

MCP SERVER

NO CODE

CLOUD HOSTED

# Hotjar MCP

Analyze Site Behavior with Natural Conversation.

Hotjar connects user experience analytics directly into your agent. It lets you analyze heatmaps, session recordings, and survey data through natural conversation. Instead of clicking through dashboards to find drop-off points or confusing feedback, simply ask your AI client about specific users, funnels, or site interactions to get immediate, actionable insights.

**A+** Quality Score 98.33/100

heatmaps

session-recordings

user-feedback

conversion-funnels

nps-surveys

ux-research



# 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.

# Hotjar MCP

15 tools available

Cloud-hosted on Vinkius

Hotjar brings user behavior analytics into the chat window. You can stop manually compiling reports from different tabs just to understand why people leave your site. This MCP lets your agent talk directly to your Hotjar data, giving you deep insights without ever seeing a dashboard.

Need to know what users are clicking on a specific page? Ask for the heatmaps. Want to see exactly where they got stuck? You can review conversion funnels and identify those drop-off points instantly. If you're investigating a bad user journey, your agent can even find session recordings or pull up all feedback submissions—complete with screenshots.

By connecting Hotjar through Vinkius, you treat your site data like any other source of information: something you just ask questions about. You get answers on NPS scores, survey trends, and specific user activity for a given ID. It's simple conversation that delivers complex product insights.

---

## Core Capabilities

### 01 — Identify user drop-off points

Examines conversion funnels to pinpoint exactly where visitors are leaving your site.

### 03 — Review individual user sessions

Retrieves session recordings and detailed behavior history for specific user IDs.

### 05 — Check site performance metadata

Lists all tracked sites within your Hotjar account to manage scope and context.

### 02 — Analyze click and scroll patterns

Generates heatmap data, showing which parts of a page users pay attention to most.

### 04 — Summarize feedback trends

Gathers all incoming survey responses, rating breakdowns, and written comments into digestible reports.

# One Click on Vinkius — From Prompt to Execution

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

- 01 Subscribe to this MCP and enter your unique Hotjar Client Secret via the API settings.
- 02 Your AI client connects, authenticating access to all your user behavior data streams.
- 03 You ask a question like, 'What's wrong with the checkout funnel?' and get an immediate, summarized answer backed by raw data.

The bottom line is that you get expert-level site analytics without having to click through any dashboard yourself.

---

## Built For

This is for the Product Manager who needs proof points, or the UX Researcher tired of stitching together data from five different tabs. You're the person whose job depends on knowing *\*why\** users fail at 2 AM.

### Product Manager

Tracks conversion funnels and monitors NPS scores to identify which features are causing drop-off.

### UX Researcher

Reviews session recordings and heatmaps to watch how users actually interact with the UI, finding behavioral patterns that reports miss.

### Growth Marketing Lead

Analyzes feedback widgets for emerging pain points or optimizing page engagement based on click data.

---

## What Changes When You Connect

- 01 Stop digging through raw data. Instead of manually reviewing `list_funnels` to find the biggest drop-off point, your agent instantly tells you where 60% of users are bailing out.

- 
- 02 Get deep user context immediately. With `lookup_user`, you can feed your AI client a specific ID and get a full summary of that person's journey across all pages.

---

  - 03 Review qualitative feedback without logging in. You can ask the agent to summarize or list responses using `list_feedback_responses` right from the chat, complete with screenshots.

---

  - 04 Understand visual attention instantly. Instead of guessing where users are looking on a page, use `get_heatmap` to see precise click and scroll data for any element.

---

  - 05 Keep track of your setup easily. Use `list_sites` to confirm which URLs Hotjar is actively tracking before running an analysis.
- 

---

## Real-World Applications

### Investigating a sudden drop in signups

A growth team noticed conversion rates dropped 15% yesterday. They ask the agent to check `list_funnels` and immediately see that the biggest step-down is between 'Pricing Page' and 'Plan Selection'. The agent also pulls up recent feedback from `list_feedback_responses`, showing several users complaining about mandatory signups.

### Comparing two different site layouts

A UX Researcher wants to compare old and new product pages. They use `list_heatmaps` to pull data for both versions, allowing the agent to summarize which version generates more clicks in key areas.

### Profiling a high-value, confused user

A PM identifies an important user ID. They use the agent to perform a `lookup_user` for that ID. The resulting report shows the user spent 10 minutes on the features page but never clicked the main CTA. This points directly to a friction point needing redesign.

### Summarizing market sentiment post-launch

A Product Manager asks about recent user sentiment after a major feature release. The agent reviews all `list_surveys` responses and summarizes the top three positive trends, alongside the most common negative critique.

---

# Patterns to Avoid

---

## Asking for 'all data'

### ✗ AVOID

A user asks, 'Give me all user activity.' The agent either fails or returns a massive, unusable dump of raw JSON/data points.

### ✓ INSTEAD

Instead, use specific tools. Start by running ``list_recordings`` to get the count, then ask for a summary focusing on the last 7 days. If you need details, run ``get_recording`` using a specific ID.

---

## Forgetting user context

### ✗ AVOID

A team member asks 'What's wrong with conversions?' without specifying the stage or page.

### ✓ INSTEAD

Always narrow it down. Start by listing funnels (``list_funnels``) to define which journey you mean, then ask for data on a specific step.

---

## Ignoring feedback details

### ✗ AVOID

A user gets vague complaints and just reads the summary score.

### ✓ INSTEAD

Don't stop at scores. Use ``list_feedback_responses`` to pull up actual text, screenshots, or rating breakdowns to understand *why* the score is what it is.

---

## The Right Fit

Use this MCP if your primary goal is understanding user behavior—the 'why' behind the numbers. If you need to know where users click, how long they look at a section, or if they got stuck on a specific form field, this is essential. It's perfect for cross-referencing data from `get_heatmap` with qualitative input from `list_feedback_responses`. Don't use this if your primary need is financial reporting; you won't find revenue projections here. Similarly, don't use it just to list pages; while `list_sites` helps, the core value is interpreting the data (e.g., using `get_survey_stats` to measure sentiment). You should only reach for this when your question starts with 'Why did they do that?' or 'What are they looking at?'

---

## Manually tracking user behavior across multiple dashboards is a nightmare.

Today, finding out why users drop off means opening the funnel report, checking session recordings for specific times, and then cross-referencing those findings with survey responses. You copy data from one tab into another spreadsheet just to build a picture that should be simple.

With this MCP, you skip all the clicking and copying. Just ask your agent: 'Why are users abandoning step three?' It pulls together funnel drop-off rates, related feedback submissions, and even points to specific heatmaps on the checkout flow—all in one conversation.

---

## Hotjar MCP delivers actionable insights from every single click.

You no longer have to manually compare what users *say* (via `list_surveys` responses) against what they *do* (visible in a heatmap). The agent combines those two data sets, telling you if the complaint about slow loading matches up with poor scroll depth metrics.

This means your product decisions shift from 'We think this is broken' to 'The data shows exactly why and where it's broken.' It's that directness that changes everything.

---

# Hotjar with 15 Tools

These tools allow your agent to execute specific data retrieval tasks, giving you granular access to every aspect of user behavior in Hotjar.

#	TOOL	DESCRIPTION
01	<code>check_hotjar_status</code>	Verifies that the connection to your Hotjar account is working properly.
02	<code>get_feedback_widget</code>	Retrieves specific details about a feedback widget you've set up on your site.
03	<code>get_heatmap</code>	Pulls the click and scroll data for a particular page to visualize user attention.
04	<code>get_recording</code>	Gets specific metadata about an individual recorded session.
05	<code>get_survey_stats</code>	Retrieves aggregated statistics for a survey, including NPS scores and completion rates.
06	<code>get_survey</code>	Fetches the full details and settings of any specific survey you've created.
07	<code>list_feedback_responses</code>	Lists all individual user feedback submissions, including screenshots and ratings.
08	<code>list_feedback</code>	Shows a list of all the feedback widgets configured on your tracked sites.
09	<code>list_funnels</code>	Lists every conversion funnel you've set up, showing overall drop-off data for each step.
10	<code>list_heatmaps</code>	Provides a list of all the heatmap snapshots available across your tracked pages.
11	<code>list_recordings</code>	Lists all recorded user sessions, including their duration and page count.
12	<code>list_sites</code>	Retrieves a list of every single website that Hotjar is tracking for you.
13	<code>list_survey_responses</code>	Lists all individual responses submitted to your surveys, allowing review of answers and dates.
14	<code>list_surveys</code>	Retrieves a list of every survey you've created in the Hotjar account.

#	TOOL	DESCRIPTION
15	lookup_user	Looks up all recorded behavior, session history, and data for a specific user ID.

---

## See It in Action

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

### U Show the NPS survey results and the top feedback submissions this month.



NPS Survey 'Product Satisfaction Q2' (survey\_401): NPS Score: 42 (Good). Promoters: 58%, Passives: 26%, Detractors: 16%. Completion rate: 34.2% (342 responses / 1,000 views). Response trend: +8% vs last month. Top feedback submissions: 1) ★★★★★ 'Love the new dashboard redesign!' — /dashboard page. 2) ★★ 'Checkout is too slow on mobile' — /checkout (screenshot attached). 3) ★★★ 'Would be great to have dark mode' — /settings. 4) ★ 'Can't find the export button' — /reports. Would you like to see the heatmap for the checkout page or the full survey responses?

### U Show the heatmap data for our pricing page and the conversion funnel from landing to signup.



Pricing page heatmap (heatmap\_089): Click data: 'Pro Plan CTA' button — 342 clicks (hottest). 'Compare Plans' toggle — 189 clicks. 'Enterprise Contact' link — 67 clicks. FAQ accordion — 45 clicks. Scroll data: 78% of users scroll past the hero section. Only 34% reach the FAQ section at bottom. Conversion funnel 'Landing → Signup': Step 1: Landing page — 8,400 visitors (100%). Step 2: Pricing page — 3,200 (38.1%, -61.9% drop). Step 3: Plan selection — 1,100 (34.4%, -65.6% drop). Step 4: Signup form — 680 (61.8%, -38.2% drop). Step 5: Completed signup — 420 (61.8% of form, 5% overall). Biggest drop-off: Landing → Pricing (61.9%). Would you like to see session recordings from the pricing page?

### U Look up user behavior for user ID 'usr\_12345' and show their session recordings.



User 'usr\_12345': 7 sessions over 3 weeks. Device: Chrome/Desktop (5 sessions), Safari/Mobile (2). Total time on site: 48 minutes. Pages viewed: 23. Key behaviors: Visited /pricing 4 times but never clicked a CTA. Spent 8 minutes on /features reading comparison table. Submitted feedback: 'Pricing is confusing' (2 stars). Session recordings: 1) rec\_7701 (Apr 20, 12 min, 6 pages — desktop). 2) rec\_7689 (Apr 15, 8 min, 4 pages). 3) rec\_7645 (Apr 10, 5 min, 3 pages — mobile). This user shows high interest but pricing friction. Would you like to watch a specific recording?

---

# Frequently Asked Questions

---

## 01 How do I find out which pages are tracked with Hotjar MCP?

You can list all the sites by calling ``list_sites``. This tells you exactly which URLs your account is monitoring right now, so you know where to focus your analysis.

---

## 02 What does the Hotjar MCP do with user IDs?

It performs a comprehensive lookup of a specific user's activity. Using ``lookup_user`` gives you their entire session history, pages visited, and total time on site in one place.

---

## 03 Can Hotjar MCP help me analyze NPS scores?

Yes, the agent accesses your survey data using ``get_survey_stats``. It provides the full NPS score breakdown (Promoters/Detractors) and tracks trends over time for you.

---

## 04 Is Hotjar MCP better than just looking at analytics reports?

It's different. While standard reports show numbers, this MCP interprets those numbers by combining heatmaps (``get_heatmap``), funnels, and feedback into a single narrative answer for you.

---

## 05 What if I want to see the full text of survey responses?

You can list all individual submissions using ``list_survey_responses``. This allows you to review not just the aggregated stats, but the raw text and context from every user.







---

# 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>"hotjar": { "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

# Hotjar 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 Hotjar. 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	Hotjar MCP
Server ID	019dd105-2c4c-73ea-bf8f-0a3ee9c9bae6
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/hotjar](https://vinkius.com/mcp/hotjar).