MCP SERVER | NO CODE | CLOUD HOSTED

Sift (Chargeback) MCP

Automate fraud scoring and dispute resolution.

Sift (Chargeback) manages your entire fraud defense lifecycle from a single conversation. Connect this MCP to instantly check user risk scores, report suspicious chargebacks, and apply manual decisions—all without opening the Sift dashboard. It lets you run real-time dispute resolution and audit user history directly through any AI agent.

F Quality Score 3.6/100

fraud-prevention | chargeback-management | risk-analysis | dispute-resolution | fraud-detection | security-monitoring

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

SCAN TO VIEW

The connectivity layer between AI and the world's software.

Vinkius sits between AI and every application. All communication passes through Vinkius Cloud via the Model Context Protocol (MCP) — with governance, observability, and security at every layer.

Your AI Connections Run Through Vinkius Cloud

The world's largest managed MCP catalog

Vinkius is the connectivity layer where AI connects 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.

Sift (Chargeback) MCP

8 tools available
Cloud-hosted on Vinkius

This connector gives your AI client full control over fraud prevention and chargeback management. Instead of jumping between dashboards to check risk or submit a dispute, you talk to your agent and it handles the heavy lifting. You can ask for a real-time fraud score on any user, instantly seeing if they're high risk or low risk. Need to audit why an order was blocked? Ask the system to list all past decisions applied to that account. If something suspicious happens, you just report the chargeback event, and Sift updates the user profile immediately. Because Vinkius hosts this MCP, you connect your agent once, giving it access to powerful security tools like these for fraud intelligence, decision automation, and deep behavioral tracking.

Core Capabilities

01 — Check current risk score

The system fetches the latest fraud score assigned to any user.

02 — Identify user labels

It retrieves the specific flags or labels Sift has applied to a user's account history (e.g., $bad).

03 — Apply manual actions

You can instruct your agent to manually accept, block, or change the status of a user's account.

04 — Report disputes

The MCP sends formal chargeback reports to Sift, updating the risk profile for the involved party.

05 — Track behavior and events

It logs custom activity like a user logging in or completing a transaction to refine Sift's machine learning model.

One Click on Vinkius — From Prompt to Execution

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

01 Subscribe to this MCP and provide your specific Sift REST API Key and Account ID.

02 Authorize the connection within your preferred AI client (Claude, Cursor, etc.).

03 Ask your agent a question like, 'What is the fraud score for user X?' and get an immediate answer.

The bottom line is you use natural language to perform complex security operations that usually require multiple logins and dashboard navigations.

Built For

This connector is built for anyone whose job involves vetting transactions, managing payment fraud, or handling disputes. If your team spends time cross-referencing logs or manually checking user histories across different systems, this MCP saves you hours.

Risk Analyst

They use the tool to monitor current user scores and audit historical chargeback patterns using simple chat prompts.

Trust & Safety Manager

They apply manual decisions (like blocking an account) or review decision history without ever leaving their main communication interface.

E-commerce Operations Lead

They verify transaction risk and report disputes for high-value orders directly from the chat window during peak hours.

What Changes When You Connect

01 Instant Risk Checks: Instead of logging into Sift just to see a score, you ask your agent for the `get_user_fraud_score` and get an immediate risk assessment in natural conversation. This cuts down on triage time drastically.

02 Full Audit Trail Access: You can request the full history of actions by running `list_user_decision_history`. Your team gets a clear, simple log of every decision made about a user's account.

03 Proactive Fraud Logging: Use `track_sift_event` to automatically feed data—like new logins or transactions—back into Sift. This continuously sharpens their detection models without any manual effort from your side.

04 Immediate Dispute Reporting: When fraud hits, you simply run `report_sift_chargeback`. The agent handles the required data submission and ensures the user's risk profile is updated instantly for review.

05 Decision Enforcement: If a user needs to be blocked or their status changed, running `apply_user_decision` executes that action securely via chat, eliminating the need to click through complex UI forms.

Real-World Applications

Handling an unknown fraud risk during checkout

A customer service agent receives a ticket for a high-value order. Instead of asking the customer for their account details and then opening Sift, they ask their agent: 'What is the fraud score for this user?' The agent runs `get_user_fraud_score` and sees it's 92% (Very High Risk). They immediately run `apply_user_decision` to block the order before processing.

Investigating a sudden spike in chargebacks

A risk analyst notices an unusual cluster of disputes. They ask their agent: 'List all recent chargeback events for this merchant.' The agent runs `report_sift_chargeback` and then uses `list_user_decision_history` to see if any past decisions correlate with the spike, finding a pattern they missed.

Training the model on new behavior

The development team needs to verify that Sift is tracking all relevant activity. They ask their agent: 'Log this specific login and transaction pair.' The agent executes `track_sift_event`, ensuring the data feeds directly into Sift for better future detection.

Auditing compliance after an incident

A manager needs to prove that all necessary steps were taken following a breach. They ask their agent: 'Show me every decision made on user X.' The agent compiles the report using `list_user_decision_history`, providing immediate, auditable proof.

Patterns to Avoid

Manual data aggregation

✗ AVOID

A user reads a fraud score in one dashboard, copies it into a spreadsheet, and then manually types the details into a ticket system to report the chargeback.

✓ INSTEAD

Don't copy anything. Use your agent to run `get_user_fraud_score` for the risk assessment, and then use `report_sift_chargeback` to submit the event directly from your chat interface.

Guessing required actions

✗ AVOID

A user tries to block a user but doesn't know the exact action name or parameters needed for Sift.

✓ INSTEAD

First, run `list_sift_decisions` to see all possible actions. Then, use your agent to execute that specific command via `apply_user_decision`.

Ignoring workflow visibility

✗ AVOID

A team suspects a process is broken but doesn't know which automated rules are running in Sift.

✓ INSTEAD

Use `list_sift_workflows`. This tool shows you exactly what fraud prevention pipelines are configured, giving you immediate oversight of the system.

The Right Fit

Use this MCP if your primary pain point is operating on critical security data (like user scores and dispute logs) but you hate context switching. You need to automate decision-making and reporting in a conversational way. This is ideal for Trust & Safety teams or analysts who operate across multiple systems. Don't use it if you are building an entirely new backend service; that requires direct API integration outside of your agent client. Also, don't rely on this MCP to replace dedicated BI tools—it reports data, but you still need those external tools for deep trend analysis and visualization. Use the `get_user_fraud_score` tool when you only need a single metric, but use the full conversation flow to handle complex tasks like reporting a chargeback or reviewing history.

Handling fraud disputes shouldn't feel like logging into three different dashboards at 2 AM.

Right now, when a customer claims fraud or an order is flagged, the process is manual misery. You open Sift to check the user score; you switch to your ticketing system to read the dispute details; then you open a spreadsheet to find the history of decisions—and finally, you copy all that information into an email for review. It's tedious, slow, and prone to human error.

With this MCP connection, you simply ask your agent to check the risk score or list the decision history for the account in one chat thread. The agent runs the required tools on the backend and hands you a single, unified answer right where you are working. It turns a 10-minute multi-system chore into a two-second conversation.

Sift (Chargeback) MCP: Automated Fraud Decisioning

The key manual steps that disappear are the repetitive data pulls. You never have to copy scores from one screen and paste them into another, nor do you need to manually run through different tabs just to find a decision's timeline.

What's different now is speed, visibility, and accuracy. You get real-time security intelligence directly embedded in your workflow, letting you act on fraud events the moment they happen—not when your morning report runs.

Sift (Chargeback) MCP: 8 Available Tools

Use these tools to manage every aspect of fraud prevention, from checking user scores to reporting complex chargeback events.

#	TOOL	DESCRIPTION
01	`apply_user_decision`	Manually applies a decision, such as blocking or approving, to a specific user account.
02	`list_user_decision_history`	Retrieves a chronological list of all decisions and actions ever applied to a given user.
03	`get_user_fraud_labels`	Fetches the specific fraud labels (like $bad or $good) Sift has assigned to a user's profile.
04	`get_user_fraud_score`	Checks and returns the real-time numerical fraud risk score for any specified user.
05	`list_sift_decisions`	Shows all the available decision types or actions that Sift recognizes.
06	`list_sift_workflows`	Lists and provides visibility into your currently configured fraud prevention workflows within Sift.
07	`report_sift_chargeback`	Submits a formal chargeback event report to Sift, triggering an update of the user's risk profile.
08	`track_sift_event`	Logs general events (like logins or transactions) into Sift so they can train their machine learning model.

See It in Action

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

U What is the fraud score for user 'user_abc_123'?

Checking risk for 'user_abc_123'... The user has a fraud score of 85.2% (High Risk). Sift has labeled this user as '$bad' based on recent transaction patterns.

U Report a chargeback for order #999 from user 'user_789' as '$fraud'.

Reporting chargeback for order #999... The event has been successfully sent to Sift with the reason '$fraud'. This will update the user's risk profile immediately.

U Show me the last 5 decisions applied to user 'user_456'.

Retrieving decision history for 'user_456'... I found 2 recent decisions: 'watch_user' (Applied on 2024-03-15) and 'approve_order' (Applied on 2024-02-10).

Frequently Asked Questions

01 **How do I check a user's risk score using the Sift (Chargeback) MCP?**

You ask your agent to run `get_user_fraud_score` and provide the username. The system returns the current numerical fraud score, along with any associated labels like $bad or $good.

02 **Can I use Sift (Chargeback) MCP to block a user?**

Yes. You can execute an action by calling `apply_user_decision` and specifying the required decision, such as 'block_user,' which immediately updates their account status.

03 **What if I need to update Sift after a dispute?**

You run `report_sift_chargeback`. This tool sends the official chargeback event details to Sift, ensuring that the user's risk profile is updated and reviewed by their models.

04 **Does Sift (Chargeback) MCP track basic activity?**

It does. You can use `track_sift_event` to log specific events, like a transaction or login, directly into Sift for behavioral analysis and ML training.

05 **How do I see what actions are available? (Sift (Chargeback) MCP)**

Run the `list_sift_decisions` tool. This shows you every recognized action or decision type that can be applied to a user within Sift.

06 **Is it possible to check past decisions on a user?**

Absolutely. Use the `list_user_decision_history` tool, and the system will retrieve all historical records of actions taken against that specific user account.

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 `"sift-chargeback": { "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

Sift (Chargeback) 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 Sift (Chargeback). 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	Sift (Chargeback) MCP
Server ID	`019d756c-885a-72e3-abc5-4d7ca68160e8`
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/sift-chargeback.