MCP SERVER | NO CODE | CLOUD HOSTED

AirOps MCP for AI Agents

Orchestrate complex LLM workflows and manage data pipelines

AirOps handles professional AI workflow orchestration and agent management. It lets your AI client execute complex, multi-step pipelines, interact with specialized agents, and query private knowledge bases—all through simple conversation. Build sophisticated LLM applications without writing boilerplate code.

A+ Quality Score 100/100

ai-workflows | agent-management | llm-ops | automation | model-deployment | memory-management

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

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.

AirOps MCP

10 tools available
Cloud-hosted on Vinkius

Running advanced AI models used to mean managing dozens of separate APIs, checking status endpoints, and manually feeding data between services. Now, you talk to your agent, and it handles the whole process.

This MCP lets you treat complex automation like a natural conversation. You can instruct your agent to execute multi-step workflows, passing custom parameters as if you were talking to a human coworker. Need to reference internal policies? The agent searches managed memory stores—your private knowledge base—and uses that context to answer. If the job is huge and takes minutes, you just ask it to run in the background, and your client tracks the status until it's done.

Whether you need to upload a file for data extraction or simply chat with a niche expert agent, this MCP manages all the connections through your preferred AI client. It’s designed to make building reliable, production-grade LLM applications feel less like engineering and more like talking.

Core Capabilities

01 — Run Structured Workflows

You execute multi-step data pipelines quickly or run them in the background for long tasks.

02 — Interact with Specialized Agents

You chat directly with niche AI agents built for specific business functions, like legal analysis or content summarizing.

03 — Query Internal Knowledge Bases

The agent searches and retrieves information from your private document repository to inform its answers.

04 — Manage Files for Inputs

You upload source files, allowing the AI to use them directly for data extraction or analysis.

05 — Track Long-Running Tasks

The system monitors execution progress and lets you cancel tasks that stall or take too long.

One Click on Vinkius — From Prompt to Execution

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

01 Subscribe to the AirOps MCP and provide your API key.

02 Use your AI client to initiate a task, such as asking it to run an application or search memory.

03 Your agent executes the operation in the background, providing real-time status updates until the result is ready for you.

The bottom line is that you speak naturally to your agent, and it manages all the complex backend orchestration automatically.

Built For

This MCP targets AI Engineers building production-grade LLM systems. It's for Data Specialists who are tired of manual data retrieval and Product Managers needing to quickly test or adjust agent behaviors without deep coding knowledge.

AI Engineer

They use the MCP to automate complex LLM chains, monitoring performance and coordinating multiple services from a single chat interface.

Data Specialist

They feed the system private knowledge documents and query them via memory stores to build data-informed AI applications.

Product Manager

They test specialized agent configurations on the fly, adjusting prompts and workflows to see how outputs change before committing code.

What Changes When You Connect

01 The AirOps MCP lets you execute multi-step processes like an expert. You don't need to write separate code blocks; your agent handles the flow.

02 Need context? Instead of manually searching databases, simply ask the agent to search memory store and it uses the retrieved data immediately.

03 Running a job that takes 20 minutes? Use execute_workflow_async. You start the task now and come back later to check its status with get_execution_status.

04 You can manage multiple applications without switching tabs. Start by listing all apps, then dive into details for any specific tool you need.

05 If a job goes wrong or takes too long, you don't get stuck. You use cancel_execution to stop the process and debug what went wrong.

Real-World Applications

Generating Compliance Reports from Internal Docs

A data specialist asks their agent for a report on 'Q3 privacy violations.' The agent first searches memory store using search_memory_store, pulls relevant policy documents, and then executes an application to summarize the findings into a structured PDF.

Processing Batch Customer Feedback

A product manager uploads hundreds of customer transcripts via upload_file. The agent runs a workflow synchronously using execute_workflow_sync, extracting sentiment and key feature requests from every file in one go.

Debugging a Broken Marketing Chain

An AI engineer notices a scheduled task failing. Instead of guessing, they use get_execution_status to check the failure point, then cancel_execution if it's stuck looping on an error.

Creating a Niche Content Summarizer

A marketer needs quick summaries for specific topics. They chat with agent using chat_with_agent and guide the conversation to produce several structured content outlines in minutes.

Patterns to Avoid

Over-relying on single-pass prompts

✗ AVOID

Asking the AI, 'Summarize this document and then tell me three key action items.' The AI often mixes up summary points with actionable advice.

✓ INSTEAD

Break it into steps. First, use upload_file to give context. Then, ask the agent to run a specific workflow using execute_workflow_sync that is trained just on summarization. Follow up in a second turn asking for action items.

Manually managing data inputs

✗ AVOID

Copying and pasting document sections into the chat window because you can't figure out how to feed it multiple sources.

✓ INSTEAD

Use upload_file to centralize all your source materials. Then, use search_memory_store so the agent knows exactly where to pull context from.

Ignoring job status

✗ AVOID

Starting a massive data extraction task and walking away without checking if it actually finished or just timed out.

✓ INSTEAD

Always use execute_workflow_async for big jobs. Then, regularly check the progress using get_execution_status until you confirm completion.

The Right Fit

Use this MCP if your AI application involves state management; that is, when the process requires multiple steps—like reading a document (upload_file), then querying private knowledge (search_memory_store), and finally generating an output (execute_workflow_sync). Don't use it if you only need simple single-query retrieval or basic text generation. For pure chat interactions without structured data needs, your agent client might suffice. But if the task requires reliability, background job handling (execute_workflow_async), or accessing a defined knowledge base, this is what you need.

AirOps MCP for AI Agents: Managing LLM Workflow Data Pipelines

Think about the current process. To build an advanced data analysis tool today, you're stuck jumping between tools: you upload a file to one service, pull the resulting JSON into another, run it through a specialized agent, and then manually search your internal wiki for compliance rules. It’s painful copy-pasting hell.

With this MCP, the process changes entirely. You simply ask your agent to 'Analyze this report against our Q3 guidelines.' The system handles the file upload, runs the data extraction workflow in sequence, searches the memory store for the right policies, and delivers one integrated answer. It’s automation you talk to.

AirOps MCP for AI Agents: Controlling Agent Memory and Knowledge Retrieval

The biggest headache is knowledge drift. You need the LLM to be smart, but its intelligence relies on context you have to constantly feed it or risk hallucination.

This tool fixes that by letting you manage memory stores directly. Instead of relying only on what was in its training data, your agent pulls specific, indexed facts from documents you control via search_memory_store and add_memory_document. Your AI gets context-aware.

AirOps: 10 Tools for LLM Workflow Management

Use these tools to manage workflow execution, retrieve memory from knowledge bases, and interact with specialized agents.

#	TOOL	DESCRIPTION
01	`upload_file`	Allows you to upload source files that serve as inputs for AI processing tasks.
02	`list_apps`	Retrieves a list of all available AI applications configured in the workspace.
03	`get_app_details`	Fetches specific metadata and details about a particular AI application.
04	`execute_workflow_sync`	Runs an entire multi-step workflow immediately, best for quick tasks that need instant results.
05	`execute_workflow_async`	Starts a long-running workflow in the background so you can continue working while it processes.
06	`get_execution_status`	Checks if an initiated task is complete, failed, or still running.
07	`cancel_execution`	Stops a workflow that has started but needs to be halted before completion.
08	`chat_with_agent`	Enables conversational interaction with a specialized AI agent for specific tasks.
09	`search_memory_store`	Searches the knowledge base, finding relevant document snippets based on your query.
10	`add_memory_document`	Adds new documents or information to the memory store to expand the AI's domain knowledge.

See It in Action

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

U I need to understand our new vendor agreement terms quickly.

Vendor Agreement Summary

- **Effective Date:** 2024-11-01
- **Key Clause:** Indemnity is mutual and capped at $5M.
- **Action Item:** Must be approved by legal before signing.

I found the relevant sections in the knowledge base. Would you like me to draft a summary email for review?

U Run the 'Client Data Extractor' app on this file and give me the JSON.

Workflow Execution Complete

✅ **Status:** Success (Sync)

```json
{
"client_id": "CUST-9012",
"name": "Jane Doe",
"service_tier": "Premium"
}
```

Do you want to save this result?

U Check the status of the large report I started an hour ago.

Job Status Check

- **Task:** Q4 Market Analysis Report (Async)
- **Status:** Processing
- **Progress:** 65% Complete
- **ETA:** Approximately 12 minutes remaining.

You can use the chat interface to cancel this job if you need to switch focus.

Frequently Asked Questions

01 **How does AirOps help me build complex AI workflows without coding?**

AirOps lets you orchestrate multi-step processes using natural conversation. You don't write API calls; you simply tell your agent to perform a sequence of actions, like 'fetch data' then 'summarize it,' and the MCP manages the flow.

02 **Can AirOps connect my private company documents so my AI can answer questions?**

Yes. You use the memory management tools to upload your internal documents. The agent searches this dedicated knowledge base, ensuring the answers it provides are based on *your* specific policies and data.

03 **What if I need to run a job that takes hours? Does AirOps handle that?**

The MCP supports asynchronous workflow execution. You start the long task, and your agent monitors it in the background until it's done, giving you status updates without freezing your chat session.

04 **Is AirOps only for data extraction? Can I use it for general tasks?**

Not at all. While it excels at structured data and workflows, you can also use the agent to interact with specialized conversational agents (chat_with_agent) for niche Q&A or content generation.

05 **How do I make sure my AI uses up-to-date information?**

You enrich your AI's context by managing memory stores. You can add new documents or update policies, and the agent will use that most current knowledge when responding to queries.

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 `"airops": { "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

AirOps 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 AirOps. 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	AirOps MCP
Server ID	`019d754b-13a8-73b7-8011-cd73705ecde2`
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/airops.