MCP SERVER | NO CODE | CLOUD HOSTED

Matillion (Cloud Data Integration & ELT) MCP

Audit, Orchestrate, and Monitor Your Entire Data Stack

Matillion (Cloud Data Integration & ELT) gives your AI client direct control over complex data workflows. Audit pipelines, check execution statuses across multiple cloud environments like Snowflake and BigQuery, and manage all aspects of your enterprise ELT orchestration using natural conversation.

A+ Quality Score 100/100

elt-pipelines | data-integration | workflow-orchestration | data-transformation | cloud-data-warehouse

MAINTAINED BY VINKIUS CLOUD FOR AI AGENTS · JUNE 2026
vinkius.com/mcp/matillion-cloud-data-integration-elt

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.

Matillion (Cloud Data Integration & ELT) MCP

6 tools available
Cloud-hosted on Vinkius

This MCP connects your AI agent directly to the Matillion Data Productivity Cloud. You take full command over your entire data integration lifecycle without needing to navigate complex dashboards or write API calls. Instead of logging into a web console just to check if 'Customer-360' ran successfully, you ask your agent and get an instant status report. It allows you to list every managed ETL pipeline, audit their structures, monitor active runtime agents on your local network, and even verify which cloud data warehouse environments are configured—whether that's Redshift, BigQuery, or Snowflake. This level of infrastructure visibility is crucial for any serious data team. By connecting through Vinkius, you give your AI client the comprehensive view it needs to treat your entire data stack like one unified system.

Core Capabilities

01 — View all ETL pipelines

List every managed pipeline container and its associated ID within Matillion.

02 — Deep-dive into pipeline structure

Retrieve the complete underlying orchestration definition for a single, specified data flow.

03 — Track run history and status

Monitor recent pipeline executions to see which ones succeeded, failed, or are currently running.

04 — List configured cloud destinations

Enumerate every data warehouse environment attached to the Matillion instance (e.g., Snowflake, BigQuery).

05 — Audit local runtime infrastructure

Check the operational status and count of all active Matillion runtime agents across your network.

06 — Manage project groupings

List high-level project containers that bind related pipelines and environments together for organization.

One Click on Vinkius — From Prompt to Execution

Available at **vinkius.com/mcp/matillion-cloud-data-integration-elt** — connect your AI agent in three steps.

01 Subscribe to this MCP on Vinkius.

02 Input your Matillion API URL, Client ID, and Client Secret into the connection settings.

03 Ask your AI client a question—like 'What are my active cloud environments?'—and it executes the necessary data calls.

The bottom line is you use plain language to manage complex enterprise data infrastructure that usually requires dedicated developer tools.

Built For

This MCP is for anyone who spends too much time clicking through dashboards or manually checking logs across multiple cloud systems. It's the Data Engineer whose day is spent debugging failures, the BI Analyst needing immediate environment confirmation, and the Ops Lead who needs a single source of truth for agent health.

Data Engineer

Debugging pipeline failures by checking run history with list_executions or verifying structural components using get_pipeline.

BI Operations Analyst

Confirming data reliability by listing all environments (list_environments) and auditing the status of local runtime agents (list_agents).

Analytics Lead

Gaining an overview of project health and identifying which pipelines need updating without manual hub navigation.

What Changes When You Connect

01 Instant failure analysis. Instead of clicking into a dashboard to find out why data didn't move, use list_executions to see if the pipeline failed and get immediate status reports.

02 Full infrastructure visibility. Easily audit your local network health by listing active agents with list_agents, ensuring that critical ELT processes can run without interruption.

03 Data mapping accuracy checks. List all environments (list_environments) to confirm every destination—whether it's Snowflake or Redshift—is correctly configured and mapped for the data.

04 Effortless structure review. Need to know exactly what a pipeline does? Use get_pipeline to retrieve its full underlying orchestration definition without ever touching the Matillion UI.

05 Project overview at a glance. List all projects (list_projects) so you can immediately understand the scope and boundaries of your data transformation efforts.

Real-World Applications

Debugging a sudden data outage

An engineer notices dashboard data is missing. They ask their agent to run list_executions for the 'Sales-Sync' pipeline, immediately identifying that the last run failed and retrieving specific error codes, solving the problem in seconds instead of hours.

Preparing a new cloud destination

A BI lead needs to know if their new BigQuery dataset is ready. They ask the agent to list_environments, confirming the target credentials exist before running any data migration tests.

Auditing compliance and governance

An operations analyst must prove that all transformation logic for a key dataset is documented. They use get_pipeline on the relevant ID to pull the full schema definition, satisfying audit requirements instantly.

Scaling out data processing capacity

A team needs to confirm if their local compute resources are sufficient. They ask the agent to list_agents, checking the count and status of all runtime agents before initiating a large-scale batch job.

Patterns to Avoid

Manually navigating complex UIs

✗ AVOID

Logging into Matillion, clicking through 'Pipelines,' then selecting the correct ID, and finally clicking 'Run History' to check status.

✓ INSTEAD

Just ask your agent directly. Say: 'Show me the recent execution statuses for all pipelines.' The MCP handles the clicks using list_executions.

Assuming environment readiness

✗ AVOID

Running a data load job and getting an obscure error code about missing connection parameters, forcing manual database checks.

✓ INSTEAD

Always start by asking to list_environments. This confirms the necessary cloud destination credentials are set up before you run any loads.

Forgetting which pipelines exist

✗ AVOID

Spending an hour searching through multiple project folders trying to find the specific 'Customer-360' pipeline ID.

✓ INSTEAD

Simply ask the agent to list_pipelines. It gives you a comprehensive, searchable inventory of every workflow available.

The Right Fit

Use this MCP if your primary pain point is observability and operational control over complex, multi-cloud data workflows. You need an AI layer that can read the state of your ELT infrastructure—like checking agent status via list_agents or verifying every environment via list_environments. Don't use it if you just need to query raw data; this MCP manages *how* the data moves, not the data itself. If all you need is a simple API endpoint for one single piece of information, a direct API connector might suffice. But if you need visibility across pipelines, environments, and agents—the whole picture—this is your tool.

The Hidden Cost of Dashboard Hopping

Today, checking the health of a data pipeline means logging into Matillion, finding the project container, clicking on the specific dataset's ID, then navigating to the 'Execution Logs' tab. If you have five different pipelines running across three clouds, that’s at least 30 minutes of painful clicking and copy-pasting IDs just to get a status report.

With this MCP, all that vanishes. You simply tell your agent, 'What's the status of my core data flows?' The system instantly aggregates execution statuses from every pipeline and environment in one conversational response.

Get Complete Control with list_pipelines

Before, figuring out which pipelines existed meant navigating nested folder structures and manually checking project documentation to find the correct IDs. If a team member changed a pipeline name or moved it, you were blind until they told you.

Now, asking the agent to list_pipelines gives you an instant inventory of every single data flow. You know exactly what exists, where it lives, and who owns it—all in one simple request.

Matillion (Cloud Data Integration & ELT) - 6 Tools

Use these tools to list pipelines, check execution status, audit environments, and monitor the physical agents running your data workflows.

#	TOOL	DESCRIPTION
01	`list_pipelines`	Lists all ETL pipelines managed within your Matillion instance.
02	`get_pipeline`	Retrieves the detailed structural components for a single, specified pipeline ID.
03	`list_executions`	Shows a list of recent data pipeline runs and their current status (success or failure).
04	`list_environments`	Displays all configured cloud destination environments attached to the Matillion hub.
05	`list_agents`	Enumerates and reports on the status of active Matillion runtime agents across your network.
06	`list_projects`	Lists all major project containers that group related pipelines together for organization.

See It in Action

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

U List all Matillion ETL pipelines in my account

I've retrieved your ETL pipelines from the Data Productivity Cloud. Highlights include 'Sales-Data-Sync' (ID: pipe-123), 'Customer-360-View' (ID: pipe-456), and 'Finance-Consolidation' (ID: pipe-789). Would you like to see the structural components for the sales sync?

U Show me the last 5 pipeline executions and their status

Retrieving recent executions... I've identified the 5 latest runs: 3 succeeded (including 'Inventory-Load'), 1 is currently 'Running' ('Marketing-Attribution'), and 1 'Failed' ('Web-Log-Ingestion'). Would you like me to investigate the error for the failed run?

U What cloud environments are configured in my Matillion instance?

I've identified 3 destination environments: 'Snowflake-Prod' (Targeting: SF_WH_01), 'Redshift-Staging' (Targeting: AWS_NODE_ALPHA), and 'BigQuery-Analytics' (Targeting: GCP_DATASET_V1). I can provide the project mappings for these environments if you'd like.

Frequently Asked Questions

01 **Can Matillion (Cloud Data Integration & ELT) MCP list all my pipelines?**

Yes. The list_pipelines tool gives you a complete inventory of every ETL pipeline container managed within your account, allowing you to see everything at a glance.

02 **How do I check if a pipeline failed using Matillion (Cloud Data Integration & ELT) MCP?**

Use list_executions. This tool retrieves recent run statuses and will clearly indicate which deployments succeeded, are running, or have outright failed.

03 **What is the difference between listing environments and checking agents with Matillion (Cloud Data Integration & ELT) MCP?**

list_environments shows where your data *goes* (Snowflake, BigQuery), while list_agents tells you about the actual local *compute power* that runs the job.

04 **Does Matillion (Cloud Data Integration & ELT) MCP help me with project grouping?**

Yes. The list_projects tool lets you view and understand how related pipelines and environments are grouped together within logical containers.

05 **Can I get the structure of a pipeline using Matillion (Cloud Data Integration & ELT) MCP?**

You can use get_pipeline. This tool doesn't just give you the name; it pulls the actual, deep structural definition of the data flow for detailed review.

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 `"matillion-cloud-data-integration-elt": { "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

Matillion (Cloud Data Integration & ELT) 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 Matillion (Cloud Data Integration & ELT). 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	Matillion (Cloud Data Integration & ELT) MCP
Server ID	`019d75d0-1c58-7144-bfe8-0251e27acbe3`
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/matillion-cloud-data-integration-elt.