MCP SERVER | NO CODE | CLOUD HOSTED

DoiT MCP for AI Agents

Monitor Multi-Cloud Spending and Asset Usage

DoiT connects your AI agent directly to cloud cost data across AWS, GCP, and Azure. It lets you monitor real-time spending anomalies, check budget limits, and audit every connected asset without opening a dashboard. You can get instant answers on total spending or pinpoint the exact service causing a spike.

A+ Quality Score 100/100

cloud-cost-management | finops | multi-cloud | cost-optimization | budget-tracking | cloud-governance

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

SCAN TO VIEW

The connectivity layer between AI and the world's software.

Vinkius sits between AI and every application. All communication passes through Vinkius Cloud via the Model Context Protocol (MCP) — with governance, observability, and security at every layer.

Your AI Connections Run Through Vinkius Cloud

The world's largest managed MCP catalog

Vinkius is the connectivity layer where AI connects 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.

DoiT MCP

10 tools available
Cloud-hosted on Vinkius

Managing multi-cloud costs used to mean jumping between three different vendor dashboards just to figure out where your money was going. This MCP changes that. It lets you manage all your cloud assets and budgets—whether they're on AWS, Google Cloud, or Azure—using natural conversation with your AI agent. You can ask it things like, 'What accounts are over budget?' or 'Show me the biggest cost spikes from last week.' The system pulls the data, identifies anomalies instantly, and gives you a clear summary, letting you focus on optimization instead of spreadsheet clicking. When you connect this MCP through Vinkius, your AI client gets access to one single point of truth for all cloud spending, making FinOps an actual conversation.

Core Capabilities

01 — Check total expenditure across platforms

Get a high-level summary showing how much you spent in total, regardless of which cloud platform (AWS, GCP, or Azure) it came from.

02 — Identify unexpected cost spikes

List all detected anomalies and spending increases that went over budget, so you know exactly what needs immediate attention.

03 — Audit specific cloud assets

Retrieve detailed configuration and cost data for a single asset or search across your entire inventory of connected accounts.

04 — Review budget health status

List all configured spending budgets, showing which ones are set up and which ones have already been exceeded.

05 — Find critical cost issues fast

Immediately identify high-severity cost spikes that require urgent investigation or action from the team.

One Click on Vinkius — From Prompt to Execution

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

01 Authorize your AI client using your DoiT API Key. You find this key in your organization settings.

02 The MCP connects to all connected cloud accounts (AWS, GCP, Azure) and pulls the latest cost metadata.

03 You prompt your agent with a question—like 'Are we over budget?'—and it processes the data to give you an immediate answer.

The bottom line is: instead of running reports manually, you just ask the AI what's wrong with your cloud spending.

Built For

This is for anyone who feels stress when looking at a cloud bill. FinOps Engineers need it to report accurate data quickly. Cloud Architects rely on it to audit multi-cloud assets before deployment, and Operations Leads use it to catch cost overruns the moment they happen.

FinOps Engineer

Running month-end reports or preparing budget forecasts by querying total cloud spending summaries.

Cloud Architect

Auditing asset configurations and platform associations across AWS, GCP, and Azure to ensure compliance before a new service goes live.

Operations Lead

Monitoring for critical cost anomalies or unexpected spending spikes across the entire organization in real-time.

What Changes When You Connect

01 Pinpoint cost overruns instantly: Instead of guessing, use the 'list_critical_cost_spikes' capability to immediately identify high-severity spending issues.

02 Track budget adherence easily: Run queries using 'list_exceeded_cost_budgets' and get a clean list of every budget that has been breached.

03 Get full visibility across vendors: The 'get_billing_cost_summary' tool aggregates total spending from AWS, GCP, and Azure into one view.

04 Audit specific services quickly: Need details on just one service? Use 'get_asset_details' to pull the exact configuration and cost data for that asset.

05 Maintain an accurate inventory: The combination of 'list_cloud_assets' and 'search_cloud_assets' gives you a complete, searchable map of every resource you own.

Real-World Applications

Discovering the source of unexpected charges

The Ops Lead notices a massive spike in costs. They ask their agent to 'list_cost_anomalies.' The agent returns an alert detailing the service and account, allowing them to fix it before the next billing cycle.

Auditing pre-deployment infrastructure

A Cloud Architect needs to check if a new staging environment complies with budget rules. They use 'list_cost_budgets' and then 'get_asset_details' on the specific resource group to confirm compliance.

Generating an executive spending report

The FinOps Engineer needs a quick overview for leadership. Prompting the agent with 'total cloud cost summary' instantly pulls data, providing a clear breakdown of total spend across all platforms.

Identifying under-governed accounts

A manager wants to know which departments are running unchecked resources. They use 'list_connected_cloud_accounts,' then cross-reference that list with assets they need to audit.

Patterns to Avoid

Reading the raw billing report

✗ AVOID

Opening a massive CSV file and trying to manually compare spending across AWS, GCP, and Azure to find the highest cost service.

✓ INSTEAD

Instead, ask your agent to 'get_billing_cost_summary.' It pulls all three platforms' data into one conversational summary, saving you hours of comparison.

Manually tracking asset status

✗ AVOID

Trying to keep a spreadsheet updated with every new resource ID or configuration change across multiple cloud consoles.

✓ INSTEAD

Run 'list_cloud_assets.' This instantly updates your understanding of the entire inventory, giving you a real-time view without manual data entry.

Ignoring spending spikes

✗ AVOID

Seeing an unexplained jump on the bill and assuming it's just normal usage, delaying investigation.

✓ INSTEAD

Prompt for 'list_critical_cost_spikes.' The agent immediately surfaces these urgent issues, telling you exactly where to look first.

The Right Fit

Use this MCP if your primary pain point is cloud cost visibility across multiple vendors. If you need to track spending budgets, audit resource configurations, or find anomalous charges, this is the right tool. Don't use it just because you want a report; use it when you need actionable intelligence. For instance, if you only care about listing assets but don't need the cost context, other inventory tools might suffice. But since your goal is optimization and governance, leveraging capabilities like 'list_exceeded_cost_budgets' makes this MCP essential.

DoiT MCP for Cloud Cost Oversight: Automating FinOps Audits

Right now, calculating your total cloud spend is a nightmare. You have to open the AWS console, switch over to GCP's billing page, then log into Azure just to get a basic idea of where all the money went last month. It’s click-heavy, time-consuming, and you always end up with fragmented data.

With this MCP, your agent handles the complexity. You ask it for a total cost summary, and it pulls combined expenditure from every connected platform into one clean answer. It doesn't just give numbers; it tells you where those funds came from.

DoiT MCP for Budget Monitoring: Proactive Cost Control

The manual process of budget tracking involves setting up alerts in multiple places, then having to constantly check each console dashboard. If you miss one alert, the cost keeps climbing until the next painful meeting.

This MCP lets you ask it to 'list_exceeded_cost_budgets.' It finds those overruns and tells you exactly which budget is breached, giving your team the foresight they need to act before spending goes out of control.

DoiT MCP: 10 Tools for Cloud Cost Analysis

Use these tools to analyze spending patterns, list assets, find cost anomalies, and audit budgets across all connected cloud platforms.

#	TOOL	DESCRIPTION
01	`get_doit_account_metadata`	Retrieves basic information and metadata about your current DoiT organization setup.
02	`get_asset_details`	Gathers deep configuration details and associated costs for a specific cloud asset you identify.
03	`get_billing_cost_summary`	Pulls together a summary of your total spending across all connected cloud platforms.
04	`list_cost_anomalies`	Lists every detected cost anomaly and unexpected spike in spending activity.
05	`list_cloud_assets`	Provides a comprehensive list of all managed cloud assets across AWS, GCP, and Azure.
06	`list_cost_budgets`	Displays every spending budget that has been configured within the system.
07	`list_connected_cloud_accounts`	Lists all AWS, GCP, or Azure accounts that are currently connected and tracked by DoiT.
08	`list_critical_cost_spikes`	Pinpoints high-severity cost anomalies that need immediate attention from the engineering team.
09	`list_exceeded_cost_budgets`	Identifies and lists all budgets whose spending limits have been surpassed.
10	`search_cloud_assets`	Searches your asset inventory using a simple name keyword to quickly locate the right resource.

See It in Action

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

U What's our total cloud cost summary for this quarter?

Total Spending Summary

The current period's total cloud spending is **$45,200**. This figure includes:
- AWS: $22,500 (Hosting & Compute)
- GCP: $18,700 (Data Processing)
- Azure: $4,000 (Networking)

Do you need a breakdown of the top three cost drivers on AWS?

U Are there any critical cost anomalies right now?

⚠️ **2 Critical Anomalies Detected:**

1. **BigQuery (GCP):** A 150% spike detected in the last 4 hours. *Immediate review required.*
2. **S3 Transfer (AWS):** Unexpected data transfer increase, up 75%. Needs investigation into source bucket.

I recommend checking the detailed telemetry for the BigQuery spike first.

U List all budgets that have exceeded 100% consumption.

🚨 **Budgets Exceeded:**

- **'Dev-Sandbox':** Currently at **115%** of its $2,000 limit. (Needs immediate attention.)
- **'Marketing-Analytics':** At **105%** of $5,000 limit.

Would you like me to check the assets associated with the 'Dev-Sandbox' account?

Frequently Asked Questions

01 **How does DoiT MCP help me track multi-cloud spending?**

DoiT aggregates your costs from AWS, GCP, and Azure into a single view. You get one central report instead of having to switch between three different vendor consoles.

02 **Can the DoiT MCP find unexpected cost spikes?**

Yes, it monitors for anomalies in real time. If spending suddenly jumps—say, a 150% spike on one service—the MCP flags it immediately so you can investigate.

03 **Is DoiT MCP only for big companies?**

No. It works by connecting to your existing accounts and tracking spending against budgets, regardless of size. You set the rules, and the MCP monitors them.

04 **What kind of assets can I audit with DoiT MCP?**

You can list and get details on virtually any cloud asset—from compute instances to storage buckets—across all connected platforms.

05 **Does the DoiT MCP help me manage budgets automatically?**

It allows you to list all configured spending limits and shows which ones have been exceeded, giving you clear visibility into your financial guardrails.

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 `"doit": { "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

DoiT 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 DoiT. 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	DoiT MCP
Server ID	`019d7588-25a7-726c-814d-596d50ee8659`
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/doit.