MCP SERVER | NO CODE | CLOUD HOSTED

Hightouch (Reverse ETL) MCP

Control your entire data sync process via conversation.

Hightouch (Reverse ETL) lets you manage data pipelines and synchronize data from your warehouse back out to SaaS tools entirely through conversation. You list all sync schedules, monitor historical runs for success rates, audit the underlying SQL models, and map every connected source or destination without ever opening a dashboard.

A+ Quality Score 100/100

reverse-etl | data-sync | warehouse-integration | data-modeling | pipeline-orchestration | saas-integration

MAINTAINED BY VINKIUS CLOUD FOR AI AGENTS · JUNE 2026
vinkius.com/mcp/hightouch-reverse-etl

SCAN TO VIEW

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.

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.

Hightouch (Reverse ETL) MCP

6 tools available
Cloud-hosted on Vinkius

Managing enterprise data flows means connecting Snowflake or BigQuery to external apps like Salesforce or HubSpot, and that process is complex. This MCP gives your agent full control over those Reverse ETL workflows using natural language. You don't need to click through multiple tabs or read error codes in a restrictive UI. Instead, you simply ask questions: "Show me the run history for sales data," or "Which SaaS apps are receiving objects from our warehouse?" Your agent answers immediately, providing detailed metrics and schemas on demand.

This ability to audit complex data models and sync performance via chat changes how quickly your team can validate data integrity. When you connect this MCP through Vinkius, you gain instant access to the entire catalog of tools needed for robust data operations right from your preferred AI client.

Core Capabilities

01 — Audit Data Destinations

List every connected SaaS application (like Salesforce or HubSpot) that is currently receiving synchronized objects.

02 — Check Synchronization Schedules

Retrieve a list of all active data sync jobs and get details on their schedules and underlying schemas.

03 — Review Run Performance History

Extract detailed records of past execution states, allowing you to monitor overall reliability and success logs for any job.

04 — Inspect Data Models

Audit the SQL definitions used in your warehouse analytics pipelines that drive data synchronization.

05 — Map Internal Sources

Enumerate and securely map all active internal data warehouse connections, such as Snowflake or BigQuery.

One Click on Vinkius — From Prompt to Execution

Available at **vinkius.com/mcp/hightouch-reverse-etl** — connect your AI agent in three steps.

01 Subscribe to this MCP and provide your Hightouch API Key.

02 Connect the key to any MCP-compatible client (Claude, Cursor, etc.).

03 Ask your agent specific questions about sync health or data models.

The bottom line is you manage complex data synchronization tasks by talking to it, not by clicking through dashboards.

Built For

This MCP is for the Data Engineer who spends hours digging into dashboard logs or the Growth Ops Manager who needs immediate confirmation that critical sales metrics hit every destination. It's for people whose daily job requires auditing complex, moving data pipelines.

Data Engineer

Uses this MCP to check sync run histories and audit SQL model definitions without opening the graphical interface.

Analytics Engineer

Verifies warehouse source connections and ensures data models are mapped correctly to outbound destinations before deployment.

Growth Operations Manager

Checks sync run histories and success rates to guarantee that key marketing or sales data is current across all connected apps.

What Changes When You Connect

01 Stop manually opening dashboards. You can check run history and get success logs simply by asking your agent for it, saving significant time.

02 Audit every connected destination—from Salesforce to Slack—using list_destinations. Know exactly where your synchronized data is going without guessing.

03 When a sync fails, don't wait for alerts. Use the run history tools (list_sync_runs) to immediately see past execution states and diagnose performance issues.

04 Verify data source connections instantly. You can list all sources using list_sources to confirm if your warehouse link is active and mapped correctly.

05 Deeply inspect pipeline logic by listing models. This lets you audit the underlying SQL definitions that drive complex data transformations.

Real-World Applications

Investigating a Missing Metric in Salesforce

The Growth Ops Manager notices sales numbers are stale in HubSpot. They ask their agent to check the run history for the specific sync, using list_sync_runs. The agent confirms the last 10 runs succeeded but points out that the underlying model (list_models) was recently updated and needs review.

Onboarding a New Data Warehouse Connection

The Analytics Engineer needs to confirm if BigQuery is ready for outbound sync. They use list_sources to enumerate all active connections, confirming the schema map from the new platform before writing any data flows.

Auditing Compliance and Data Flow

A compliance officer asks their agent: "Which apps receive PII?" The agent uses list_destinations to identify all connected SaaS nodes, providing an instant audit trail of where sensitive data is flowing.

Troubleshooting a Broken Sync

The Data Engineer notices the 'Marketing' sync hasn't updated. They ask the agent for details on the specific job (get_sync) and find that while the sync was scheduled, the underlying data model definition needs adjustment.

Patterns to Avoid

Treating it like a simple database query

✗ AVOID

Trying to ask for raw data records (e.g., 'Give me all user emails') instead of asking about the flow's health.

✓ INSTEAD

Remember this MCP manages *flows*, not raw storage. Use list_syncs to understand which schedules are running, or get_sync to check if a specific connection is active.

Assuming data flows automatically connect

✗ AVOID

Thinking that because Snowflake is connected, all external destinations like Slack will receive data without explicit mapping.

✓ INSTEAD

Always confirm the flow path. Use list_destinations to see exactly which SaaS apps are receiving outbound objects and verify they're configured correctly.

Ignoring run context

✗ AVOID

Seeing a green light on a dashboard but having no record of when that data was actually pulled or if it passed through necessary transformations.

✓ INSTEAD

Don't trust the surface status. Use list_sync_runs to get the detailed history and confirm the timestamps and success metrics for critical runs.

The Right Fit

Use this MCP if your primary pain point is managing, auditing, or troubleshooting complex data movement between a warehouse (like Snowflake) and multiple SaaS destinations. You need an AI layer that can read the entire state of your ETL process—the models, the schedules, the sources, and the destinations—and answer questions about it in plain language.

Don't use this MCP if you just need to run a simple ad-hoc query (use a standard SQL connection tool) or if your data flow is linear and doesn't involve multiple complex models. You also won't need it if all your data resides in one system; this is specifically for *reverse* ETL—sending data back out.

If you are constantly opening the Hightouch UI to check status, then yes, this MCP saves you hours of clicking.

The Data Dashboard Drilldown Nightmare

Today, checking if your critical sales data made it from Snowflake to HubSpot means opening the Hightouch UI. Then you navigate to 'Syncs,' click on the relevant job, find the run history tab, and finally check the status logs—all while hoping nothing is broken in between.

With this MCP, you just ask your agent: "What's the status of our HubSpot sync?" It pulls all that complex diagnostic information into one conversational answer. You get an immediate performance summary without touching a single dashboard tab.

Hightouch (Reverse ETL) Gives You Full Data Flow Control

You no longer have to manually cross-reference source connections, model definitions, and destination lists across three different menus just to validate a data point. The agent aggregates all this metadata for you.

It's not about reading logs; it's about getting answers. You gain immediate, conversational visibility into the entire lifecycle of your reverse ETL process.

Hightouch (Reverse ETL) MCP - 6 Tools

These tools let you manage every part of the data synchronization cycle: listing sources, checking schedules, monitoring runs, and auditing models.

#	TOOL	DESCRIPTION
01	`list_destinations`	Lists every SaaS application node that receives synchronized objects from your warehouse.
02	`list_syncs`	Retrieves a list of all current data synchronization schedules configured in Hightouch.
03	`get_sync`	Fetches specific, detailed information about an individual sync job.
04	`list_sync_runs`	Retrieves historical records of all data synchronization runs for performance monitoring.
05	`list_models`	Provides a list and audit of the SQL definitions used in your analytics pipelines.
06	`list_sources`	Enumerate all active internal data warehouse connections, such as Snowflake or BigQuery.

See It in Action

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

U List all my current Hightouch syncs

I've found 5 active syncs. Highlights include 'Snowflake to Salesforce' (ID: sync-001), 'BigQuery to HubSpot' (ID: sync-004), and 'Postgres to Slack' (ID: sync-007). Would you like the status of the most recent run for any of these?

U Show me the last 5 runs for sync ID 'sync-001'

Retrieving run history for 'sync-001'... The last 5 runs all finished with a 'Success' status. A total of 12,450 rows were processed in the most recent run (finished 2 hours ago). Would you like to see if any errors were logged in previous weeks?

U Which data sources are connected to our Hightouch account?

I've identified 3 connected data sources: 'Production Snowflake', 'Stripe BigQuery Dataset', and 'Internal Postgres'. All connections are showing as 'Active'. Would you like to see the available models for 'Production Snowflake'?

Frequently Asked Questions

01 **How do I check if a specific data sync is running with Hightouch? (list_syncs)**

You use list_syncs to get an immediate overview of all active schedules. This lets you see which pipelines are currently scheduled or need attention across your entire account.

02 **What does the MCP help me audit about my data models? (list_models)**

It allows you to list and inspect the underlying SQL definitions that power your syncs. This is critical for knowing exactly how your raw warehouse data gets shaped before it leaves.

03 **Can I see if a destination like Salesforce is connected? (list_destinations)**

Yes, you use list_destinations to identify all connected SaaS nodes. This confirms that the endpoint is ready and receiving objects from your warehouse.

04 **How can I check if a data source connection is working? (list_sources)**

You ask the agent to list sources, which enumerates all active connections like Snowflake or BigQuery. This verifies that the initial internal link feeding your process is stable.

05 **If I have a sync error, how do I find out why? (list_sync_runs)**

Use list_sync_runs to pull historical execution states and success logs. This gives you the performance metrics needed to pinpoint exactly when and where an issue started.

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 `"hightouch-reverse-etl": { "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

Hightouch (Reverse ETL) 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 Hightouch (Reverse ETL). 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	Hightouch (Reverse ETL) MCP
Server ID	`019d75b1-8a27-72a0-9de9-cad223388e37`
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/hightouch-reverse-etl.