

MCP SERVER

NO CODE

CLOUD HOSTED

Hugging Face Discussions MCP

Find the conversation happening right now, wherever it is.

Hugging Face Discussions intercepts high-value conversations across the world's largest machine learning developer hub. Connect your agent to monitor trending models, scan active discussions for specific keywords, and engage with AI developers in real time. It turns passive observation into actionable intelligence, letting you find exact threads discussing tool calling limitations or deployment errors.

A+ Quality Score 100/100

hugging-face

machine-learning

community-management

issue-tracking

reconnaissance



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.

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.

Hugging Face Discussions MCP

8 tools available

Cloud-hosted on Vinkius

This MCP lets your agent manage all community interactions on the Hugging Face platform. Instead of manually checking multiple dashboards or sifting through endless model pages, your agent performs deep reconnaissance automatically. You can point it at trending repositories and have it scan active discussions for specific keywords—say, 'tool-calling' or 'deployment errors'. When you find a topic that matters to your business, the agent pulls up the full details of the conversation so you know exactly what's being talked about. Need to jump in? You can even draft replies or open new issues programmatically, making it look like a real community manager is working 24/7. It's all accessible through Vinkius, letting you keep your AI client connected across thousands of services without switching catalogs.

Core Capabilities

01 — Find hot topics

Scans currently trending models and projects to surface high-interest conversations.

03 — Track keywords across projects

Searches multiple active repositories for keywords related to your stack, finding exactly where developers are running into problems.

05 — Report issues or ask questions

Creates brand new discussions or issues when you need to report a bug or suggest a feature formally.

02 — Audit specific discussions

Reads the full history of a single conversation thread so you understand the entire context before replying or acting.

04 — Engage and reply

Posts comments or replies directly to discussions, giving technical answers right in the conversation stream.

One Click on Vinkius — From Prompt to Execution

Available at vinkius.com/mcp/hugging-face-discussions — connect your AI agent in three steps.

- 01** Subscribe to the integration and provide your Hugging Face Access Token, ensuring read/write permissions for discussions.
- 02** Direct your agent to monitor specific repositories, trend lists, or keywords you care about (e.g., 'RAG pipeline setup').
- 03** Your AI client executes the search, identifies relevant conversations, and returns actionable summaries ready for engagement.

The bottom line is that you get a single pane of glass view into massive amounts of distributed developer conversation data.

Built For

This MCP is built for ML infrastructure providers and technical marketing teams. If your job involves monitoring the health, trends, or pain points within the AI developer community, this tool saves you hours of manual searching.

ML Developer Advocate

Uses the MCP to monitor discussions across popular models, finding technical questions about specific frameworks and responding with expert advice.

Technical Product Manager

Runs rapid audits on trending repositories using intelligence reports to gauge community interest in new features or model types before defining a roadmap.

Growth Lead / Community Manager

Scans the most active projects for keywords related to competitor technology or pain points, ensuring your brand is visible when users are stuck.

What Changes When You Connect

-
- 01** Never miss a critical user pain point. Use `scan_target_discussions` to find keywords like 'deployment errors' across dozens of trending models automatically.

 - 02** Understand context immediately. Before you reply, use `get_discussion_details` to read the full thread history, ensuring your technical solution is perfectly targeted.

 - 03** Know what's hot before it explodes. Run `get_trending_repositories` to see which models and datasets are generating the most community buzz right now.

 - 04** Act like a real expert. Use `comment_discussion` or `create_discussion` to provide solutions or report bugs directly where developers are looking for help.

 - 05** Quickly assess market viability. Run `get_repo_intelligence` on a model of interest to get stats on likes, downloads, and current open discussions in minutes.
-

Real-World Applications

Competitor monitoring

A marketing manager needs to know if developers are discussing flaws in a competitor's tool. They ask their agent to `scan_target_discussions` using the keyword 'competitor-name limitation.' The agent finds three active discussions and summarizes the common thread: poor API documentation.

Launching new infrastructure

A company is launching an AI platform. They ask the agent to `get_trending_repositories` and then use that list to find discussions related to 'tool-calling.' The resulting insights show a clear gap in market support for their specific framework.

Feature gap identification

A product team suspects users are struggling with a specific workflow. They instruct their agent to `get_repo_discussions` on their core model's page, filtering for 'latency'. The agent identifies a recurring pattern of complaints about response time.

Community support scaling

Support staff need visibility into all open issues. They instruct the agent to `get_repo_discussions` across five key client repositories, allowing them to triage and prioritize fixes without manual checking.

Patterns to Avoid

Treating discussions like simple searches

X AVOID

A user just types 'tool calling' into the search bar and gets a list of results, but has no idea which thread is active or what the actual conversation flow is.

✓ INSTEAD

To get the full picture, first use `scan_target_discussions` to find the relevant threads. Then, use `get_discussion_details` on the specific issue number to read the entire context.

Over-relying on model stats

X AVOID

A developer sees a repo with 10,000 likes and assumes it's perfect, ignoring any active community complaints.

✓ INSTEAD

Always pair `get_repo_intelligence` with `get_repo_discussions`. The discussions show the *real* problems that the stats hide.

Posting blind replies

X AVOID

An agent blindly posts a reply to an old thread, missing crucial context or realizing the issue was fixed last week.

✓ INSTEAD

Always use `get_discussion_details` first. This confirms you read the entire conversation and that your response is timely and relevant.

The Right Fit

Use this MCP if your goal is deep, actionable intelligence drawn from ongoing human technical conversations. If you need to know *why* developers are struggling with a model or framework, this is essential. Use `scan_target_discussions` when you're hunting for specific keywords like 'data leakage' across many projects simultaneously. Don't use it if your only goal is to find the top 10 most downloaded models; that requires `get_trending_repositories`. Also, don't use this MCP if you just need to read static documentation—that's a simple API call. This tool handles active, messy, human conversations.

The manual process of tracking community sentiment is a nightmare.

Today, finding out what the ML developer community is actually talking about requires jumping between pages: checking trending lists, clicking into repositories, reading the issue board, and then manually scanning comments for keywords like 'latency' or 'tool-calling'. It's copy-pasting links and losing context in a dozen browser tabs.

With this MCP, you let your agent perform that entire reconnaissance sweep instantly. You get a consolidated view of where developers are stuck, allowing you to focus only on the highest-value conversations without ever leaving your dashboard.

Hugging Face Discussions gives you immediate community context.

The manual steps—checking model statistics (via `get_repo_intelligence`) and then separately checking the issue board (using `get_repo_discussions`)—are now combined. You get a single, continuous feed of both quantitative metrics and qualitative pain points.

You don't just know *that* people are talking about something; you know *what* they are saying, allowing you to respond with authority immediately.

Hugging Face Discussions: 8 Tools

These tools let your agent perform every action needed to monitor, audit, or contribute to community discussions across the entire Hugging Face platform.

#	TOOL	DESCRIPTION
01	<code>comment_discussion</code>	Posts a formatted comment directly into an existing Hugging Face discussion thread.
02	<code>create_discussion</code>	Starts a brand new issue or discussion within any specified Hugging Face repository.
03	<code>get_discussion_details</code>	Retrieves the complete history and comments for one specific, identified Hugging Face discussion.
04	<code>get_repo_discussions</code>	Retrieves all current discussions and open issues associated with a single Hugging Face repository.
05	<code>search_repositories</code>	Finds repositories (models, datasets) based on keywords to identify high-value areas for engagement.
06	<code>get_repo_intelligence</code>	Generates a quick report summarizing key statistics and current activity levels for any given repository.
07	<code>get_trending_repositories</code>	Lists the most popular repositories currently gaining traction across Hugging Face.
08	<code>scan_target_discussions</code>	Searches trending models and projects for discussions that contain a specific keyword or phrase.

See It in Action

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

U Find discussions about 'tool-calling' in trending models.



I scanned 10 trending models and found 2 active discussions about 'tool-calling'. In 'meta-llama/Meta-Llama-3-8B' (Issue #142), a user is asking how to implement tool calling. Should I prepare a reply suggesting our Vinkius infrastructure?

U Get an intelligence report on the Qwen/Qwen2-72B model.



Here is the intelligence report for Qwen/Qwen2-72B: It has over 15,000 likes and 2M downloads this month. There are currently 8 active open discussions. The top discussion is about context window limitations.

U Reply to discussion #42 in mistralai/Mistral-7B-v0.1.



I've successfully posted your comment to discussion #42. The community will now see your technical breakdown of the infrastructure solution.

Frequently Asked Questions

01 How can I find out what developers are discussing about 'tool-calling' in trending models using Hugging Face Discussions?

You use the `scan_target_discussions` tool. It searches all currently popular repositories for that exact keyword, giving you a list of active threads where developers are asking questions or reporting issues.

02 Can I get an overview of a model's performance and discussion activity with Hugging Face Discussions?

Yes. The ``get_repo_intelligence`` tool provides a rapid report covering key statistics, download numbers, and the count of open discussions for immediate situational awareness.

03 What is the difference between using ``search_repositories`` and ``get_trending_repositories``?

``get_trending_repositories`` shows you what's currently hot and gaining buzz. ``search_repositories`` lets you specifically look for models, datasets, or spaces based on criteria you define.

04 If I want to reply to a conversation, should I use ``comment_discussion`` or ``create_discussion``?

Use ``comment_discussion`` if the conversation already exists and you are providing input there. Use ``create_discussion`` when the topic is new, and you need to open an entirely fresh issue.

05 Does Hugging Face Discussions help me monitor all my models?

Yes, by using ``get_repo_discussions``, you can pull discussion feeds for multiple specific repositories. This centralizes your community monitoring without logging into each model page.

06 How do I find my Hugging Face Access Token?

Go to your Hugging Face account settings, under **Access Tokens**, and create a new fine-grained token. Ensure it has permissions to read and interact with discussions on the specific repositories you want to target.

07 Can the agent post replies automatically?

Yes. Using the ``comment_discussion`` tool, your agent can write and post Markdown replies directly into community threads.

08 How does the target scanning work?

The ``scan_target_discussions`` tool fetches the top trending repositories (models, datasets, etc.) and then iterates through their active discussions, filtering for the exact keyword you specify. It's highly efficient for finding interception points.

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 `"hugging-face-discussions": { "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

Hugging Face Discussions 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](https://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 Hugging Face Discussions. 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	Hugging Face Discussions MCP
Server ID	019eeddf-d013-720f-ab40-73ed0aeb1614
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/hugging-face-discussions.