MCP SERVER | NO CODE | CLOUD HOSTED

Datadog MCP for AI Agents

Monitor infrastructure, APM, and logs in natural conversation

Datadog provides full observability over your entire infrastructure, applications, and logs through natural conversation. Your AI client can query raw metrics, search structured error logs, track incidents, and audit service health without you ever having to open a dashboard.

A+ Quality Score 98.33/100

apm | infrastructure-monitoring | incident-response | metrics-querying | alerting | cloud-ops

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

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

Managing complex systems means jumping between dashboards, log viewers, and metric graphs—it's exhausting. This MCP connects your existing Datadog account directly to your AI agent, giving it the power to act as a dedicated Site Reliability Engineer (SRE). Instead of clicking through tabs, you just talk to your client. You can ask about resource bottlenecks across specific hosts or check if a recent deployment broke an endpoint. The tool lets you search logs using complex filters, audit service level objectives (SLOs), and even list all active alerts so you know exactly what needs attention. If you're already managing observability tools in the Vinkius catalog, adding this MCP means consolidating your entire operational knowledge base into one conversation with your AI agent.

Core Capabilities

01 — Audit Infrastructure Inventory

See a list of every host monitored by Datadog, along with its current CPU, memory usage, and custom tags.

02 — Query Metric Trends Over Time

Analyze raw time-series data for any metric type—like system CPU or custom business metrics—using specific query syntax to understand performance trends.

03 — Search Detailed Log Events

Filter through structured and unstructured log entries using advanced queries, narrowing results by service, host, or status code.

04 — Manage Alerting Monitors

Create new alerts or modify existing ones (like changing a threshold or setting the notification message) to ensure your systems are properly monitored.

05 — Review Service Health Objectives

List and audit all defined SLOs, letting you check compliance rates for critical services over specific time periods.

One Click on Vinkius — From Prompt to Execution

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

01 First, subscribe to this MCP on Vinkius and provide your Datadog API Key and Application Key.

02 Your AI client authenticates with the keys, establishing a secure connection to your live monitoring data.

03 You then ask conversational questions—like 'What was the average CPU usage for the staging environment last week?'—and receive real-time answers from the platform.

The bottom line is that it turns complex dashboard navigation into simple, actionable conversations with your AI client.

Built For

This MCP is for operations engineers and developers who are tired of context-switching. If you spend half your day clicking between the metrics dashboard, the log viewer, and the alerting console, this tool saves time. It's built for people whose job requires deep, real-time knowledge of system health.

Site Reliability Engineer (SRE)

Triaging active incidents by instantly listing monitors or searching specific error logs without opening the Datadog dashboard.

DevOps Engineer

Auditing overall system health by running metric queries on resource usage and checking SLO compliance across multiple services.

Software Developer

Debugging code issues by querying specific metrics or inspecting log events directly from their IDE, keeping focus on the code itself.

What Changes When You Connect

01 Instantly triage alerts: Use the `list_monitors` tool to see all active monitors without opening the dashboard. You'll know exactly what needs attention right now.

02 Deep log analysis on demand: Instead of manually clicking through Log Explorer filters, use `search_logs` to find specific error patterns across services in seconds.

03 Analyze historical performance trends: The `query_metrics` tool lets you pull raw metric timeseries data for deep dives, identifying the root cause before it becomes a major incident.

04 Audit service reliability easily: Check compliance by calling `list_slos` or review your automated coverage using `list_synthetics_tests`. No more guessing on SLA adherence.

05 Control alert lifecycle: Use `mute_monitor` during planned maintenance windows, and then call `unmute_monitor` when the work is done. It keeps alerts from becoming noise.

Real-World Applications

Finding a P99 Latency Spike

The agent notices an alert spike on API latency. You ask, 'What was our average response time for the payment service last night?' The agent runs `query_metrics`, providing a graph showing the exact minute and magnitude of the performance dip.

Investigating a Service Outage

A user reports an outage. You ask, 'Search for errors related to payment failure.' The agent uses `search_logs` and returns 20 matching entries, pointing immediately to the failing host and providing the stack trace.

Auditing Alerting Coverage

Before a major release, you ask, 'List all service monitors.' The agent uses `list_monitors`, allowing you to quickly spot any critical services that lack an alert definition. You can then use `create_monitor` to fix it.

Onboarding New Team Members

The Engineering Manager asks, 'Who owns the inventory monitoring?' The agent runs `list_teams`, showing team membership and ownership for both monitors and SLOs, streamlining knowledge transfer.

Patterns to Avoid

Manual Dashboard Navigation

✗ AVOID

Trying to check error rates by opening the Datadog dashboard, navigating to the Logs tab, selecting a host filter, and then applying time ranges manually.

✓ INSTEAD

Instead, ask your agent: 'Search logs for status code 503 from the web service in the last two hours.' This uses `search_logs` directly via conversation.

Checking Status One by One

✗ AVOID

Remembering to check if every single monitor is running, or opening the API documentation just to find the correct query syntax for a metric.

✓ INSTEAD

Simply ask: 'List all active monitors.' The agent uses `list_monitors` to give you an immediate overview of your entire alerting footprint.

Guessing Metric Names

✗ AVOID

Trying to figure out the correct metric name for CPU usage or memory utilization without knowing Datadog's specific query syntax.

✓ INSTEAD

Ask the agent: 'What was the average resource consumption on host web01 yesterday?' The tool uses `query_metrics` and handles the complex syntax behind the scenes.

The Right Fit

Use this MCP if your primary pain point is context switching. If you find yourself opening Datadog, then your IDE, then a separate documentation site just to answer one question, you need this. This tool excels when you must correlate data across logs, metrics, and alerts in a single conversation flow. Don't use it if you only need to view static dashboards; simply viewing widgets is faster natively. If you only want to manage users or teams without checking system health, another directory MCP might be better suited. But for deep operational observability, this is the tool.

Datadog and Observability: Managing Infrastructure Alerts with Datadog

Currently, diagnosing an issue requires a painful process of jumping between systems. You check the dashboard for alerts, then copy relevant IDs to your log viewer, filter by time manually, and finally run a metric query to see the resource bottleneck. This constant context switching is where hours disappear.

With this MCP, you talk to your agent instead. If an alert fires, just ask it about the incident; it automatically checks monitoring status and pulls relevant error logs for you. You get actionable intelligence in one chat window.

Datadog and SLOs: Auditing Service Level Objectives with Datadog

Manually auditing service compliance means pulling reports on availability percentages for different time windows (30 days, 90 days). You have to cross-reference these numbers across multiple spreadsheets to see if you're hitting your goals.

Now, just ask: 'What is our SLO compliance rate?' The MCP uses `list_slos` and summarizes the data instantly. It doesn't just report a number; it flags where teams need to take ownership.

Datadog: 16 Tools for Observability, Metrics, and Incident Response

These tools let your AI client run specific operations like querying metrics time series, listing hosts, or searching detailed logs when you need precision.

#	TOOL	DESCRIPTION
01	create_monitor	Creates a new alert monitor based on specified criteria like metric thresholds, anomaly detection, or service checks.
02	list_dashboards	Retrieves a list of all available dashboards so you can identify the right view for your investigation.
03	get_dashboard	Fetches specific details about one dashboard using its unique ID.
04	get_monitor	Gets detailed information for a single, specified monitor by its numeric ID.
05	list_hosts	Lists all monitored servers and hosts, providing key metrics summary and filtering options by tags.
06	list_incidents	Retrieves a record of current or resolved incidents, showing severity, responder assignments, and postmortem status.
07	list_monitors	Lists every active monitor, allowing you to audit your overall alerting coverage across different types (metric, log, etc.).
08	mute_monitor	Temporarily silences a specified alert monitor during planned maintenance or known issues.
09	query_metrics	Runs detailed queries against time-series metric data, analyzing trends for specific metrics within defined time ranges.
10	search_logs	Searches through large volumes of log events using advanced filters like service name or status code to find contextual errors.
11	list_slos	Retrieves all defined Service Level Objectives, which track the target availability percentage for your services.
12	list_synthetics_tests	Lists automated synthetic tests to verify that key endpoints and user journeys are actively monitored.
13	list_teams	Shows the organizational structure by listing all teams responsible for specific monitors, SLOs, or dashboards.
14	unmute_monitor	Reactivates a previously muted alert monitor.

#	TOOL	DESCRIPTION
15	update_monitor	Modifies an existing monitor's details, such as changing the query string or updating the notification message.
16	list_users	Provides a directory of user accounts and their access permissions within the platform for security audits.

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 in an alert state.

Active Alerts Found:

- ⚠️ **Web01 CPU Spike** (Metric Alert) - Critical. Triggered 12 minutes ago. Check `query_metrics` for context.
- ❌ **Payment Gateway Down** (Service Check) - Error. Last checked 5 minutes ago. Requires immediate action.
- ⚙️ **Cache Miss Rate High** (Query Alert) - Warning. Running since 4 hours ago. Review recent logs.

U I need to check the error rate for the user profile service in the last hour.

Error Log Summary: User Profile Service (Last Hour)

We found 45 entries. The primary issue is `DatabaseConnectionFailed` originating from host `user-db-02`.

Top Error Codes:

- `DB Connection Failed` : 31 occurrences
- `Timeout` : 14 occurrences

Would you like me to run a metric query on the connection pool utilization for that host?

U What are our Service Level Objectives and who owns them?

Service Level Objective Audit:

SLO NAME	TARGET	PERIOD	CURRENT STATUS
Checkout Availability	99.9%	30d	**Green** (99.8%)
API Latency P95	< 200ms	7d	**Yellow** (210ms)

Ownership: The `Payments Team` is responsible for the Checkout SLO, and the `Platform Core` team owns the API latency target.

Frequently Asked Questions

01 **What's the difference between Datadog API Key and Application Key?**

The **API Key** authenticates your requests to the Datadog platform and is required for all endpoints. The **Application Key** is an additional layer of authorization that controls what actions your integration can perform. Both are generated in Organization Settings > API and Application Keys. Most Datadog API endpoints require both keys.

02 **Can I mute a monitor during a maintenance window?**

Yes! Use the `mute_monitor` action with the monitor ID. You can optionally set an `end` timestamp (ISO 8601) for the mute to automatically expire, or specify a `scope` to mute only certain sub-alerts (e.g. 'env:staging'). Use `unmute_monitor` to re-enable notifications.

03 **What query syntax does the metrics endpoint use?**

Datadog uses a specific query format: `[function]:[metric]{[tags]}`. For example: `avg:system.cpu.user{host:web01}` returns the average CPU user time for host web01. Common functions include `avg`, `sum`, `max`, `min`, `count`. Time windows are specified in the query as `avg(last_5m):...` or passed as `from`/`to` Unix timestamps to the tool.

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-alternative": { "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	`019d842c-f8d9-7005-9e17-9859020b4ded`
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-alternative.