

MCP SERVER

NO CODE

CLOUD HOSTED

# Loop MCP

Manage Returns, Refunds, and Feedback Sentiment in One Conversation.

Loop brings e-commerce returns management and customer feedback collection together. Use this MCP to track product exchanges, monitor refunds, and gather actionable sentiment data via micro-surveys—all without disrupting the user experience. Connect your AI agent to Loop for comprehensive insights into why customers are leaving or what they want next.

**A+** Quality Score 98.33/100

returns-management

refund-automation

exchange-tracking

customer-feedback

nps-surveys



# 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.

# Loop MCP

10 tools available

Cloud-hosted on Vinkius

When a product fails or a customer changes their mind, you need immediate answers. This MCP connects your AI client directly to your e-commerce return and feedback data. You can manage entire lifecycles—from tracking an initial return request to analyzing the root cause of dissatisfaction. Instead of jumping between separate systems for order status, refund history, and NPS scores, your agent handles it all in a single conversation. Need to figure out if 'wrong size' is costing you more than 'late delivery'? You just ask. The system pulls together the raw feedback, calculates overall sentiment metrics, and even helps build developer tickets based on recurring product issues. By connecting through Vinkius, your AI agent gets immediate access to this entire operational suite, making complex logistics management simple enough for a chat window.

---

## Core Capabilities

### 01 — Analyze Customer Sentiment

Get overall sentiment analytics from customer feedback data.

### 02 — Process Returns and Exchanges

Track specific return requests, monitor product exchanges, and manage new order creation based on returns.

### 03 — Audit Refund Status

View refund history, including amounts processed and current status for accounting reconciliation.

### 04 — Identify Problematic Products

Access return rates, top reasons for returns, and trend data to spot product weaknesses.

### 05 — Review Feedback Sources

List all integrated sources where customers provide feedback (e.g., surveys, checkout forms).

# One Click on Vinkius — From Prompt to Execution

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

- 01 Subscribe to this MCP and enter your Loop API key.
- 02 Your AI agent accesses the connection via Vinkius.
- 03 You prompt the agent with a task (e.g., 'Show all pending refunds for Q3'), and it executes the required data retrieval.

The bottom line is, your AI client turns complex operational questions into simple conversational queries.

---

## Built For

This MCP is essential for Operations Managers and Customer Support Leads who are tired of juggling multiple spreadsheets. If you need to connect product dissatisfaction (NPS scores) directly with order logistics (refund status), this is your tool.

### E-commerce Operations Manager

Analyzes return rates and top reasons for returns, using the data to adjust inventory or update product listings.

### Customer Support Lead

Processes customer requests by checking individual return history and initiating exchanges directly through conversation.

### Product Manager

Gathers recurring feedback themes from micro-surveys to inform the next product iteration or feature build.

---

## What Changes When You Connect

- 01 Stop manually cross-referencing order logs. Your agent checks the refund history and return status instantly, giving you a single source of truth for customer accounts.

- 02 Analyze why products are failing. The MCP lets you pull detailed return analytics to pinpoint top failure reasons or most returned items, guiding product improvements.

---

- 03 Turn feedback into action. Instead of just reading comments, your agent can generate developer tickets based on recurring themes found in the data.

---

- 04 Understand customer sentiment immediately. By using `get_sentiment_metrics`, you get a quantified measure of customer satisfaction that goes far beyond simple star ratings.

---

- 05 Keep support history clean. You can use the ability to add an internal note directly to any feedback item, giving context to your team without modifying the public record.

---

---

## Real-World Applications

### Investigating a high-value customer's complaint

A support agent needs to know why Sarah returned an item. They ask their AI client, and it pulls up her entire return history using `list_feedback`, checks for any pending refunds, and identifies if the product she bought is also showing a high return rate across other customers.

### Product team identifying a flaw

The Product Manager wants to know if the new shirt design is bad. They ask their agent to list recurring feedback themes and check the general sentiment metrics, immediately confirming that 'Wrong size' is the dominant complaint.

### Quarterly operations review

The Ops Manager needs to know the biggest cost driver. They prompt their agent to show top return reasons and calculate total refunds processed last quarter, allowing them to present clear, data-backed arguments for logistics changes.

### Handling an exchange request

A customer calls about a damaged item. The support agent uses the MCP to track the product exchange status and confirm if the initial return was processed correctly, guiding the customer through creating the new order immediately.

---

## Patterns to Avoid

---

### Treating feedback data like static reports

#### X AVOID

A user downloads a spreadsheet of 500 comments and spends hours manually tallying up every instance of 'defective' or 'wrong size' to find the main problem.

#### ✓ INSTEAD

Use your AI client with this MCP. Ask it to list general feedback, then instruct it to calculate sentiment metrics and identify top return reasons in one prompt. The agent handles the counting.

---

### Ignoring process flow

#### X AVOID

A support agent only sees the initial return request but can't tell if a refund was actually issued or if an exchange was processed.

#### ✓ INSTEAD

Use the MCP to check both `list\_feedback` and then specifically ask for refund history. This confirms whether the transaction is complete, not just initiated.

---

## The Right Fit

Use this MCP when your core business challenge involves connecting customer emotion (sentiment/NPS) directly to physical product logistics (returns/refunds). If you are only focused on collecting survey data without linking it to order status, a pure feedback tool might suffice. But if the issue is 'Why did they return it?' and that reason impacts inventory or refunds, this MCP provides the necessary cross-functional view. Don't use this if your main problem is managing internal knowledge bases; stick to dedicated documentation tools. However, if you need to analyze customer input (e.g., finding themes) AND check the financial outcome of those inputs (refunds), then Loop is exactly what you need.

---

---

## The Gap Between Feedback and Fulfillment

Today, gathering feedback means logging into survey software. Checking returns status requires jumping to your order management system. To understand the full picture—like determining that a

With this MCP, you eliminate those manual hops. You tell your agent what question needs answering, whether it's total refunds processed last month or the overall sentiment score related to

high volume of 'wrong size' complaints is leading to massive refund totals—you have to copy data from three or four different dashboards and manually cross-reference everything in an Excel sheet.

a specific product line. The agent pulls that disparate data together for you instantly, giving you immediate operational clarity.

---

## Loop MCP: Insight into Product Issues

The manual process of identifying recurring issues involves reviewing hundreds of comments and then trying to categorize them—was it a defect? A size issue? Late shipping? This requires dedicated staff time just for categorization.

Now, the MCP helps you automatically identify these patterns. By listing feedback themes and checking sentiment metrics, your agent doesn't just tell you 'people are unhappy'; it tells you *why* they're unhappy, allowing you to act immediately.

---

# Loop: 10 Tools for E-commerce Operations

These tools let your AI client interact with every part of the Loop platform, from tracking specific returns to running high-level sentiment analytics.

#	TOOL	DESCRIPTION
01	<code>list_projects</code>	Retrieves a list of active projects managed within Loop.
02	<code>list_feedback_sources</code>	Lists all the external sources where customer feedback is gathered.
03	<code>list_feedback_themes</code>	Retrieves a list of common, recurring themes found in customer comments.
04	<code>list_dev_tickets</code>	Lists developer tickets that were automatically generated from feedback data for tracking.
05	<code>add_internal_note</code>	Attaches a private internal note to any specific piece of customer feedback for team context.
06	<code>get_feedback_details</code>	Retrieves all detailed information about one specific piece of submitted customer feedback.
07	<code>get_me</code>	Pulls basic account and user profile information for the connected service.
08	<code>get_sentiment_metrics</code>	Calculates and retrieves overall sentiment analytics across all collected customer feedback.
09	<code>get_ticket_details</code>	Retrieves the full details for a specific developer ticket created from product issues.
10	<code>list_feedback</code>	Lists multiple customer feedback submissions, allowing review of recent activity.

---

## See It in Action

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

### **U** Show return requests from this week and top return reasons.



This week: 23 returns. Pending: 8, Approved: 12, Rejected: 3. Exchanges: 6 (26%). Refunds: 14 (61%). Store credit: 3 (13%). Top reasons: 'Wrong size' (9, 39%), 'Not as described' (5, 22%), 'Defective' (4, 17%), 'Changed mind' (3, 13%), 'Late delivery' (2, 9%). Return rate: 4.2%.

### **U** Show return analytics and products with highest return rates.



Return analytics (30 days): 89 returns, 4.2% rate. Exchange rate: 28% (25 exchanges saved \$2,100 in revenue). Top return products: 'Slim Fit Shirt M' (12 returns, 15% rate ⚠️). 'Running Shoes 10' (8, 8%). 'Wireless Headphones' (6, 3%). Refund total: \$4,580. Average processing: 2.1 days. Customer retention post-return: 72%.

### **U** Show return history for customer sarah@company.com and pending refunds.



sarah@company.com: 2 returns (last 6 months). 1) 'Slim Fit Shirt L' — exchanged for XL ✅ (Mar 15). 2) 'Wireless Mouse' — refunded \$45 ✅ (Feb 8, defective). Lifetime orders: 12 (\$890). Return rate: 17% (above avg ⚠️). Pending refunds (all customers): 5 total, \$234. Oldest: 3 days (Sarah, \$45). Processing: automatic.

---

## Frequently Asked Questions

### 01 How does the Loop MCP help with refund tracking?

The MCP allows your agent to monitor refunds by accessing specific amounts and processing status. You can quickly confirm if a customer's money has been fully returned or is still pending.

---

**02 Can I use the Loop MCP to find out why customers are returning things?**

Yes, you can. The tool provides return analytics and top reasons for returns, giving you quantifiable data on what's driving your product returns.

---

**03 Does this MCP handle sentiment analysis from surveys?**

It does. You can call `get\_sentiment\_metrics` to get an overall numerical score of customer satisfaction across all collected feedback sources.

---

**04 What is the difference between listing feedback and getting details?**

Listing feedback shows you a summary view of multiple submissions. Getting feedback details pulls every single piece of information for just one specific item, giving deep context.

---

**05 How do I track exchanges using the Loop MCP?**

The MCP enables tracking product exchanges and can even help with creating new orders based on those exchange records, streamlining the fulfillment process.

---

# 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



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 `"loop": { "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

# Loop 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 Loop. 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	Loop MCP
Server ID	019dd11a-ac5b-73cc-8e20-83b1a282c027
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/loop](https://vinkius.com/mcp/loop).