

MCP SERVER

NO CODE

CLOUD HOSTED

# Fireflies.ai MCP

Turn recorded meetings into actionable data.

Fireflies.ai connects your AI client directly to meeting intelligence. It transcribes calls from Zoom or Google Meet and turns hours of audio into searchable data points, summarized action items, and detailed conversation analytics. Stop re-watching recordings; ask your agent questions about your entire meeting history.

**A+** Quality Score 100/100

meeting-transcription

note-taking

conversation-analytics

voice-assistant

meeting-intelligence

nlp



# 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](https://cloud.vinkius.com) — connect your AI agent in under 60 seconds.

# Fireflies.ai MCP

12 tools available

Cloud-hosted on Vinkius

Connecting Fireflies.ai through Vinkius gives you deep insight into everything that happens during virtual meetings. Instead of dealing with disorganized notes or massive video files, you can talk to your AI client and simply ask what happened. It transcribes conversations across multiple platforms, pulls out key action items, and even analyzes who talked more than whom—giving you hard metrics like the talk-to-listen ratio for any team discussion. You can query every single transcript in your history with natural language prompts. Need to track down a specific decision made three months ago? Just ask. Plus, you manage the bot directly; you can invite it to an ongoing meeting using just the URL or rename old transcripts so your database stays clean and organized.

---

## Core Capabilities

### 01 — Manage Transcripts

You can list recent meetings, fetch detailed summaries, rename transcripts, or delete them entirely.

### 03 — Analyze Conversation Metrics

The system delivers aggregate data, including talk-to-listen ratios and words-per-minute counts, for deep team communication analysis.

### 05 — Manage Team Data

The agent can list all team members and retrieve user-specific meeting details.

### 02 — Control Live Meetings

You tell your agent to invite the Fireflies bot into a live Zoom or Google Meet session using only the meeting link.

### 04 — Query Meeting Content

You use the AskFred integration to ask specific questions about past meetings or start brand new threads of analysis.

# One Click on Vinkius — From Prompt to Execution

Available at [vinkius.com/mcp/firefliesai](https://vinkius.com/mcp/firefliesai) — connect your AI agent in three steps.

- 01 Subscribe to this MCP on the Vinkius Marketplace.
- 02 Input your Fireflies API Key into your AI client settings.
- 03 Start asking your agent questions about meetings, or tell it to join a live meeting using its tools.

The bottom line is you get immediate access to structured, searchable data from all your recorded conversations.

---

## Built For

This is for anyone drowning in post-meeting cleanup. It helps the Project Manager who hates re-watching syncs just to find a single action item. It's perfect for Sales Reps who need to analyze call sentiment, or EAs who spend hours organizing and renaming files.

### Project Manager

You use this MCP to quickly pull action items and summaries from recent project syncs without having to re-watch the full recordings.

### Sales Representative

You run analytics on discovery calls, checking talk-to-listen ratios and sentiment data to immediately refine your pitch for next time.

### Executive Assistant

You use the system to automatically organize, rename, or delete meeting transcripts using simple AI commands, keeping the archive clean.

---

## What Changes When You Connect

- 01 Stop wasting time re-watching hours of video. Instead, you can ask your agent to pull out specific action items or keywords directly from the transcript using the `get_transcript` tool.

- 
- 02** Improve team communication immediately by getting conversation analytics. The `get_analytics` function delivers hard metrics, like talk-to-listen ratios, so teams know what needs work.

---

  - 03** Never miss a follow-up task again. Use the `ask_fred_thread` capability to query your entire meeting history and retrieve specific decisions made months ago.

---

  - 04** Keep your workspace clean and searchable. You can list transcripts and use `update_meeting_title` or `delete_transcript` when you're done with old data, keeping only what matters.

---

  - 05** You don't need to copy links manually. Use the `add_to_live_meeting` tool to invite the bot directly into a Zoom or Google Meet call using just the URL.

---
- 

---

## Real-World Applications

### Finding Decisions After a Large Sync

A PM needs to know who was assigned the Q3 budget report from a meeting last month. Instead of searching through 40 minutes of video, they ask their agent to `get_transcript` for that meeting and specify 'action items related to budgeting,' instantly getting a clear list.

### Quickly Onboarding a New Team Member

An executive assistant wants to give a new hire context. They ask the agent to `list_users`, get details on key team members using `get_user`, and then summarize all past project discussions for the new person.

### Analyzing Sales Performance

A sales manager wants to benchmark two reps. They use the `get_analytics` tool on both reps' discovery calls, allowing them to compare talk-to-listen ratios and identify where one rep needs pitch refinement.

### Handling Meeting Overload

A user has hundreds of old transcripts cluttering their system. They use `list_transcripts` to see what they have, then rename irrelevant ones using `update_meeting_title` and delete the rest via `delete_transcript`.

---

---

# Patterns to Avoid

---

## Summarizing without context

### X AVOID

Asking your agent for a summary of 'the meeting' when you have 50 transcripts. The response is vague, missing key details and action items.

### ✓ INSTEAD

First, use `list_transcripts` to narrow it down. Then, ask the `get_transcript` tool specifically for the meeting ID that matters. This ensures the agent only summarizes the right conversation.

---

## Manual data extraction

### X AVOID

Copying and pasting action items from a transcript into a separate spreadsheet to track them.

### ✓ INSTEAD

Ask your agent to `get_analytics` or use `ask_fred_thread`. The AI handles the complex pattern matching for you, delivering structured data instantly.

---

## Ignoring live meetings

### X AVOID

Forgetting to record a critical brainstorming session because it's happening in Google Meet.

### ✓ INSTEAD

Before the meeting starts, use `add_to_live_meeting`. It invites the bot using just the URL and ensures you capture every word.

---

## The Right Fit

Use this MCP if your biggest time sink is post-meeting cleanup or having to re-watch recordings for a specific piece of data. You need conversational intelligence applied across hundreds of hours of audio, not just one document.

Don't use it if you simply need a calendar reminder or a basic contact list—other simple tools will handle that. Also, don't rely on this to manage your email inbox; it only deals with meeting data.

If you need deep analysis (talk-to-listen ratios, sentiment), you must use the `get_analytics` tool. If you just want a quick summary of what happened, start by listing transcripts and then getting the specific transcript details.

---

---

## The biggest time drain is always figuring out who said what.

Today, after any major sync, you're left with a mix of Slack threads, disorganized notes in Notion, and recordings that are too long to watch. You have to click through five different tabs, skim pages of text, and copy-paste potential next steps into a master document just to get the core decisions.

With this MCP, you simply tell your agent what you need—'What were the three key action items from the Q2 review?' Your agent goes straight to the source material, pulls out only the necessary data, and delivers it in plain text. You get immediate answers without ever opening a video player.

---

## Fireflies.ai gives you full control over your meeting insights.

You no longer need to manually rename or archive transcripts into specific folders across different cloud services. You can use the `list_transcripts` tool to see everything, and then `update_meeting_title` to organize them right inside the system.

The process is clean: you ask for it, you get it. Your meeting knowledge moves from being a static video file to an active, queryable database.

---

# Fireflies.ai: 12 Tools for Meeting Intelligence

These tools let you manage live meeting invites, retrieve detailed transcript summaries, pull conversation analytics, and organize your entire library of recorded discussions.

#	TOOL	DESCRIPTION
01	<code>add_to_live_meeting</code>	Invites Fireflies.ai to join an ongoing live video conference call.
02	<code>create_ask_fred_thread</code>	Allows you to ask a specific question directly to the integrated AI assistant, AskFred.
03	<code>delete_transcript</code>	Removes an entire meeting transcript from your archive.
04	<code>get_analytics</code>	Pulls comprehensive metrics about a past meeting, like who spoke the most.
05	<code>get_ask_fred_thread</code>	Retrieves the back-and-forth message history from your AskFred analysis threads.
06	<code>get_transcript</code>	Fetches all summarized details and key points for a specific, named meeting transcript.
07	<code>get_user</code>	Retrieves detailed profile information about a specified team member.
08	<code>list_active_meetings</code>	Shows you which meetings are currently being recorded and transcribed in real-time.
09	<code>list_ask_fred_threads</code>	Lists all the analysis threads you've created with AskFred.
10	<code>list_transcripts</code>	Retrieves a list of your most recently recorded meetings and their basic summaries.
11	<code>list_users</code>	Outputs a roster listing all team members connected to Fireflies.ai.
12	<code>update_meeting_title</code>	Changes the name of an old meeting transcript for better organization.

---

## See It in Action

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

**U** List my 5 most recent meeting transcripts.



Fetching transcripts... I found 5 recent meetings, including 'Product Roadmap Sync' and 'Q1 Review'. Would you like the AI summary or action items for any of these?

**U** Invite Fireflies to join my current meeting at <https://zoom.us/j/123456789>



Request sent! I've invited the Fireflies bot to your Zoom meeting. It should join shortly to start transcribing the conversation.

**U** Show me the action items from meeting ID 'trans\_987'.



Retrieving summary... For meeting 'trans\_987', the action items are: 1) Update the landing page design, 2) Send the final proposal by Friday, and 3) Schedule a follow-up call with the engineering team.

---

## Frequently Asked Questions

### 01 How do I use Fireflies.ai MCP to start transcribing a live call?

Use the `add_to_live_meeting` tool. Just provide the meeting URL and your agent handles inviting the bot to join and begin recording immediately.

### 02 Can I find action items from old meetings using Fireflies.ai MCP?

Yes. You list recent transcripts first, then use `get_transcript` or `create_ask_fred_thread` to query the content for specific deliverables and owners.

---

**03 What if I want to analyze team talk-to-listen ratios? Do I need Fireflies.ai MCP?**

The `get_analytics` tool does this. It pulls aggregate metrics from your meeting history, giving you hard data on who spoke and for how long.

---

**04 Does Fireflies.ai MCP let me organize my transcripts?**

Yes. You can list transcripts to see what's available, then use `update_meeting_title` to rename them or `delete_transcript` to clean up the archive.

---

**05 Is there a way to ask specific questions about past meetings with Fireflies.ai MCP?**

Use `create_ask_fred_thread`. This tool lets you engage AskFred directly, allowing for deep, focused analysis on any meeting content.







---

# Go Live in 60 Seconds

Get your connection token from [cloud.vinkius.com](https://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
 <b>Claude AI</b>	Profile → Customize → Connectors → "+" → Add custom connector → Paste endpoint
 <b>Cursor</b>	Settings → Features → MCP Servers → "+ Add New MCP Server" → Type: SSE → Paste endpoint
 <b>VS Code</b>	Ctrl/Cmd+Shift+P → "MCP: Add Server" → add <code>"firefliesai": { "url": "..."</code>
 <b>Windsurf</b>	MCP Settings → <code>mcp_settings.json</code> → Add endpoint URL
 <b>ChatGPT</b>	Settings → Tools & plugins → Add MCP server → Paste endpoint
 <b>Gemini</b>	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

# Fireflies.ai 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](https://vinkius.com) · [support@vinkius.com](mailto: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 Fireflies.ai. 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	Fireflies.ai MCP
Server ID	019d7599-d0b3-7271-8658-c5159f765ba1
Platform	Vinkius Cloud for AI Agents
Endpoint	<a href="https://edge.vinkius.com/{token}/mcp">https://edge.vinkius.com/{token}/mcp</a>

### 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/firefliesai](https://vinkius.com/mcp/firefliesai).