MCP SERVER | NO CODE | CLOUD HOSTED

Integrate.io MCP

Manage Data Pipelines Through Conversation

Integrate.io (ETL & Data Integration) lets you manage complex data pipelines and ETL jobs using natural conversation. List all your active packages, track job runs for failures, audit every connection, or check account credit limits—all without leaving your AI client.

A+ Quality Score 100/100

etl-pipelines | data-integration | pipeline-monitoring | data-transformation | data-warehouse | job-scheduling

MAINTAINED BY VINKIUS CLOUD FOR AI AGENTS · JUNE 2026
vinkius.com/mcp/integrateio-etl-data-integration

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.

Integrate.io (ETL & Data Integration) MCP

6 tools available
Cloud-hosted on Vinkius

You can take full control of your automated data workflows through this MCP. Instead of logging into a web dashboard to manage data pipelines, you talk to your agent and get the answers instantly. Need to know what data moved last night? Your agent runs `list_jobs` and tells you if the sync succeeded or failed. Want to verify that the 'Stripe' connection is still pointing to the right database? You simply ask it to list connections, checking credentials in seconds. It also lets you check account limits with a simple call to get your overall status, helping you manage your budget while running complex data transformations. This MCP connects deep infrastructure actions—like listing all data transformations or retrieving specific pipeline details—directly into your workflow via Vinkius. You use this for everything from checking job history to inspecting schemas.

Core Capabilities

01 — List active pipelines

View every scheduled data package in the Integrate.io account with one command.

02 — Get specific pipeline details

Retrieve a deep dive into the structure, nodes, and variables of any single data pipeline by its ID.

03 — Monitor job history and status

Track the success or failure status of past and current automated jobs to confirm data warehouse updates ran correctly.

04 — Audit data connections

Enumerate all linked database credentials and API sources used across your entire data infrastructure.

05 — Inspect transformation logic

List and review the detailed mapping rules for every data transformation you've set up in your account.

06 — Check account limits

Get real-time status on your workspace credits, remaining usage, and overall account metrics.

One Click on Vinkius — From Prompt to Execution

Available at **vinkius.com/mcp/integrateio-etl-data-integration** — connect your AI agent in three steps.

01 First, subscribe to this MCP and provide your Integrate.io API Key.

02 Next, connect the credentials from any compatible AI client (like Cursor or Claude).

03 You can then use natural conversation prompts to run commands like listing pipelines or checking job statuses.

The bottom line is you control complex data infrastructure directly through chat, eliminating dashboard hopping.

Built For

This MCP is for the Data Engineer tired of jumping between monitoring dashboards. It's for the Analytics Lead who needs to prove data integrity quickly, and the Operations Analyst needing instant visibility into cost centers like API usage.

Data Engineer

Uses this MCP to check job history or get details about a specific pipeline when debugging an ETL failure.

Analytics Lead

Audits data transformations and connections before running major reports to guarantee data quality for business stakeholders.

Operations Analyst

Tracks account status, checking remaining credits or listing all API connections to optimize usage and manage budget.

What Changes When You Connect

01 Stop jumping between tabs. Instead of logging into the dashboard just to check if your nightly 'Stripe Sync' succeeded, ask your agent to run `list_jobs` and get the status immediately.

02 Audit data integrity instantly. Need to confirm which sources feed a report? Use `list_connections` to inventory every database and API source before running any major analysis.

03 Control your budget in real-time. Instead of guessing where your credits are going, use `get_account` to see exactly how many processing units you have left for the month.

04 Understand complex data flows quickly. When a pipeline fails, don't just get an error code; ask for details using `get_pipeline` to see which specific node or variable caused the issue.

05 Verify your setup before deployment. Use `list_transformations` to inspect every mapping logic and ensure the source data is being correctly converted into the target schema.

Real-World Applications

The Data Engineer needs a failure root cause.

A nightly job fails, leaving the warehouse empty. Instead of manually clicking through logs and dashboards for hours, the engineer asks the agent to run `list_jobs` then uses `get_pipeline` on that specific pipeline ID. The agent immediately pulls up the schema details and shows where the connection variable is failing.

The Operations Analyst is managing cost overruns.

Billing seems high. The analyst asks the agent for the account status using `get_account`. The response shows low remaining credits, prompting them to review all active sources by running `list_connections` and find an unused API key.

The Analytics Lead must prove data lineage.

A VP asks, 'How did we get this revenue number?' The lead doesn't know the exact path. They ask the agent to run `list_transformations` and then use `list_connections`. This quickly maps out every source and rule used to calculate the final metric.

Patterns to Avoid

Checking data flow manually.

✗ AVOID

A user has to log into the web UI, navigate to 'Jobs,' filter by date range, download a CSV of results, and then manually check connection names one by one.

✓ INSTEAD

Instead, simply ask your agent to `list_jobs` for the last 7 days. If you need details on the data source, run `list_connections`. This keeps everything in the chat interface.

Debugging pipelines with partial info.

✗ AVOID

The user gets an error message 'Invalid variable' but has no idea which pipeline or schema is involved, forcing them to re-read documentation and guess at the source ID.

✓ INSTEAD

First, use `list_pipelines` to confirm the correct package name. Then, run `get_pipeline` on that specific ID to view all associated schemas and variables in one place.

Assuming connections are active.

✗ AVOID

The data runs fine for a month, but suddenly fails with an 'Auth Failed' error. The user wastes time checking the source application instead of the credentials.

✓ INSTEAD

Before running anything critical, always run `list_connections`. This confirms that the stored database and API keys are still valid and accessible to your agent.

The Right Fit

Use this MCP if you manage data flow through structured ETL/ELT pipelines that involve multiple steps, schemas, and external connections. If tracking job run history, auditing credentials, or listing complex transformations is a regular part of your day-to-day work, this tool saves time. Don't use it if you just need to perform a simple API call (like fetching a single user record) or move a file from Point A to Point B without transformation logic; for those tasks, a specialized messaging MCP would be better. However, if your task involves checking the status of orchestrated, multi-step data packages—that's where this excels.

Data Infrastructure is a Black Box Without Central Visibility

Right now, managing automated data pipelines means jumping through hoops. You open the dashboard to check job status, then click another tab to see connection credentials. If you need to verify transformation logic, you often have to pull up a separate documentation sheet and cross-reference everything manually. It’s slow, it's tedious, and when something breaks at 2 AM, you waste valuable time just figuring out *where* to start looking.

With this MCP, all that data—the jobs, the connections, the transformations—is accessible through natural conversation. You simply ask your agent about what happened or why something broke, and it gives you a full report in text format. The whole process happens without ever leaving your AI client.

Integrate.io (ETL & Data Integration) Gives You Full Control

Manual data checks vanish when you use the `list_pipelines` tool to see every job package at a glance, and then run `get_pipeline` to understand its exact schema structure without clicking through multiple pages.

You get an immediate, conversational overview of your entire data stack. It’s not just monitoring; it's managing the whole lifecycle—from connection audit (`list_connections`) to job history review (`list_jobs`).

Integrate.io (ETL & Data Integration) MCP - 6 Tools

These tools let you interact with your data infrastructure to list pipelines, track jobs, check connections, and audit transformations using natural language commands.

#	TOOL	DESCRIPTION
01	list_pipelines	Shows you a list of every data pipeline currently configured in your Integrate.io account.
02	get_pipeline	Pulls all the technical details, schemas, and variables for one specific pipeline using its ID.
03	list_jobs	Retrieves a log of both successful and failed jobs to see when your data ran last.
04	list_connections	Provides an inventory of all database credentials and API connections you've set up.
05	list_transformations	Lists every data transformation rule in place, allowing you to audit how data is mapped from source to target.
06	get_account	Checks your overall Integrate.io account status, including limits and remaining processing credits.

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 data pipelines in Integrate.io

I've found 5 active pipelines: 'Shopify to Snowflake', 'Stripe Data Sync', 'Customer CRM Export', 'Marketing Analytics', and 'Inventory Hub'. Which one would you like to see the details or run history for?

U Show me the status of the last 3 jobs

Retrieving job history... The last 3 jobs were: 1) 'Shopify Sync' (Success, 12 min ago), 2) 'CRM Export' (Failed, 1 hour ago - connection timeout), 3) 'Stripe Sync' (Success, 3 hours ago). Would you like to inspect the error for the failed job?

U What data sources are currently connected to my account?

I've identified 4 connected data sources: 'Production PostgreSQL', 'Cloud Snowflake DW', 'Stripe API', and 'HubSpot CRM'. All connections are active and ready for pipeline use. Would you like to see the transformation models mapped to 'Snowflake'?

Frequently Asked Questions

01 **How do I check if a specific ETL job ran successfully using Integrate.io (ETL & Data Integration)?**

You use the `list_jobs` tool to see the history of runs. This shows you success status, failure times, and which pipelines were involved in the run.

02 **Can I list all my data sources with Integrate.io (ETL & Data Integration)?**

Yes, running `list_connections` pulls an inventory of every database and API connection you have set up for your pipelines.

03 **How do I see the details of a specific data pipeline?**

Use the `get_pipeline` tool. You must provide the unique ID, and the agent will return all technical specifics like schemas and variables associated with that package.

04 **What is the best way to check my remaining Integrate.io credits?**

The `get_account` tool provides a real-time view of your account status, including current usage and remaining processing credits so you don't hit a spending limit.

05 **Does Integrate.io (ETL & Data Integration) help me audit data transformations?**

Yes, running `list_transformations` shows you every mapping rule established in your account. This is crucial for verifying data quality logic.

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 `"integrateio-etl-data-integration": { "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

Integrate.io (ETL & Data Integration) 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 Integrate.io (ETL & Data Integration). 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	Integrate.io (ETL & Data Integration) MCP
Server ID	`019d75ba-5d47-72ac-af7b-40902c618d34`
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/integrateio-etl-data-integration.