MCP SERVER | NO CODE | CLOUD HOSTED

Datadog MCP for AI Agents

Monitor Infrastructure Health and Query Performance Logs

Datadog connects your AI client to full-stack observability data. You get conversational control over metrics, infrastructure health, and logs in real time. Instead of clicking through complex dashboards, you talk to your system to list active incidents, query specific performance metrics, or search error logs across every service.

A+ Quality Score 100/100

full-stack-monitoring | infrastructure-metrics | log-analysis | incident-management | cloud-monitoring | alerting

MAINTAINED BY VINKIUS CLOUD FOR AI AGENTS · JUNE 2026
vinkius.com/mcp/datadog-extended

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

16 tools available
Cloud-hosted on Vinkius

Monitoring a large application stack shouldn't require deep knowledge of the Datadog UI. This MCP connects any AI client to your entire observability setup, letting you manage infrastructure health using natural conversation. You can ask your agent to find out why latency spiked yesterday or check if a specific host is running low on disk space without opening a single dashboard tab. It lets you run time-series queries with precise Datadog syntax and search across all indexed logs immediately. Need to control noise? You can even list and mute monitors during maintenance windows, keeping your team focused on actual issues. Connect this MCP through Vinkius to gain unified visibility into everything from Service Level Objectives (SLOs) down to individual host metadata.

Core Capabilities

01 — Check system connectivity

Verify the connection status between your AI client and Datadog.

02 — List, search, or inspect monitors

Review all defined alerts to see what's firing or mute noisy ones during planned maintenance periods.

03 — Get dashboard layouts and variables

Retrieve the full structure of any operational dashboard, including widget details and template variables.

04 — Run custom metric queries

Execute specific time-series queries using Datadog syntax to analyze performance data across custom time ranges.

05 — Search detailed log events

Find specific error or warning events by querying logs using standard Datadog query language.

06 — Manage platform and custom events

List existing system events, check out host inventory details, or create new operational tags.

07 — Track active incidents and SLOs

See a list of current high-severity incidents, including who is responding and the timeline. You can also review Service Level Objectives for error budget compliance status.

One Click on Vinkius — From Prompt to Execution

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

01 Subscribe to this MCP and provide your Datadog API Key along with the correct site URL (e.g., `https://api.datadoghq.com`).

02 Your AI client authenticates with Vinkius, allowing it to send structured commands directly to the monitoring platform.

03 You simply ask your agent a question—like 'Why did the API latency spike last night?'—and it runs the necessary metric queries or log searches for you.

The bottom line is that instead of navigating complex UIs, your AI client talks directly to your monitoring data via structured tools.

Built For

This MCP is built for the engineering trenches. It's for DevOps and SRE engineers who are tired of context switching between dashboards, logs, and alert screens during an outage. If your day involves triaging incidents at 2 AM by clicking through multiple tabs, this tool saves you time.

SRE / DevOps Engineer

You use this MCP to query monitors or search error logs instantly without opening the main Datadog dashboard. You can validate SLO compliance and mute noisy alerts on demand.

Platform Team Lead

You leverage it to run complex metric queries against historical data, validating service health and tracking host metadata across your entire fleet.

On-Call Site Reliability Engineer

When an incident hits, you use this MCP to list active incidents, check the severity, and search for correlated error logs—all in a single conversational thread.

What Changes When You Connect

01 Instantly triage alerts. Instead of listing every monitor manually, use the `list_monitors` tool to quickly see which alerts are firing and check their status.

02 Deep dive into performance data. Run complex time-series analysis using `query_metrics` with specific Datadog syntax, getting granular results without writing a query language script.

03 Reduce alert fatigue. Use the `mute_monitor` tool to silence noisy alerts for planned maintenance periods, ensuring your team only gets notified of real issues.

04 Pinpoint root causes fast. The `search_logs` tool lets you search across all indexed log sources using natural queries, correlating errors with specific hosts or services.

05 Full visibility into service commitments. Review Service Level Objectives and check error budget compliance via the SLO tools to ensure your application meets its goals.

Real-World Applications

Investigating a sudden spike in checkout latency

The agent can run `query_metrics` for P95 latency over the last four hours. It then uses `search_logs` to correlate the exact time window of high latency with error events found in the payment service logs, identifying 'TimeoutException' as the root cause.

Preparing for a major system update

Before deploying new code, the agent can use `list_hosts` to generate a current inventory list of all reporting hosts and their tags. It can then run `get_monitor` on key services to ensure alerts are configured correctly before the change.

Handling an active outage incident

A user asks, 'What's going wrong right now?' The agent uses `list_incidents` for a summary, then checks `get_incident` details to see who is responding and the current status of the service.

Auditing system reliability targets

The team needs an overview. They ask the agent to check all SLOs via `list_slos`. The agent identifies which objectives are nearing their error budget limit, flagging services that require immediate attention.

Patterns to Avoid

Copy-pasting API endpoints

✗ AVOID

Manually navigating to the Datadog dashboard and copying complex query syntax from one tab to another takes minutes, risking human error with dates or tags.

✓ INSTEAD

Just ask your agent. It runs `query_metrics` directly against the data using natural language instructions, generating the exact time-series graph you need instantly.

Missing context on hosts

✗ AVOID

Seeing an alert about 'Host X' but having no idea if it's a development sandbox or production hardware. The manual process requires checking separate inventory sheets.

✓ INSTEAD

Use the `list_hosts` tool to get all reporting host metadata and tags in one place, immediately telling you what environment that asset belongs to.

Over-relying on dashboards

✗ AVOID

Dashboards are great summaries, but they often hide the raw details. When a spike happens, you're left with 'Unknown Cause' and have to manually search logs.

✓ INSTEAD

Ask your agent to `search_logs` for the specific metric that spiked in the same timeframe. This gives you the raw error messages needed for true root cause analysis.

The Right Fit

Use this MCP if your team needs to treat observability data like a conversation: querying metrics, logs, and alerts with natural language commands. It's perfect for SRE teams who need to correlate disparate pieces of information quickly during an outage. Don't use it if you only need simple viewing—for instance, listing basic dashboard titles is fine through `list_dashboards`. However, if your workflow requires complex UI interaction like drag-and-drop widget adjustments or building highly customized visualization templates, this MCP won't help because its focus is on data retrieval and action, not visual design. It’s about asking 'Why?' and getting a specific answer.

Datadog MCP: Simplifying Infrastructure Monitoring Tasks

Today, checking system health means clicking through half a dozen tabs in the dashboard—one for CPU, one for latency, another for logs. You copy-paste tags from the host list into your query builder, then you run the search and wait for results. It’s tedious, it's slow, and when an incident is happening, every second counts.

With this MCP, you just tell your agent what to look for—for example, 'Show me all hosts where disk space dropped below 20%.' The agent executes the necessary checks using `list_hosts` or `get_monitor`, pulls the data, and presents a clean, actionable summary. You get immediate answers without opening a single GUI panel.

Datadog MCP: Understanding Host Inventory and Alerts

Before, checking the health of your fleet meant running inventory reports or manually listing all monitors. If a host went offline, finding its status required navigating through multiple dashboards and tags.

With this MCP, you can simply ask for the current `list_hosts` metadata or use `get_monitor` to check specific alerts by name. It gives you immediate, structured data on your entire physical and virtual environment.

Datadog: 16 Tools for Cloud Monitoring & Log Analysis

These tools give your agent the power to list hosts, query metrics, search logs, manage alerts, and check SLOs across your entire Datadog environment.

#	TOOL	DESCRIPTION
01	`check_datadog_status`	Verifies that your AI client can successfully connect to Datadog.
02	`create_event`	Allows you to programmatically create a new system event with specific tags and priority levels.
03	`get_dashboard`	Fetches the complete layout, widgets, and template variables for a specified dashboard.
04	`get_incident`	Retrieves all details about an active incident, including responders and timeline information.
05	`get_monitor`	Gets the full configuration and status of a single alert monitor.
06	`list_dashboards`	Shows all available dashboards within your account.
07	`list_events`	Retrieves a list of recent platform events and custom system activity.
08	`list_hosts`	Lists all reporting hosts, providing metadata, tags, and agent version details for inventory checks.
09	`list_incidents`	Shows a comprehensive list of currently open incidents with their severity and status.
10	`list_metrics`	Lists the available metric types that can be queried.
11	`list_monitors`	Retrieves a list of all defined alert monitors for review.
12	`list_slos`	Shows a summary of Service Level Objectives, including their targets and current compliance status.
13	`mute_monitor`	Temporarily silences an alert monitor to prevent notification noise during maintenance periods.
14	`query_metrics`	Executes specific time-series queries using Datadog syntax on metric data.
15	`search_logs`	Searches through indexed logs across all sources to find specific error or warning messages.

#	TOOL	DESCRIPTION
16	search_monitors	Allows you to search for monitors using keywords or filters.

See It in Action

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

U Show me all monitors that are currently alerting and mute the noisiest one.

Alert Summary

Found 4 monitors in Alert state:
1. 'High CPU - web-prod-01' (ID: 58291, triggered 47 times today).
2. 'Disk Space Low - db-replica-03' (ID: 58344, triggered 12 times).
...

I muted the noisiest one, Monitor 58291. No further CPU alerts will fire until I unmuted it.

U Search for error logs in production from the last hour.

Log Search Results

Query executed: `status:error env:production` ...

Found **234** error events in the last hour. Top sources:
• api-gateway (89 errors)
• payment-service (67 errors)
The most common issue is 'TimeoutException: upstream connection timed out' (124 occurrences).
Do you want to check related incidents?

U List all SLOs and tell me which ones are at risk of breaching their error budget.

SLO Health Check

Reviewed 8 Service Level Objectives. Two are flagged as 'At Risk':
1. **API Availability (99.9%):** Current status: 99.82%. Budget remaining: 12% (7-day window). Expect exhaustion in 2.3 days.
2. **Checkout Latency P95 < 800ms:** Status: 99.1%. Budget remaining: 28%.
The other six are healthy (>60% budget).
Which monitor should I inspect first?

Frequently Asked Questions

01 **How does the Datadog MCP help me query performance metrics?**

It lets you run time-series queries using specific syntax, so you don't have to manually build complex metric queries. You just ask for the data point—like 'P95 latency over 4 hours'—and get the graph.

02 **Can I use this Datadog MCP to manage my alerts and monitors?**

Yes, you can list all defined monitors and even mute them. This is useful for reducing alert noise when your team knows maintenance or testing is happening across the infrastructure.

03 **What if I need to check logs from a specific host?**

The MCP lets you access the full `list_hosts` inventory details, giving you metadata and tags. You can then use this context when searching for error logs via `search_logs`.

04 **Does connecting Datadog MCP improve my incident response time?**

Yes, because it aggregates all critical information—incidents, SLOs, and logs—into a single conversational flow. You spend less time jumping between tabs and more time fixing the problem.

05 **Is this MCP only for viewing data or can I perform actions?**

It does both. You can read detailed reports on SLOs, but you can also take action, like muting a monitor or creating a new system event directly through your AI agent.

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-extended": { "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	June 2026
MCP Server	Datadog MCP
Server ID	`019dd0dc-ff4e-7209-abf7-b03ef00e7665`
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-extended.