MCP SERVER | NO CODE | CLOUD HOSTED

Airbyte MCP for AI Agents

Monitor and audit data pipeline connections in real time

Airbyte MCP lets your AI agent monitor entire data integration pipelines. Check sync job status, list all active sources (like Postgres or Stripe), and audit configured destinations (such as Snowflake or BigQuery) instantly via conversation. It keeps your modern data stack running without you touching a dashboard.

F Quality Score 3.6/100

etl-pipelines | data-integration | data-warehousing | pipeline-monitoring | data-sync | data-engineering

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

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.

Airbyte MCP

7 tools available
Cloud-hosted on Vinkius

Your AI agent can talk directly to your Airbyte instance, giving you conversational visibility into every part of your ETL/ELT process. Instead of logging into the dashboard and clicking through pages just to see if everything ran overnight, your agent handles the audit automatically. You tell it what you need—like checking yesterday's Postgres sync rate or listing all destinations pointing to Snowflake—and get a clean answer back immediately. It's like having a dedicated data ops engineer on standby 24/7. This MCP connects that oversight capability directly into Vinkius, making your whole data flow visible through any compatible AI client.

Core Capabilities

01 — List all configured sources

Retrieves a full list of every data origin (sources) you've connected in Airbyte.

02 — Get details for a specific source

Pulls detailed configuration and status information for one particular data source.

03 — List all destinations

Provides a comprehensive list of every target warehouse or destination configured in Airbyte.

04 — List active sync connections

Shows all the established data pipelines (connections) that move data from sources to destinations.

05 — Get connection details

Fetches specific details, configuration, and status for a single data synchronization connection.

06 — Track job history and success rates

View historical records of sync jobs for any given connection, detailing success or failure.

07 — List Airbyte workspaces

Retrieves a list of all active workspace environments within your Airbyte account.

One Click on Vinkius — From Prompt to Execution

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

01 Subscribe to this MCP and provide your specific Airbyte API URL and API Key.

02 Your AI client runs diagnostic queries, asking for pipeline status or connection details via the exposed tools.

03 The agent returns structured data—like job history or source lists—which it presents back to you in plain language.

The bottom line is your AI can act as a constant monitor, querying Airbyte's operational state without needing manual dashboard interaction.

Built For

This MCP is built for the data team. Data Engineers who spend too much time debugging failed sync jobs, and Analytics Engineers who need quick verification of warehouse paths—you're the target user.

Data Engineer

Needs to check yesterday's sync job success rate or debug a failing database connection across multiple sources.

Analytics Engineer

Quickly lists configured warehouse destinations, like Snowflake and BigQuery, and verifies infrastructure paths for new reports.

Data Analyst

Needs a high-level summary of all active data sources feeding into the main data lake without diving into technical dashboards.

What Changes When You Connect

01 Stop checking dashboards manually. Your agent directly queries the job history using `list_jobs` to tell you instantly if a nightly sync failed.

02 Get a full inventory of your infrastructure by running `list_sources` and `list_destinations`, giving you immediate visibility into all data origins and targets.

03 Quickly troubleshoot connectivity issues. Use `get_connection` to pull detailed status for a specific pipeline, saving minutes of dashboard clicking.

04 Understand the whole scope of your setup by calling `list_connections`. You see every active path from source to warehouse at a glance.

05 Maintain environment oversight by running `list_workspaces`, confirming that all operational environments are correctly configured.

Real-World Applications

The nightly Postgres sync failed

A data engineer asks their agent, 'What was the status of the Postgres source connection last night?' The agent runs `list_jobs` and tells them exactly which job failed, why it timed out, and when the previous run succeeded.

I need to audit our warehouse targets

An analytics engineer asks, 'Show me every destination we've pointed data towards.' The agent uses `list_destinations` and returns a clean list of all configured endpoints like Snowflake and BigQuery.

Which sources are currently connected?

A manager needs to know what systems feed the data lake. They ask, 'List all active data origins.' The agent uses `list_sources` to provide a clean count and list of everything from Postgres to Stripe.

Verify connection paths for new projects

A team member asks the agent to summarize current pipelines. The agent calls `list_connections`, providing a comprehensive overview of all data movement paths currently running.

Patterns to Avoid

Treating Airbyte like a simple spreadsheet

✗ AVOID

Assuming that just knowing the source name is enough. You might only run `list_sources` and miss critical job status details.

✓ INSTEAD

Always follow up by running `list_connections` to see how that source is actually used, then use `list_jobs` on a specific connection ID for real-time operational data.

Ignoring the workspace context

✗ AVOID

Running diagnostics without knowing which environment you're in. This could lead to checking the wrong set of credentials.

✓ INSTEAD

First, always call `list_workspaces` to confirm the active workspace before running any diagnostic tools like `get_connection`.

Focusing only on connectivity

✗ AVOID

Just listing sources and destinations without checking if data actually moved. You could have a perfect setup, but no running jobs.

✓ INSTEAD

After confirming all connections exist using `list_connections`, immediately run `list_jobs` to validate that the pipelines are actively syncing data.

The Right Fit

Use this MCP if your pain point is knowing *why* a pipeline failed or needing an up-to-date inventory of all connected systems. You need continuous, auditable operational visibility into your ELT jobs, not just the static setup details. Don't use it if you only need to change credentials; for that, you still have to manually update the Airbyte UI. If you are struggling with *which* source connects to *which* destination, start by running `list_sources` and `list_destinations` together. This helps narrow down the scope before calling `list_connections`.

Airbyte MCP for AI Agents: Monitoring Data Sync Job History

Right now, checking if your data pipelines ran correctly means jumping into a dashboard. You have to manually click through connection names, look at the job status column, and then cross-reference timestamps to figure out which runs succeeded and which failed. It's slow, tedious work that happens every time there's an alert.

With this MCP, your AI agent handles the entire audit. You simply ask it about a specific connection or source, and it uses tools like `list_jobs` to pull historical success rates and failure reasons directly into the chat window. It gives you immediate, actionable status reports without ever touching the Airbyte UI.

Airbyte MCP for AI Agents: Auditing Data Source and Destination Connections

Manually tracking your data architecture involves logging into multiple places—one place to list sources (Postgres, Stripe), another to see destinations (Snowflake, BigQuery). Keeping a centralized inventory of these endpoints is a constant chore that leads to outdated documentation.

This MCP fixes that by consolidating the view. You ask for all available origins and targets using `list_sources` or `list_destinations`, getting a single, verified list right where you're working with your agent. It gives you total control over your data footprint.

7 Tools for Monitoring Airbyte Data Pipeline Sync Jobs

Use these tools to list all sources, destinations, and track the job history of any connection in your data stack.

#	TOOL	DESCRIPTION
01	`list_sources`	Lists every available source system connected to Airbyte.
02	`get_source`	Retrieves the detailed configuration for a single data source.
03	`list_destinations`	Lists all configured target destinations where data is sent.
04	`list_connections`	Lists every active pipeline connection between sources and destinations.
05	`get_connection`	Gets the specific details for one established data sync connection.
06	`list_jobs`	Lists job history, showing when a connection last ran and whether it succeeded or failed.
07	`list_workspaces`	Retrieves the list of separate workspaces within your Airbyte environment.

See It in Action

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

U What were the sync results for my Stripe connection last week?

Stripe Connection Job History

- **Date:** 2026-06-12 (Yesterday)
* Status: **SUCCESS**
* Records Sync'd: 1,450
* Duration: 3m 12s
- **Date:** 2026-06-11 (Day Before)
* Status: **FAILED**
* Reason: Missing authentication token for the billing API.
* Records Sync'd: 0

The failure on June 11th was due to a credential issue. You need to update your Stripe credentials in Airbyte.

U List all active destinations and check if Snowflake is configured.

You have **4** active data destinations:

- ✅ **Snowflake Data Cloud**: (ID: ...a12) Used for core reporting.
- ✅ **Google BigQuery Prod**: (ID: ...b3c) Used for marketing analytics.
🟡 ***Amazon S3 Bucket****: (ID: ...d4e) Used for raw JSON dumps.* Status check needed.*
- ❌ **Legacy Redshift Cluster**: (ID: ...f5g) Connection is listed but appears inactive.

Snowflake looks good to go, but I recommend checking the status on your S3 bucket.

U Show me all data origins and check their configuration IDs.

Airbyte Sources Found (Total: 5)

1. **Postgres DB**: ID `pg-prod-001`. Status: Active.
2. **Stripe API**: ID `stripe-live`. Status: Active, last sync 3 hours ago.
3. **Salesforce CRM**: ID `sf-uat-test`. Status: Warning (Credentials nearing expiry).
4. **Google Sheets**: ID `gsheet-leads`. Status: Active.
5. **Internal API**: ID `internal-users`. Status: Active.

Frequently Asked Questions

01 **How do I check if my data pipeline ran successfully using the Airbyte MCP for AI Agents?**

Your agent checks the job history directly. You simply ask it about a connection, and it tells you the status (success/fail) of specific runs, saving you from clicking through dashboards.

02 **Can I use the Airbyte MCP for AI Agents to see all my data sources?**

Yes. You can ask the agent to list all your connected data origins (like Postgres or Stripe) instantly. It gives you a clean, comprehensive inventory of everything feeding your data lake.

03 **Does the Airbyte MCP for AI Agents help me find my warehouse endpoints?**

Absolutely. You can list all configured destinations—whether it's Snowflake or BigQuery—so you always know exactly where every piece of data is going.

04 **What if I need to debug a failed sync job with the Airbyte MCP for AI Agents?**

You tell your agent which connection failed, and it retrieves the detailed job history. It often includes the error reason (like a missing credential) so you know exactly what needs fixing.

05 **Is the Airbyte MCP for AI Agents better than just checking the dashboard?**

It's faster and more reliable. Instead of manual clicking, your agent performs automated audits, giving you a summarized report in plain language that highlights exactly what needs attention.

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 `"airbyte": { "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

Airbyte 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 Airbyte. 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	Airbyte MCP
Server ID	`019d754a-987d-72d0-8004-b3bb6a4d7810`
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/airbyte.