MCP SERVER | NO CODE | CLOUD HOSTED

Datadog MCP for AI Agents

Query performance metrics and search logs across your cloud infrastructure.

Datadog MCP. It lets you query performance metrics, search application logs, and manage active monitors directly from your AI agent. You can check system health, find specific error traces, and audit your alerting rules without ever leaving your chat or IDE. It turns your observability data into a conversational tool for faster troubleshooting and more efficient incident response across your entire cloud infrastructure.

A+ Quality Score 100/100

infrastructure-monitoring | log-analysis | performance-metrics | cloud-observability | alerting | real-time-monitoring

MAINTAINED BY VINKIUS CLOUD FOR AI AGENTS · JULY 2026
vinkius.com/mcp/datadog

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.

Datadog MCP

11 tools available
Cloud-hosted on Vinkius

You can connect your Datadog account to your AI agent and take control of your infrastructure monitoring and log management through natural conversation. Instead of jumping between tabs to find out why a service is lagging, you can just ask your agent to pull the latest metrics or search for specific error codes in your logs. This MCP gives your agent the eyes it needs to see what's happening across your cloud environment. You can check if your monitors are firing, find out which hosts are running out of disk space, or see how your SLOs are performing. Because Vinkius makes it easy to connect these types of tools, you can set up your workspace and start asking questions about your stack immediately. It's about moving from manually digging through dashboards to getting answers in real-time while you're already working in your preferred environment. You can audit your monitor configurations to ensure they're hitting the right thresholds or quickly check for scheduled maintenance to avoid chasing ghost issues. Whether you're troubleshooting a production outage or just doing a routine health check, you can get specific data points without the friction of navigating a complex UI. It turns your observability data into a conversational resource that's always ready when you need it.

Core Capabilities

01 — Query live performance metrics

Pull time-series data for infrastructure or application metrics instantly.

02 — Search and filter application logs

Find specific error traces or status codes using natural language queries.

03 — Check monitor statuses and health

See which monitors are in alert, warning, or ok states across your stack.

04 — View dashboard configurations

Retrieve layout structures and widget details for your existing dashboards.

05 — Audit service level objectives

Check your SLO compliance and target percentages for key services.

06 — Identify scheduled maintenance

Look up planned downtime windows to avoid chasing false alerts.

07 — Mute active alerts

Silence specific monitors for set periods during maintenance or known issues.

One Click on Vinkius — From Prompt to Execution

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

01 Add the Datadog integration to your AI client.

02 Provide your Datadog API Key, APP Key, and Site.

03 Ask your agent to query metrics, search logs, or list monitors.

The bottom line is you get a conversational interface for your entire Datadog observability stack.

Built For

The DevOps engineer who's tired of clicking through dozens of dashboards at 2am to find a single spike. The developer who needs to see logs without switching windows.

DevOps Engineer

Audits alerts and checks system health during active incidents without leaving the chat.

Software Developer

Searches application logs for 500 errors while debugging features in the IDE.

SRE

Monitors active alerts and tracks performance trends during high-pressure response windows.

System Admin

Verifies monitor configurations and host attributes through natural language queries.

What Changes When You Connect

01 Stop tab-switching by checking Datadog metrics directly in your chat or IDE.

02 Find errors faster by using search_logs to filter application logs for specific status codes.

03 Manage your alerting strategy with mute_monitor to silence known issues during maintenance.

04 Track your uptime and reliability using list_slos to see if you're meeting your targets.

05 Get instant visibility into system health by using list_monitors to see active alerts.

06 Audit your infrastructure quickly with list_hosts to see active cloud provider attributes.

Real-World Applications

Finding the root cause of a spike

A developer notices a latency spike. They ask the agent to query_metrics for the web-server and then use search_logs to find 500 errors.

Checking maintenance status

An SRE sees an alert but wants to know if it's planned. They ask the agent to list_downtimes to see if maintenance is active.

Managing alert fatigue

A DevOps engineer wants to silence a flapping monitor. They tell the agent to use mute_monitor on the specific monitor ID.

Auditing SLO compliance

A manager wants to know if the API is healthy. They ask the agent to list_slos to see the current compliance status.

Patterns to Avoid

Trying to guess monitor IDs

✗ AVOID

Trying to remember every specific ID for every monitor in your stack.

✓ INSTEAD

Use list_monitors first to see what's active, then use get_monitor for the specific details of the one you need.

Querying logs without filters

✗ AVOID

Asking the agent to show all logs, which can be overwhelming and slow.

✓ INSTEAD

Use search_logs with specific query syntax to find exactly what you need, like a specific error code.

Ignoring downtime

✗ AVOID

Treating a planned maintenance window as a critical production outage.

✓ INSTEAD

Use list_downtimes to check for scheduled windows before escalating any alerts.

The Right Fit

Use this if you need to interact with Datadog data using natural language during your normal workflow. It's best for quick checks, log hunting, and alert management. Don't use it if you need to build complex, multi-widget dashboards from scratch; you should still use the Datadog UI for that. Use this when you want your AI agent to have eyes on your infrastructure without making you jump between windows.

Datadog Infrastructure Monitoring and Log Analysis

Most DevOps teams spend their day jumping between dozens of browser tabs. You check a dashboard, then switch to the logs, then go back to your IDE to fix the code. It's a constant cycle of

This MCP changes that by bringing the data to you. You can stay in your chat or editor and ask your agent to pull specific metrics or find error

context switching that breaks your focus and slows down incident response.

traces. You get the answers you need without ever leaving your workspace.

Datadog Alert Management and SLO Tracking

Managing alert fatigue is a constant battle. You have to manually mute monitors, check if a downtime is scheduled, and keep an eye on SLO compliance across multiple services.

You can now automate those checks. Just ask your agent to list_monitors to see what's broken or check list_slos to see your current uptime. It turns a manual audit into a quick conversation.

Datadog 11 Tools for Infrastructure Monitoring

Use these tools to query metrics, search logs, and manage monitors via your AI agent.

#	TOOL	DESCRIPTION
01	list_dashboards	Get a list of all your dashboard IDs and titles. This helps you find the right visualization quickly.
02	query_metrics	Pull time-series data for specific infrastructure or app metrics. Use it to see how CPU or memory usage changes over time.
03	list_downtimes	See planned maintenance windows and their current status. This helps you avoid chasing ghost issues during scheduled work.
04	list_slos	View your Service Level Objectives and their current compliance. Use it to see if your app is meeting its uptime targets.
05	search_logs	Find specific log entries that match your query. It's great for spotting errors or specific status codes across your apps.
06	list_monitors	See all your monitors and their current states like alert or ok. This gives you a quick overview of your system health.
07	get_monitor	Get the full details of a specific monitor, including its thresholds. Use it to understand exactly why a monitor is firing.
08	mute_monitor	Silence a specific monitor for a set period. This is perfect for ignoring known issues during maintenance.
09	list_events	Pull a list of events like alerts and deployment logs. Use it to see what happened in your environment recently.
10	get_dashboard	See the full configuration and widgets for a specific dashboard. This helps you understand how a dashboard is built.
11	list_hosts	Get a list of your infrastructure hosts and their tags. Use it to see which cloud resources are currently active.

See It in Action

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

U Show me the CPU usage for 'web-server' over the last 30 minutes

Querying Datadog metrics...

METRIC	VALUE
Average CPU	45%
Peak CPU	78%

Observation: Performance peaked about 10 minutes ago but has since stabilized. Everything looks normal now.

U Find logs with '500 Internal Server Error' from the last hour

I found 12 logs with '500 Internal Server Error' in the last hour.

- **Primary Service:** auth-service
- **Status:** Critical
- **Trace IDs:** `tr_9921`, `tr_9925`, `tr_9930`

Would you like me to pull the full stack traces for these specific IDs?

U Are there any active monitors in 'Alert' state?

Checking Datadog monitors... I found 2 monitors currently in **Alert** state:

1. **Free Disk Space Low** on `db-node-1`
2. **High Request Latency** on `api-gateway`

All other monitors are reporting 'OK' or 'Warning'.

Frequently Asked Questions

01 **Can I use the Datadog MCP to see my server health?**

Yes, you can check your system health by asking the agent to list monitors or query specific metrics. It gives you a real-time look at what's running and what's failing.

02 **How do I connect my Datadog account to my AI agent?**

You'll need your Datadog API Key, APP Key, and Site. Once you provide those to your AI client, the agent can start querying your infrastructure.

03 **Can it help me find specific errors in my logs?**

Yes, it can search logs using specific queries. You can ask for logs containing specific error codes, status levels, or structured attributes.

04 **Can I silence alerts using this MCP?**

Yes, you can use the mute_monitor tool to silence specific monitors for set periods. This is great for ignoring known issues during maintenance.

05 **Does this MCP work for cloud infrastructure?**

Yes, it works with any host or service you have in Datadog, including cloud provider attributes and agent versions.

06 **Can I see my SLO status?**

Yes, you can ask the agent to list your Service Level Objectives. It will return your target percentages, time windows, and current compliance status.

07 **Can my agent query specific Datadog metrics using DDQL?**

Yes. Use the 'query_metrics' tool. Provide your DDQL query string and the target time range. The agent will fetch the numeric timeseries data directly from Datadog's telemetry datastores.

08 **How do I search for a specific error message across my application logs?**

Use the 'search_logs' tool. Provide a query matching your error string and an ISO time boundary. The agent will retrieve the structural extraction of logs matching those parameters to help you identify failures.

09 **Can I see which monitors are currently in an alert state?**

Absolutely. The 'list_monitors' tool allows you to filter by group state (e.g., 'alert,warn'). The agent pulls the explicitly configured system triggers to show you which services are currently unhealthy.

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 `"datadog": { "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

Datadog 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 Datadog. 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	July 2026
MCP Server	Datadog MCP
Server ID	`019d7581-c015-7220-b99a-6852b938fd83`
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/datadog.