MCP SERVER | NO CODE | CLOUD HOSTED

Datadog Cloud SIEM MCP for AI Agents

Audit cloud activity and security signals across all environments

Datadog Cloud SIEM connects your security module to any AI agent, giving you full control over threat hunting and cloud auditing. Your agent can search critical security signals matching MITRE ATT&CK vectors, update alert statuses, and build new detection rules using raw log data—all through natural conversation.

A+ Quality Score 100/100

cloud-siem | threat-detection | security-signals | vulnerability-scanning | incident-response | mitre-att-ck

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

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 Cloud SIEM MCP

10 tools available
Cloud-hosted on Vinkius

Managing cloud threats used to mean jumping between dashboards, running complex queries in a terminal, and manually tracking down logs from AWS or Kubernetes. This MCP changes that. You connect Datadog Cloud SIEM via Vinkius, giving your AI agent the deep access needed for true security operations. Instead of writing dense query language, you just talk to it. Your agent can hunt through raw log data over specific timeframes, find critical indicators—like an unauthorized S3 bucket access attempt—and even manage the detection rules themselves. You tell it, 'Find me all instances where a user attempts root escalation,' and it executes that logic instantly, providing structured results so you know exactly what's wrong. It’s like having a highly specialized security analyst always ready to take your verbal instructions.

Core Capabilities

01 — Triage Active Security Alerts

Change the status of an alert signal, marking it as archived or re-opening it, and adding official documentation for why you made the change.

02 — Audit Cloud Detection Rules

View the exact logic used by existing security rules, or retrieve raw information about global log exclusion policies to verify what data isn't being seen by your SIEM.

03 — Perform Deep Threat Hunts

Query massive amounts of raw Datadog logs directly, allowing you to look back at specific IP addresses or application traces related to a potential breach.

04 — Create Custom Detection Rules

Write and activate new Cloud SIEM detection rules by specifying the necessary log fields, query bindings, and desired severity levels.

One Click on Vinkius — From Prompt to Execution

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

01 Subscribe to this MCP on Vinkius and provide your required Datadog API Key and APP Key.

02 Authorize your agent's access using your preferred AI client (like Cursor or Claude).

03 Start by asking your agent a security question, such as 'List all detection rules that monitor Kubernetes root escalations,' to begin managing cloud security.

The bottom line is you use natural language conversation to interact with complex, structured security data and operational tools.

Built For

This MCP is built for the people who live in a state of constant alert: Security Analysts and Incident Responders. If your day involves diving into dense logs or managing dozens of active alerts, this tool cuts down hours of repetitive clicking.

Security Analyst

Uses the MCP to search for high-severity security signals and triage existing alerts by confirming if they are false positives.

Incident Responder

Runs detailed raw log context queries against suspicious source IPs or timestamps immediately after detecting an active threat.

Security Engineer

Manages the security posture by listing, retrieving logic for, and deploying new detection rules using natural language prompts instead of a dedicated console.

What Changes When You Connect

01 **Instant Alert Status Updates:** You don't need to manually change alert statuses. Use `triage_signal` to move alerts from 'open' to 'archived' with a simple conversation, logging the reason automatically.

02 **Deep Threat Context:** Stop guessing what happened. The `get_raw_log_context` tool lets you pull 100 raw log messages right after verifying an attacker footprint, providing immediate context for your report.

03 **Proactive Rule Management:** Instead of digging through consoles, use the MCP to list all rules (`list_detection_rules`) or get a specific rule's logic via `get_detection_rule` , letting you know exactly what coverage you have.

04 **Targeted Threat Hunting:** You can run focused queries using `search_raw_logs` over the last 15 minutes to find contextual VPC Flow Logs related to an active breach, far faster than a manual search.

05 **Rule Deployment via Chat:** Need to monitor for a specific type of IAM usage? Use `create_detection_rule` to build and activate new Cloud SIEM rules using raw name/message fields right from your chat interface.

Real-World Applications

A user suspects an administrator account was compromised.

The agent runs a query via `search_raw_logs` targeting the administrator's source IPs. It finds 15 suspicious events in the last hour, showing multiple failed login attempts and unusual access to administrative endpoints. The analyst then uses this raw context to build a new detection rule using `create_detection_rule` to catch similar patterns immediately.

The team needs to confirm if an alert is a false positive.

A signal pops up in the queue, but it's known maintenance activity. The agent runs `triage_signal`, marking the signal as 'archived' and logging the reason: 'scheduled system testing,' ensuring the record is clean for compliance.

Compliance requires auditing log exclusion policies.

The Compliance Officer asks the agent to run `list_security_filters`. The MCP returns a list of all global filters, allowing the officer to verify that low-value logging vectors aren't accidentally being blocked from critical review.

A new AWS service was deployed and needs monitoring.

The Security Engineer asks the agent to run `list_detection_rules`. Seeing no coverage, they use the MCP to retrieve the specific logic for a similar existing rule using `get_detection_rule`, then modify it via `create_detection_rule` to cover the new service's unique event types.

Patterns to Avoid

Deleting critical rules by mistake

✗ AVOID

A user, trying to clean up old data, asks the agent to delete a rule they think is deprecated. The system might execute `delete_detection_rule` on a core policy that was only meant to be disabled.

✓ INSTEAD

Always check first. Instead of deleting, ask the agent to run `list_detection_rules` or use `get_detection_rule` to view the current logic before making any changes. If you need it gone, confirm the rule is user-created and not a system default.

Relying on high-level summaries

✗ AVOID

The agent reports that 'Brute Force Attempt' was detected by searching signals. The analyst accepts this without checking *how* it happened, missing the source IP or payload details.

✓ INSTEAD

Never stop at the signal level. After receiving a critical alert via `search_signals`, immediately ask the agent to run `get_raw_log_context` against that specific signal ID for the full story.

Treating logs as a single stream

✗ AVOID

Trying to search all data using only general keywords. This results in thousands of irrelevant, low-value records, wasting time and compute budget.

✓ INSTEAD

Always narrow your focus. When hunting for threats, use `search_raw_logs` and provide specific parameters like a known malicious source IP or a precise time window (e.g., the last 15 minutes).

The Right Fit

Use this MCP if your primary workflow involves deep, conversational interaction with structured cloud logs, security signals, and detection rule logic. You need to rapidly transition from 'alert found' to 'context understood' to 'rule updated.' Don't use it if

you simply need a general dashboard view of log volume; for that, stick to native logging tools. If your job is purely ticketing or user management, this MCP won't help because its focus is on the highly technical domain of cloud security auditing. You must be comfortable translating complex concepts like MITRE ATT&CK vectors into plain language commands for your agent.

Datadog Cloud SIEM: Streamlining Threat Signal Analysis with the MCP

Right now, tracking a potential breach means juggling several dashboards. You spot an alert in one place, jump to a log viewer for context, and then switch to a rule management console to see if you need to adjust anything. This manual process is slow; it takes hours of copy-pasting fields and switching between tabs just to understand the full scope of the threat.

With this MCP, your agent handles that entire sequence conversationally. You ask it to find all critical signals and then follow up with, 'Now get the raw log context for Signal XYZ.' The system pulls together the alert, the logic, and the deep logs into one coherent answer. It's immediate visibility, without leaving your chat window.

Datadog Cloud SIEM: Governing Cloud Detection Rules via AI Agents

Setting up detection rules is usually a painful, highly technical process. You have to consult documentation to understand Lucene query bindings and manually test if the rule correctly captures an AWS CloudTrail deviation or a Kubernetes escalation. One wrong binding, and the whole thing fails silently.

Now, you just tell your agent what pattern you want to catch—'I need to detect unauthorized IAM usage on this service.' The MCP handles constructing the complex query, listing existing rules via `list_detection_rules`, and deploying the new rule using `create_detection_rule`. It turns a day-long engineering task into a three-minute conversation.

Datadog Cloud SIEM: 10 Tools for Threat Detection & Log Auditing

These tools let your agent search alerts, manage detection rules, and query raw logs across your entire cloud environment.

#	TOOL	DESCRIPTION
01	`create_detection_rule`	Builds and activates a new Cloud SIEM Log Detection Rule by specifying fields, queries, and severity.
02	`security_system_ping`	Tests the API connection to confirm that your agent can communicate with the Datadog Security Module.
03	`delete_detection_rule`	Permanently removes user-created custom detection rules from the system (use with caution).
04	`get_raw_log_context`	Extracts a deep set of raw log entries immediately after verifying an attacker's operational footprint.
05	`get_detection_rule`	Retrieves the precise query logic for any specific detection rule currently running in your SIEM.
06	`list_security_filters`	Lists all global exclusion policies, showing which low-value log vectors are intentionally blocked from evaluation.
07	`list_detection_rules`	Retrieves a list of every configured proactive detection rule monitoring your cloud environment.
08	`search_raw_logs`	Directly queries raw log data over defined time periods for deep threat hunting, useful for finding context during an active breach.
09	`search_signals`	Searches high-level security signals (alerts) using query language to filter by severity or user ID matching MITRE ATT&CK vectors.
10	`triage_signal`	Changes the status of a signal from open to archived, requiring you to provide an official reason for the change.

See It in Action

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

U Can you list all my current detection rules and tell me what they monitor?

Cloud SIEM Detection Rules Overview

RULE NAME	STATUS	TARGET EVENT	MONITORS FOR
AWS_Trail_Deviations	Active	CloudTrail	Unauthorized read/write actions outside defined scope.
K8s_Root_Escalate	Active	Kubernetes	Attempts to escalate user permissions above service account level.
S3_Access_Filter	Disabled	S3 API	Blocked access attempts on non-public buckets (currently disabled).

I found 3 rules configured. Which one do you want the full logic for?

U I think someone logged into the console from a weird IP. Can you check the raw logs around that time?

🔎 **Raw Log Context (Past 15 Minutes)**

Source IP: `203.0.113.42`
Timestamp: `2026-06-12T10:15:22Z`

```json
{ "event": "auth_attempt", "user": "admin_svc", "status": "failure", "details": "Invalid credential set.", "source_ip": "203.0.113.42" }
{ "event": "api_call", "user": "unknown", "status": "success", "endpoint": "/v1/profile", "source_ip": "192.168.1.5" }
```

The logs show 4 failed attempts from that IP, followed by a successful call from an internal source.

U That critical alert 'Brute Force Attempt' looks like it was cleared up last week. Can you mark it as false positive?

✅ **Signal Triage Confirmation**

Signal `sig_12345` has been successfully updated.

- **Old State:** Open (Critical)
- **New State:** Archived
- **Reason Applied:** False Positive: Maintenance testing on network segment 7B.

This change is logged in your compliance audit trail.

Frequently Asked Questions

01 **How can I use Datadog Cloud SIEM MCP to find suspicious activity in my cloud logs?**

You can ask your agent to search raw logs directly, specifying a timeframe or an IP address. The system pulls the contextual log data and presents it conversationally, allowing you to immediately spot indicators of compromise without running complex queries.

02 **What if I want to change an alert status from open to closed?**

You can use the MCP to manage your signals. You simply tell your agent which signal needs updating and provide a reason (like 'false_positive'). This action archives the signal while creating a permanent, auditable record of the decision.

03 **Can I write new security rules using this Datadog Cloud SIEM MCP?**

Yes. You can define and deploy completely new detection rules by giving your agent raw field names, query bindings, and severity levels. This lets you adapt your threat monitoring to brand new services or attack vectors.

04 **Does the Datadog Cloud SIEM MCP only search alerts, or can it look at logs too?**

It does both. It runs high-level searches on existing security signals (alerts) and also allows you to perform deep threat hunting by querying raw log data over specific time ranges for full context.

05 **What if I need to check which logs are being blocked from my SIEM?**

You can ask the MCP to list security filters. This tool retrieves global exclusion policies, allowing you to confirm exactly what data vectors are intentionally excluded and why they aren't reaching your evaluation engine.

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-cloud-siem": { "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 Cloud SIEM 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 Cloud SIEM. 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 Cloud SIEM MCP
Server ID	`019d7581-d73c-7308-b3bb-ab53297a95e0`
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-cloud-siem.