MCP SERVER | NO CODE | CLOUD HOSTED

RudderStack MCP

Audit global data flow with chat commands

RudderStack connects your AI agent directly into RudderStack, an enterprise Customer Data Platform. Use this MCP to audit complex marketing data pipelines, check connectivity between sources and warehouses, and map user segments in plain language. Instead of digging through dashboards, you can ask for details on every configured source or destination—all from a simple chat command.

A+ Quality Score 100/100

customer-data-platform · data-pipeline · event-tracking · segmentation · data-auditing · marketing-analytics

MAINTAINED BY VINKIUS CLOUD FOR AI AGENTS · JUNE 2026
vinkius.com/mcp/rudderstack

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.

RudderStack MCP

7 tools available
Cloud-hosted on Vinkius

This integration turns your AI client into an expert data auditor. You stop guessing if your web analytics are correctly feeding your warehouse; you just ask the agent to verify the entire flow. It reads and maps every connection point, giving you immediate answers on what user events are flowing from where, and where they end up. Need to know which segments are active for remarketing? The MCP queries those defined audience clusters instantly. If a data source changes or a pipeline breaks, your agent flags it immediately. This makes the whole process of tracking customer journey reliable and auditable. By connecting this MCP via Vinkius, you gain access to enterprise-grade data governance without needing a dedicated data engineer on call.

Core Capabilities

01 — List all active data sources

You get a simple list of every configured platform that feeds data into the system.

02 — Get details for one source or destination

The MCP retrieves deep metrics and specific information about any single data point in your pipeline.

03 — List all configured destinations

You can see a complete catalog of every final location where the customer data is sent.

04 — Check which connections exist

The agent confirms if your web tracking pipelines are properly linked to their intended data warehouses.

05 — List defined user audiences

You pull a list of all active customer segments used for personalized marketing campaigns.

06 — Review event type mappings

The MCP lists the tracking plans to confirm exactly what types of user events are being captured and tracked.

One Click on Vinkius — From Prompt to Execution

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

01 Inject the RudderStack operational connector into your secure workspace.

02 Sign in using your enterprise cloud data CDP account credentials and generate a Personal Access Token.

03 Submit a natural language query, asking for specific pipeline details like 'Show me what connects to Snowflake'.

The bottom line is, you talk conversationally about complex data plumbing, and the MCP executes the necessary API calls behind the scenes.

Built For

This is for the Marketing Operations Manager who spends hours cross-referencing dashboards to confirm data flow. It's for the Data Analyst who needs to validate campaign tracking without writing SQL scripts, and the Enterprise Architect needing a quick, reliable way to audit connections across dozens of systems.

Marketing Operations Manager

They use this MCP to verify that new marketing campaigns correctly capture all necessary user events before they launch.

Data Analyst

They check if data pipelines are healthy by using the agent to confirm connections between source platforms and final destinations.

Product Manager

They audit user tracking plans to ensure that new product features generate the right kind of event data for future analysis.

What Changes When You Connect

01 Verify Data Integrity: Instead of checking multiple dashboards to see if your web events made it through, you simply ask the agent to list connections. The MCP confirms whether your sources are correctly linked to their destinations.

02 Segment Auditing: Need to confirm which user groups are active? Use `list_audiences` to instantly pull a catalog of all defined remarketing segments without manually querying the database.

03 Pinpoint Source Issues: When data seems missing, don't waste time guessing. You can use `get_source` or `list_sources` to drill down and see exactly what metrics are available from a specific platform.

04 Understand Tracking Scope: Before launching a feature, check the rules by running `list_tracking_plans`. This shows you which event schemas are enforced for your data, preventing unexpected gaps in tracking.

05 Map Endpoints Quickly: If you're onboarding a new warehouse, use `list_destinations` to see all available endpoints and confirm where the customer data can reliably land.

Real-World Applications

Checking for Missing Campaign Data

A Marketing Manager notices that campaign data from their new mobile app isn't showing up in Snowflake. They ask the agent to check the connection between 'Mobile App Source' and 'Snowflake Destination'. The MCP runs `list_connections` and reports that the pipeline is active but warns of a specific failure point, telling them exactly where they need to adjust the tracking plan.

Validating Data Governance for a New Feature

A Product Manager adds a new signup flow. Before release, they ask the agent: 'What are my current event type mappings?' The MCP uses `list_tracking_plans` to show them the strict schemas currently in place, ensuring their new feature generates data that fits existing governance rules.

Troubleshooting a Broken Pipeline

The Data Analyst sees an alert about dropped events. They ask the agent to 'Show me all sources connected to the primary warehouse.' The MCP uses `list_sources` and cross-references it with `list_connections`, quickly identifying that one specific source is failing to connect, saving hours of investigation.

Preparing for a Major Migration

An Enterprise Administrator needs to move data from an old system. They ask the agent to list all possible final endpoints using `list_destinations` and then run `get_destination` on each one, guaranteeing they don't miss any required storage locations.

Patterns to Avoid

Manual Dashboard Cross-Referencing

✗ AVOID

The analyst opens the 'Sources' tab, then switches to 'Connections', and finally opens 'Destinations'. They manually check if the names match across three different views.

✓ INSTEAD

Instead of switching tabs, ask your agent directly: 'Show me all source-to-destination connections.' The MCP runs `list_connections` in a single step, giving you one definitive list.

Guessing the Right Tool

✗ AVOID

The user sees data is missing and tries to run generic 'Check connection' queries without specifying source or destination names.

✓ INSTEAD

First, use `list_sources` to confirm the exact platform name. Then, ask for connections using `list_connections` combined with the source name. This narrows down the search immediately.

Ignoring Segmentation Rules

✗ AVOID

A marketer launches a campaign thinking they are targeting 'iOS users' but don't know if that segment was properly defined or synced.

✓ INSTEAD

Always start by asking the agent to use `list_audiences`. This confirms the exact name and status of every user cluster before building any marketing assets.

The Right Fit

Use this MCP if your primary pain point is data visibility, specifically understanding how customer event data moves from point A (a website or app) to point B (a data warehouse). If you struggle with knowing which sources are active, whether connections exist, or what the current set of defined user segments are, this is for you. Don't use it if your problem is *transforming* the data (e.g., calculating a new metric like LTV); for that, you need to write SQL or build transformation logic. If you only want to know about the structure of events and not the flow, then simply running `list_tracking_plans` might be enough, but this MCP gives you the full context of connectivity.

The headache of tracing data lineage across dozens of tools.

Today, checking if your marketing signals made it where they were supposed to is a nightmare. You open dashboard A and see the source events; you switch tabs to the connection log to verify the pipeline status; then you jump over to the destination view just to confirm the warehouse received them. This process involves clicking through three or four separate views, copying names into comparison spreadsheets, and spending hours confirming that everything is talking to everything else.

With this MCP, you simply ask your AI agent: 'Show me all data sources connected to Snowflake.' The agent runs the necessary checks across the entire CDP framework, returning a single, audited report. You get immediate certainty about your data flow without ever leaving the conversation.

RudderStack MCP gives you full visibility into your data pipeline.

Manual checks used to require running separate reports for sources, connections, and destinations. You had to manually verify that a source listed in one place matched the connection ID in another place, leading to potential human error and wasted time.

Now you can ask your agent to combine these checks into a single narrative answer. It’s not just listing data; it's giving you an auditable story of where every piece of customer information lives.

RudderStack: 7 Tools for Data Auditing

These tools allow your agent to perform deep dives into the CDP, giving you detailed information on sources, destinations, connections, and audience segments.

#	TOOL	DESCRIPTION
01	get_destination	This tool fetches specific details about a single data destination you configured in RudderStack.
02	get_source	Use this to retrieve detailed metrics and information for one specific data source.
03	list_audiences	This tool lists every defined user segment or audience group currently maintained in the CDP.
04	list_connections	The MCP uses this to list all established connections between data sources and destinations.
05	list_destinations	This tool provides a comprehensive overview of every final destination set up in RudderStack.
06	list_sources	The MCP lists all available data sources configured within the system for immediate review.
07	list_tracking_plans	This tool reviews and displays all the tracking plans that define how event data is captured in your catalog.

See It in Action

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

U List all configured sources.

I effectively efficiently accurately parsed internal records manually flawlessly pulling the actively mapped instances naturally gracefully seamlessly organically running `list_sources`. Your configuration optimally possesses smoothly accurately exactly 3 sources properly successfully securely effectively.

U Check if the connection between our website source and Snowflake destination is active.

Verified: The connection (ID: conn_12345) from 'Production Website' (source) to 'Snowflake Data Warehouse' (destination) is currently active and healthy with 0 dropped events in the last hour.

U Show me the tracking plans currently applied to our iOS app source.

The 'iOS App' source is currently linked to the 'Mobile E-commerce v2' tracking plan. This plan enforces strict schemas for 15 standard commerce events.

Frequently Asked Questions

01 **How does RudderStack MCP list all configured sources?**

You ask the agent to use `list_sources`. The MCP pulls a comprehensive catalog, letting you see exactly which platforms are currently feeding data into your CDP.

02 **Can I check if my web analytics pipeline is connected using RudderStack MCP?**

Yes. Ask the agent to list connections. It runs `list_connections` and verifies if a specific source is successfully linked to its intended destination, confirming data health.

03 What kind of information does get_source provide in RudderStack MCP?

The `get_source` tool provides detailed metrics and technical information about one particular data source. This is useful for troubleshooting or auditing specific platform configurations.

04 How do I audit customer segments using RudderStack MCP?

You use the agent to call `list_audiences`. It lists all defined user groups, which is critical for marketing teams planning personalized campaigns.

05 Is RudderStack MCP useful if I add a new data warehouse?

Absolutely. You can use `list_destinations` to view every existing endpoint and confirm where the data needs to be routed before setting up your new connection.

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 `"rudderstack": { "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

RudderStack 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

INDEPENDENT PLATFORM DISCLAIMER

Vinkius is an independent platform and is not affiliated with, endorsed by, sponsored by, verified by, or otherwise authorized by RudderStack. 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	RudderStack MCP
Server ID	`019d7600-9424-739b-b3ce-6c1e15878308`
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/rudderstack.