

MCP SERVER

NO CODE

CLOUD HOSTED

Amazon CloudWatch Log Group MCP for AI Agents

Analyze AWS application logs and track service health metrics by grouping.

The Amazon CloudWatch Log Group MCP lets your AI agent securely query and filter log events from a single, specified CloudWatch Log Group. It provides immediate operational observability without granting broad AWS permissions, making it perfect for debugging application errors or analyzing traffic spikes safely.

F Quality Score 3.6/100

aws

cloud-logging

infrastructure-monitoring

log-analysis

security-scoping

devops



The infrastructure that powers AI agents in the real world.



Vinkius connects AI to the world's software through secure, enterprise-grade infrastructure — enabling real-world execution at scale, built on the Model Context Protocol (MCP).

Your AI Connections Run Through Vinkius Cloud

The world's largest
managed MCP catalog

Vinkius is the cloud infrastructure where AI agents connect 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.

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.

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.

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

Encryption at rest

Ed25519

PKI vault signatures

24h TTL

Ephemeral session keys

V8 Isolate

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.

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.

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.

06 — Honeypot Trap System

Phantom credentials are injected into isolated environments. If a honeypot 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.

FinOps

Cost tracking per tool, payload compression savings, budget optimization signals, and consumption trends.

Firewall & DLP

PII redaction activity, sensitive data protection counters, and security event timeline.

Agent Activity

Which AI clients are connecting, how often, and what they're doing — real-time session tracking.

Tool Health

Slowest and most error-prone tools, with actionable root-cause insights and performance baselines.

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.

Amazon CloudWatch Log Group MCP

1 tools available

Cloud-hosted on Vinkius

When production services fail, you can't afford to waste time clicking through dashboards or wading through massive console logs. This MCP gives your AI agent one specific, powerful ability: secure access to run deep searches on a single CloudWatch Log Group. The system is intentionally scoped down; it never sees your entire AWS log estate. Instead, your agent operates with surgical precision.

This means you can safely troubleshoot application errors and track infrastructure performance without the risk of accidentally viewing sensitive audit trails in other services. You simply prompt your AI client—asking for all records from a specific time frame or filtering by a unique error code—and it handles the complex data retrieval. Connecting this MCP via Vinkius's catalog lets any compatible agent immediately analyze operational metrics, turning overwhelming log streams into actionable insights.

Core Capabilities

01 — Filter Specific Log Events

The AI searches and filters for particular entries within the configured CloudWatch Log Group based on user-defined criteria.

One Click on Vinkius — From Prompt to Execution

Available at vinkius.com/mcp/amazon-cloudwatch-log-group — connect your AI agent in three steps.

- 01** You tell your AI agent what you're looking for, like 'Show me all login failures in the last hour.'
- 02** The MCP executes a targeted query, running the search only against the dedicated CloudWatch Log Group.
- 03** Your agent returns filtered log events, giving you an immediate list of relevant entries without needing to navigate AWS consoles.

The bottom line is that your AI client retrieves and filters specific log data from one defined group, turning raw logs into targeted information instantly.

Built For

This MCP is essential for SREs, DevOps Engineers, and Backend Developers who spend too much time manually digging through AWS dashboards at 2 AM. If your job requires finding specific failure patterns or tracking service health across logs, this tool saves hours of clicking.

Site Reliability Engineer (SRE)

They use the MCP to instantly check for cascading failures by running complex queries across log events in a single Log Group.

DevOps Engineer

They rely on it to validate deployment health, checking logs immediately after code pushes to detect unexpected error messages or configuration drift.

Backend Developer

They use it to debug complex user journeys by filtering log events for a specific User ID and tracking the sequence of calls that led to an issue.

What Changes When You Connect

- 01 Security-Scoped Access:** You don't risk exposing sensitive data. The agent is locked down to a single log group, providing contained observability for debugging.
- 02 Targeted Debugging:** Instead of sifting through millions of unrelated records, the `filter_log_events` tool pinpoints exactly the failure messages or user IDs you need immediately.
- 03 Saves Dashboard Time:** You eliminate the manual process of navigating AWS console dashboards. Your AI agent goes straight to the data and pulls out only what matters.
- 04 Deep Pattern Matching:** The MCP supports full query syntax, allowing your agent to aggregate log data across specific time windows and filter by JSON keys.
- 05 Actionable Insights:** It moves beyond just showing logs; it helps you identify trends, like repeated failed connection attempts or unusual traffic spikes.

Real-World Applications

Tracking a user's failing checkout process

The agent searches the log group using `filter_log_events` for a specific User ID and time range. It then shows the full sequence of events, identifying whether the failure occurred during payment processing or inventory checks.

Validating successful deployment

After rolling out new code, the agent runs a query to verify that all expected 'Service Initialized' messages appeared in the logs. This provides instant confirmation of application readiness.

Investigating intermittent API service errors

An engineer asks their agent to find all 'HTTP 503' status codes in the last four hours. The MCP queries and returns a list of instances, helping determine if the issue is localized or widespread.

Patterns to Avoid

Searching across all AWS logs

X AVOID

Trying to manually search global log groups for a single error message, which can return gigabytes of irrelevant data and cause throttling.

✓ INSTEAD

Use the Amazon CloudWatch Log Group MCP. By confining your query using `filter_log_events`` to one specific group, you get surgical results without wasting time or resources.

Relying on fragmented dashboards

X AVOID

Jumping between Metrics, Logs, and CloudWatch Dashboard tabs trying to piece together a single root cause failure.

✓ INSTEAD

Let your AI client query the log group directly. The agent handles the complexity of data retrieval, giving you a unified view of filtered logs.

Using vague search terms

X AVOID

Just searching for 'error' without specifying the time or service component, resulting in noise and missed critical details.

✓ INSTEAD

Use `filter_log_events`` to combine filters. For example: filter by 'ERROR' AND within the last 30 minutes AND for Service X.

The Right Fit

You should use this MCP if your primary need is deep, targeted analysis of operational logs from a known source. Specifically, when you need to find patterns, trace a user request across multiple log lines, or validate service health after a change, this tool works perfectly. However, don't use it if you need to monitor infrastructure metrics (like CPU utilization) or access audit trails outside the configured group; for those tasks, dedicated monitoring tools are required. If your goal is broad discovery across all AWS resources, this scoped MCP won't help—you'll need a different catalog solution.

Amazon CloudWatch Log Group MCP: Troubleshooting Application Errors in DevOps

Today, debugging an application failure means navigating multiple AWS consoles. You copy timestamps here, search for error codes there, and then cross-reference them in a separate logging dashboard. It's slow, prone to human error, and often requires combining data from several disconnected views just to find the root cause.

With this MCP, you talk to your agent like you talk to a coworker. You ask it directly: 'What happened when user 123 failed at checkout?' The agent uses `filter_log_events` to pull all relevant logs into one place and presents the narrative of failure—the precise steps that went wrong.

Amazon CloudWatch Log Group MCP: Monitoring Service Health in SRE Operations

Before this, validating a deployment meant setting up complex alert rules and manually checking the logs for specific success messages. If an alert fired, you'd still have to jump into the console, figure out which log group was relevant, and then run a manual query just to see what happened.

Now, your agent monitors service health by querying the configured Log Group. You simply ask it to 'Show me all successful heartbeat messages for Service X in the last 15 minutes.' It gives you an immediate, definitive answer—no dashboard required.

Amazon CloudWatch Log Group: 1 Tool for Filtering Log Events

Use the `filter_log_events` tool to search, filter, and extract specific log entries within the designated CloudWatch Log Group.

#	TOOL	DESCRIPTION
01	<code>filter_log_events</code>	Searches and filters log events within the configured CloudWatch Log Group based on a specified query or time range.

See It in Action

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

- U** I think there's a memory leak causing crashes. Find all logs mentioning 'OutOfMemoryError' from the past 8 hours.



System Log Analysis Report

- **Total Events Found:** 14 (Source: Service A)
- **Time Range:** Last 8 Hours
- **Primary Cause Identified:** The logs point to a recurring issue with the connection pool.

```
`` json
{
  "timestamp": "2026-06-12T14:35:11Z",
  "level": "ERROR",
  "message": "OutOfMemoryError: Connection Pool Exhausted.",
  "service": "auth_api"
}
``
```

This pattern suggests the connection pool size needs increasing.

U What were all the login attempts for user ID 456 in the last hour?



User Activity Trace: User ID 456

- **Successful Logins:** 2 (Most recent at 10:15 UTC)
- **Failed Attempts:** 3 (Potential brute force attempts detected)

TIMESTAMP	STATUS	SOURCE IP
14:58:01	FAILURE	203.0.113.4
14:57:59	FAILURE	203.0.113.4
14:56:12	SUCCESS	192.168.1.1

Action needed: Check IP address 203.0.113.4 for suspicious activity.

U Give me a summary of all critical warnings from the last three hours.



Warning Summary (Last 3 Hours)

The logs show several warnings, mostly related to external service timeouts.

- **Timeout Count:** 27 instances.
- **Affected Service:** Inventory API (`inv-svc`).
- **Root Message:** 'External dependency timeout exceeded (504).'

The pattern suggests the rate limiting on the external dependency needs review or a retry mechanism implemented.

Frequently Asked Questions

01 How do I use the Amazon CloudWatch Log Group MCP to debug an issue?

You simply tell your AI agent what you are looking for, like 'Show me all errors from the last hour.' The MCP connects and filters the logs in that specific group for you. It turns a massive data dump into a focused list of actionable events.

02 Is the Amazon CloudWatch Log Group MCP safe to use with my production environment?

Yes, it is highly secure because it only allows your agent to query logs from one pre-selected group. It doesn't give access to your entire AWS account or other sensitive log groups.

03 Can I use the Amazon CloudWatch Log Group MCP to find user activity?

Absolutely. You can ask it to track a specific User ID across all relevant logs in that group, letting you see the exact sequence of events—successes and failures alike.

04 What kind of data does this MCP analyze for me?

It analyzes standard log formats, including error messages, warning flags, request details, IP addresses, timestamps, and structured JSON data. It's designed to find patterns in operational logs.

05 If I need more than one log group, can the Amazon CloudWatch Log Group MCP handle it?







No. This MCP is intentionally scoped for maximum security; it works with only one specific CloudWatch Log Group at a time. If you need multiple groups, you'll need to connect several separate MCPs.

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](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 <code>"amazon-cloudwatch-log-group": { "url": "..." }</code>
 Windsurf	MCP Settings → <code>mcp_settings.json</code> → 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

Amazon CloudWatch Log Group 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](https://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 Amazon CloudWatch Log Group. 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	Amazon CloudWatch Log Group MCP
Server ID	019e3861-f09e-73ca-9116-76134b6f085f
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/amazon-cloudwatch-log-group.