technical conversations and emerging industry standards.

What Changes When You Connect

01 You stop relying on platform-specific analytics. By using `get_trending_hashtags`, you get usage counts for topics that are purely driven by community interest, bypassing algorithmic manipulation.

02 Don't waste time searching multiple sites. With this MCP, you can monitor the general conversation pulse of the Fediverse instantly via `get_public_timeline` in a single query.

03 You gain immediate insight into external news relevance. The `get_trending_links` tool tells you which outside articles are actually getting traction and shared by the community right now.

04 Find what's truly resonating with people. Instead of guessing, run `get_trending_posts` to identify posts that have earned high engagement through organic sharing and favoriting.

05 You maintain focus on key verticals. When you need to track specific industry conversations—say, #Fintech or #Journalism—you use `get_posts_by_hashtag` to filter the noise instantly.

Real-World Applications

Tracking competitor sentiment

A market researcher needs to know if a major tech shift is causing anxiety. They run `get_posts_by_hashtag` for terms like 'AI ethics' and use `get_trending_hashtags` to see which specific concerns are gaining traction, giving them immediate talking points.

Preparing for a major launch

A content strategist wants to write an article that hits the current cultural nerve. They run `get_trending_links` and discover that 'decentralized finance' is the most shared external topic, directing their focus and informing their outline.

Monitoring event coverage

A developer advocate needs to follow a specific conference. They use `get_public_timeline` for real-time updates while simultaneously running `get_trending_posts` to see which key speakers' posts are getting the most attention and buzz.

Assessing platform health

A communications team needs a baseline understanding of a network. They first run `get_instance_info` to confirm the scale, then use `get_trending_posts` to quickly gauge overall community engagement levels.

Patterns to Avoid

Relying on platform search

✗ AVOID

Trying to find all posts about 'quantum computing' by manually searching the Mastodon website, which often only shows recent activity and misses older viral content.

✓ INSTEAD

Instead, ask your agent to use `get_posts_by_hashtag` for #QuantumComputing. This method ensures you pull a comprehensive, topic-specific feed of public conversations.

Ignoring external context

✗ AVOID

Creating content about a hot topic without knowing which specific articles the community is referencing or arguing over.

✓ INSTEAD

Run `get_trending_links` first. Seeing what links are viral gives you the precise source material and narrative context needed to make your content accurate.

Assuming internal knowledge

✗ AVOID

Thinking that just because a topic is mentioned often doesn't mean it's currently spiking in interest or debate.

✓ INSTEAD

Use `get_trending_hashtags` to check for usage spikes. This tool tells you which topics are trending right now, providing quantifiable evidence of current community focus.

The Right Fit

You should use this MCP if your goal is genuine social listening—meaning you need unfiltered, raw data on what people are *actually* talking about across a decentralized network. If you need to know which links or topics have natural momentum, start here. Don't use it if you only need general platform analytics; for those metrics, dedicated dashboard tools are better. Crucially, don't try to scrape

private accounts (the MCP requires authenticated API access). Also, if your needs are limited to just one type of data—like only monitoring hashtags—you might get by with a more specialized tool category. But because this MCP offers six distinct ways to look at the same network activity, it covers nearly every angle required for comprehensive social intelligence gathering.

Keeping up with decentralized conversations is exhausting.

Today, tracking community pulse means clicking into a few different tabs: checking the main timeline for general buzz, running separate searches for niche topics, and then manually going off-site to see what articles everyone else is reading. It's tedious copy-pasting of data points, and you always feel like you're missing the genuinely viral stuff that happened an hour ago.

With this MCP, your agent handles the whole process. You simply ask for 'trending topics' or 'most shared links,' and it aggregates all that complex monitoring into one clean response. You get accurate, real-time intelligence without ever leaving your chat window.

Discovering Trends with Mastodon Intelligence MCP

The biggest time sink disappears when you don't have to guess where the conversation is headed. You no longer waste time trying to interpret a general timeline; instead, you can use `get_trending_hashtags` to instantly narrow your focus to only the most relevant, actively discussed topics.

It's simple: you ask for the data, and you get pure signals of interest—pure usage counts, purely derived from community action. That's how smart social listening works now.

Mastodon Intelligence with 6 Tools

These tools give your AI client multiple ways to pull raw data—from tracking specific hashtags to identifying the most popular articles shared right now.

#	TOOL	DESCRIPTION
01	`get_instance_info`	Retrieves statistics and general information about the specific Mastodon network you are monitoring.
02	`get_posts_by_hashtag`	Gathers the latest public posts from the Fediverse that contain a specified hashtag, like #DevOps or #Privacy.
03	`get_public_timeline`	Fetches a stream of recent public activity across the open, decentralized network in real time.
04	`get_trending_links`	Identifies and lists external articles or news sources that are currently being shared most often on Mastodon.
05	`get_trending_posts`	Finds the posts with the highest levels of engagement (boosts and favorites) happening right now across the network.
06	`get_trending_hashtags`	Lists currently trending hashtags on Mastodon, showing usage counts based purely on community activity.

See It in Action

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

U What are the top trending hashtags on Mastodon right now?

Here are the top trending hashtags on mastodon.social: #AI (342 uses today), #OpenSource (218 uses), #Privacy (156 uses). These trends are algorithm-free — purely usage-based.

U Show me the latest public posts about #DevOps on the Fediverse.

Found 15 recent posts with #DevOps. The most engaged post is from @sre_engineer discussing Kubernetes cost optimization with 47 boosts and 89 favourites.

U What external articles are being shared the most on Mastodon today?

The top trending links today: 1) 'The EU AI Act explained' from TechCrunch (shared by 89 accounts), 2) 'Open Source funding crisis' from The Register (67 accounts).

Frequently Asked Questions

01 **How do I know if a hashtag is actually popular using Mastodon Intelligence MCP?**

You run the `get_trending_hashtags` tool. This provides usage counts based purely on community activity, so you get a reliable measure of current interest.

02 **Does Mastodon Intelligence MCP track deleted posts?**

No, this MCP pulls public data from the live network feed and trending lists. It doesn't have access to private or deleted content.

03 **What is the difference between `get_public_timeline` and getting general trends?**

`get_public_timeline` shows a continuous, real-time stream of all public activity. Trends show you what specific topics are spiking or being shared most often right now.

04 **Can I use Mastodon Intelligence MCP to track discussions about my company?**

Yes, simply run `get_posts_by_hashtag` using a relevant hashtag like #MyCompanyName. You'll see all the public conversations around that topic.

05 **What kind of data does `get_trending_links` provide?**

This tool gives you the titles and sources of external articles or news stories that are currently receiving the highest number of shares and mentions across the network.

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 `"mastodon-intelligence": { "url": "…" }`
Windsurf	MCP Settings → `mcp_settings.json` → 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

Mastodon Intelligence 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 Mastodon Intelligence. 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	June 2026
MCP Server	Mastodon Intelligence MCP
Server ID	`019ef248-84ef-73bd-9651-681c13ed729e`
Platform	Vinkius Cloud for AI Agents
Endpoint	`https://edge.vinkius.com/{token}/mcp`

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/mcp/mastodon-intelligence.