

MCP SERVER

NO CODE

CLOUD HOSTED

Fairing MCP

Analyze customer feedback from conversations.

Fairing helps you analyze post-purchase surveys and zero-party data directly through your AI agent. You can list survey questions, track individual customer responses, and pull high-level performance metrics without leaving your workflow. It puts all your consumer insights into natural conversation.

A+ Quality Score 100/100

zero-party-data

customer-attribution

post-purchase-survey

consumer-insights

data-collection

marketing-analytics



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.

Fairing MCP

12 tools available

Cloud-hosted on Vinkius

Managing customer feedback shouldn't require jumping between dashboards or running complex reports. This MCP connects to Fairing so you can analyze everything from post-purchase surveys to zero-party data using just a chat prompt. You gain full control over understanding what customers think after they buy something. Instead of exporting raw data, your agent pulls insights on demand; for example, you can run `list_responses` to see all submitted feedback or use `get_customer_responses` to understand exactly why one specific user left a particular rating. It's about getting answers immediately. This capability fits right into the Vinkius catalog, letting any MCP-compatible client access this data alongside your other services. You can even check active integrations with platforms like Klaviyo or GA4 using `list_integrations`, giving you one place to monitor performance and customer sentiment.

Core Capabilities

01 — List all available surveys

Retrieve a list of every survey you have set up in Fairing.

03 — Find all customer feedback records

List every single submitted response across all your active surveys.

05 — Extract overall performance metrics

Get aggregated insights and high-level performance numbers across your entire suite of surveys.

02 — Get detailed question information

Fetch the full configuration and details for any specific survey question.

04 — Analyze one customer's history

Pull all survey responses and data points associated with a single, named customer.

06 — Review account details

Retrieve essential information about your Fairing account, including API token identity.

One Click on Vinkius — From Prompt to Execution

Available at vinkius.com/mcp/fairing — connect your AI agent in three steps.

- 01 Subscribe to this MCP and paste in your Fairing API Key (you find this key in the Fairing Settings > Account section).
- 02 Connect your preferred AI client, like Cursor or Claude, directly to this Vinkius catalog.
- 03 Ask a natural language question—for example, 'What were the top 3 reasons customers responded last week?' The agent runs the necessary tools and gives you the answer.

The bottom line is: you talk to your data instead of digging through spreadsheets.

Built For

Anyone who needs to connect customer feedback directly into their operational workflow. This helps marketers and analysts move past static dashboards and start asking the right questions about consumer intent.

E-commerce Marketing Manager

Uses it to check customer attribution and survey performance, ensuring every campaign touchpoint is linked back to a sale or feedback point.

Retention Specialist

Reviews individual customer responses to pinpoint the exact pain points that prevent churn, allowing them to personalize outreach messages.

Data Analyst

Pulls raw survey responses and high-level insights directly into their data workflow for deeper statistical modeling or reporting.

What Changes When You Connect

- 01 Stop compiling manual attribution reports. You can ask for aggregated insights, and the agent pulls performance metrics across all your survey streams instantly.

-
- 02 Deep dive into single users' intent using `get_customer_responses` . Instead of wading through dozens of entries, you get everything related to one person's journey in seconds.

 - 03 Keep track of your tech stack. Use `list_integrations` to monitor if Klaviyo or GA4 are syncing correctly without having to log into those platforms separately.

 - 04 Understand the data structure first by running `list_questions` . You can quickly verify exactly what questions are active before asking for complex analysis.

 - 05 Gain full visibility over your account status using `get_account_info` , ensuring your AI agent has the correct credentials and permissions to run reports.
-

Real-World Applications

Why did my recent ad campaign underperform?

The marketing manager needs to know if poor performance relates to survey feedback. They ask their agent, 'Show me the responses from customers who mentioned paid ads.' The agent uses `get_customer_responses` and filters by keywords, giving a clear reason for low conversion rates.

We're launching a new product line; what should we ask about?

The data analyst uses the MCP to run `get_survey_details` on existing surveys. They compare current questions against known market gaps and suggest additions for the next round of feedback.

We need to segment our top 10% most valuable users.

The retention specialist asks the agent to 'List all customers who left feedback mentioning premium products.' The agent runs `list_customers` and helps pull targeted groups, allowing immediate, personalized outreach.

I need a quick summary of general product sentiment.

Instead of looking at 50 different survey tabs, the analyst asks the agent to 'What are the overall performance metrics?' The agent runs `get_insights` and delivers a single, actionable summary.

Patterns to Avoid

Treating data as flat text

✗ AVOID

A user copies 50 survey responses into an LLM prompt and asks it to 'Analyze this.' The AI gets overwhelmed by the sheer volume, providing generic advice without source context.

✓ INSTEAD

Don't dump raw data. Ask your agent specifically: 'Using `list_responses`, pull all feedback from last week mentioning shipping times.' This focuses the tool on actionable subsets of data.

Forgetting which customer gave what

✗ AVOID

A user runs a broad query like 'What are people complaining about?' The AI returns 20 general points, but no way to tie them back to specific accounts.

✓ INSTEAD

Always use `get_customer_responses` when investigating. By filtering by customer ID or name, you ensure the advice is tied directly to a source and an individual journey.

Ignoring system context

✗ AVOID

A user asks for 'the latest survey results' without knowing if data has been synced recently. The agent might give outdated numbers.

✓ INSTEAD

First, run `list_integrations` to confirm the connection status and ensure your AI client knows which systems are live before asking for analysis.

The Right Fit

Use this MCP if your primary need is synthesizing actionable insights from unstructured customer feedback. If you constantly find yourself jumping between a survey tool, a CRM dashboard, and an analytics platform just to answer 'Why did they leave?', this is the solution. The ability to query specific data sets like `get_customer_responses` or pull high-level metrics via `get_insights` means you get answers in conversation, not through complex report building.

Don't use this if your only goal is basic record keeping, like simply listing user accounts (a dedicated identity service might be better). Also, if you only need to manage survey *content* and never look at the *responses*, a standalone content management tool will suffice. This MCP excels because it connects the question setup (`list_questions`) directly to the answer data flow.

Sifting through customer feedback is an all-day job.

Right now, getting a holistic view of post-purchase sentiment means logging into Fairing, running reports on attribution, exporting raw lists of responses, and then pasting those messy spreadsheets into your marketing platform just to start looking for patterns. It's clicking through six different tabs before you even get an answer.

With this MCP, that manual process evaporates. You simply tell your agent what you want—like 'Give me the top three themes from last month's submissions.' The tool runs `get_insights` and hands you a clean summary immediately.

Fairing MCP: Immediate access to customer insights.

You eliminate the need for exporting CSV files, cleaning up data in Excel, or spending hours cross-referencing feedback against marketing campaign dates. The agent handles that heavy lifting for you.

Your AI workflow becomes a direct line to consumer intent. You stop reporting on what happened and start asking your agent why it happened.

Fairing: Manage 12 Tools for Consumer Insights

These tools let you manage everything from listing customer responses to getting high-level performance insights, all through your AI agent.

#	TOOL	DESCRIPTION
01	<code>list_responses</code>	Lists every single survey response submitted by customers.
02	<code>list_surveys</code>	Retrieves a list of all the surveys you have created in Fairing.
03	<code>get_account_info</code>	Pulls general details about your connected Fairing account.
04	<code>get_customer_responses</code>	Fetches all survey responses and data points specifically for one customer.
05	<code>get_insights</code>	Gathers overall performance metrics and high-level trends across your surveys.
06	<code>get_me</code>	Checks the current identity and status of the API token being used.
07	<code>get_question</code>	Retrieves detailed information for a single, specific survey question.
08	<code>get_response</code>	Gets the full details for one particular customer survey submission.
09	<code>get_survey_details</code>	Fetches complete information about a single, defined survey.
10	<code>list_customers</code>	Lists customers who have submitted any kind of survey response.
11	<code>list_integrations</code>	Checks and displays all third-party platforms currently connected to Fairing.
12	<code>list_questions</code>	Lists every available question template used across your surveys.

See It in Action

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

U List all active survey questions on Fairing.



I've retrieved your questions. You have 3 active questions including 'How did you hear about us?' (ID: 101) and 'Who is this gift for?' (ID: 102). Would you like to see the response count for any of these?

U Show me the latest 5 survey responses.



Fetching recent responses... The latest feedback includes 'Instagram Ad', 'Friend Referral', and 'Google Search'. Would you like the details for response ID 205?

U Check my active integrations on Fairing.



Inspecting integrations... Your account is currently connected to Klaviyo (Syncing), Google Analytics 4 (Active), and Meta (Active).

Frequently Asked Questions

01 How do I find out which surveys are active using the Fairing MCP?

You call `list_surveys`. This tool gives you a definitive list of all survey titles and IDs, helping you focus your analysis on what's currently running.

02 Can I get feedback for only one customer using Fairing MCP?

Yes. Use `get_customer_responses` to pull every piece of data submitted by a single user, allowing you to trace their entire interaction history immediately.

03 Does the Fairing MCP show me which platforms are connected?

You run ``list_integrations``. This tool gives you a current status report on all third-party services like Klaviyo or GA4 that are actively syncing with your account.

04 What is the difference between `list_responses` and `get_insights` in Fairing MCP?

The difference is scope. ``list_responses`` gives you a raw, itemized record of every single submission. ``get_insights``, however, provides high-level averages, percentages, and trends across all those responses.

05 Can I check my account details with the Fairing MCP?







You use ``get_account_info``. This confirms your basic setup information, which is useful for troubleshooting or confirming API access scope before running complex queries.

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>"fairing": { "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

Fairing 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 Fairing. 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	Fairing MCP
Server ID	019d7596-2079-7220-9ea6-2b620c64550b
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/fairing.