

MCP SERVER

NO CODE

CLOUD HOSTED

Lacework MCP

Audit your cloud security posture instantly.

Lacework (Cloud Security & CNAPP) connects your AI agent to deep cloud security data. You can search behavioral alerts for anomalies like AWS IAM brute-forcing or Kubernetes breakouts. It audits cloud assets, scans container images, and checks live hosts for critical vulnerabilities using specialized query language.

A+ Quality Score 100/100

cnapp

threat-detection

vulnerability-scanning

cloud-security

kubernetes-security

iam-auditing



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.

Lacework (Cloud Security & CNAPP) MCP

10 tools available
Cloud-hosted on Vinkius

Connecting Lacework's security data directly into your AI client changes how you hunt threats in the cloud. Instead of clicking through endless dashboards trying to piece together what went wrong, you talk to your agent. Your agent handles the complex queries across your entire infrastructure footprint. You can ask it to find all running instances that might be exposed or check if any container image has a known weakness before deployment. When you run into complexity—like mapping out every single unrestricted S3 bucket—your Vinkius connection lets you access those detailed logs conversationally. It's about getting immediate, actionable answers on your cloud security posture without manual dashboard filtering.

Core Capabilities

01 — Search Behavioral Security Alerts

Find deep telemetry data related to anomalous activity, such as unusual Kubernetes processes or AWS access attempts.

03 — Identify Host Vulnerabilities

Check live VMs (like EC2 or GCE) to see which critical vulnerabilities are currently executing on the machine.

05 — Check for Specific Vulnerability Exposure

Pinpoint exactly which nodes across your entire cloud setup are exposed to a specific flaw, like Log4j.

02 — Audit Cloud Asset Inventory

Get a real-time list of every running instance and any unrestricted cloud resources across your accounts.

04 — Scan Container Image Flaws

Examine images stored in registries like ECR or DockerHub for known CVEs before they get promoted into production.

06 — Run Advanced Threat Queries

Execute custom queries using Lacework Query Language (LQL) to analyze vast datasets for patterns of abuse or unusual activity.

One Click on Vinkius — From Prompt to Execution

Available at vinkius.com/mcp/lacework-cloud-security-cnapp — connect your AI agent in three steps.

- 01** Subscribe to this MCP and provide your Lacework Account Key ID and Secret.
- 02** Direct your AI client, like Claude or Cursor, to the connection. The agent now has access to your live cloud security data.
- 03** Ask a direct question, for example: 'Show me all unrestricted S3 buckets.' Your agent runs the necessary query and returns a clean list of assets.

The bottom line is you get immediate visibility into complex cloud risks without ever having to navigate a dashboard or write a query yourself.

Built For

This connector is built for the security professional who spends too much time clicking between dashboards. It's for analysts and engineers who need to know, right now, if a vulnerability exists, where it lives, and what resources are exposed.

Security Analyst

Investigates anomalous alerts by asking the agent to fetch deep behavioral payloads for specific incidents.

DevOps Engineer

Verifies that container images and running hosts are free of critical flaws before promoting them through CI/CD pipelines.

Cloud Compliance Officer

Audits global security policies and checks for exposed, unmanaged cloud assets to maintain regulatory compliance.

What Changes When You Connect

- 01** Stop manually searching dashboards. Use the agent to run an `execute_query` for complex threat hunting, finding anomalies like API key abuse in seconds.

-
- 02** Don't wait for incidents. Run a vulnerability check using `list_host_vulnerabilities` or `list_container_vulnerabilities` to proactively find weaknesses before they are exploited.
-
- 03** Eliminate blind spots. Use `search_ccloud_inventory` to discover every single running asset, especially those unrestricted S3 buckets that should be locked down.
-
- 04** Respond faster during an emergency. With `search_cve_exposure`, you can instantly map out every vulnerable machine when a zero-day exploit hits.
-
- 05** Keep your infrastructure clean. Use the agent to review all security policies via `list_security_policies` and ensure continuous compliance auditing.
-

Real-World Applications

Finding the Source of an Outage

An engineer notices service degradation and needs to know if a recent change introduced a vulnerability. They ask their agent to run `list_host_vulnerabilities` on the affected cluster, quickly identifying two high-impact CVEs that need patching.

Compliance Audit for Public Data

A compliance officer needs proof that no sensitive data is publicly exposed. They ask the agent to run `search_cloud_inventory`, which immediately flags two unrestricted S3 buckets requiring policy lockdown.

Pre-Deployment Security Check

A DevOps team is ready to push a new microservice. Instead of manual testing, they use `list_container_vulnerabilities` via the agent to scan the image registry and confirm zero critical flaws.

Investigating a Suspicious Login Spike

The security team detects an unusual login pattern. Instead of manually sifting through logs, they use the agent to `search_alerts` for suspicious activity and then run `execute_query` for behavioral confirmation.

Patterns to Avoid

Treating it like a simple dashboard filter

✗ AVOID

Thinking you can just type 'show me all vulnerabilities' and get a basic list. This ignores context, severity, or resource group.

✓ INSTEAD

You need to be specific. Use `search_cve_exposure` to target one CVE across your whole footprint, or use `list_host_vulnerabilities` for only Critical/High issues in the 'Production' resource group.

Ignoring asset visibility

✗ AVOID

Only checking resources within a known VPC. This leaves unattached S3 buckets and cross-account assets completely blind.

✓ INSTEAD

Always start with `search_cloud_inventory` to map the entire control plane first. It finds every resource, including those outside your expected network perimeters.

Relying on manual policy checks

✗ AVOID

Manually reviewing security policies one by one to see if a specific risk is covered.

✓ INSTEAD

Use `list_security_policies` and then ask the agent to audit them against your requirements. It confirms if Lacework will even alert you for structural violations.

The Right Fit

Use this MCP if your security process requires correlating data from multiple, disparate cloud sources—for example, linking an anomalous login event (`search_alerts`) to a specific vulnerable host (`list_host_vulnerabilities`) and confirming the resource's location in the inventory (`search_cloud_inventory`). It excels at cross-functional threat hunting. Don't use it if you just need simple documentation retrieval; for that, checking `list_lql_queries` first will help define your scope. If you only care about a single type of vulnerability across one specific application stack, a dedicated scanning tool might be better. But if the problem is 'Where are we vulnerable right now?'—this MCP has the answers.

Security teams spend hours clicking through tabs just to map risk.

Today, finding out what's exposed feels like a scavenger hunt. You jump into the dashboard for alerts, then switch to another tool to check inventory, and finally hop over to a console to manually list host vulnerabilities. Copying IDs from one screen and pasting them into another is how most threat hunting gets done.

With this MCP connection, you just talk to your agent. You tell it, 'Show me all critical risks in the Production group.' It runs the necessary checks—pulling data from alerts, inventory, and host vulnerability lists—and gives you one consolidated answer. No clicking required.

Lacework (Cloud Security & CNAPP) MCP: Full Visibility

You no longer have to manually verify if a resource is restricted or what the policy surrounding it actually is. The agent runs `search_cloud_inventory` and correlates that output with `list_security_policies`, giving you immediate confidence in your posture.

This isn't just viewing data; it's asking questions of your entire cloud estate and getting definitive, actionable answers back. It fundamentally changes the speed at which you can respond to a threat.

Lacework (Cloud Security & CNAPP) MCP with 10 Tools

These tools let you programmatically interact with Lacework's security data to audit cloud resources, scan vulnerabilities, and analyze behavioral alerts through your AI agent.

#	TOOL	DESCRIPTION
01	<code>list_container_vulnerabilities</code>	Checks container registries or deployment clusters to list any static image vulnerabilities found before a build goes live.
02	<code>get_alert</code>	Retrieves the detailed data payload for an alert, showing exactly what behavior deviated from the norm and which accounts were involved.
03	<code>list_host_vulnerabilities</code>	Identifies critical or high-impact vulnerabilities that are actively running on specific cloud hosts or virtual machines.
04	<code>search_cloud_inventory</code>	Queries the real-time asset inventory to dynamically list all active instances, networking perimeters, and unmanaged S3 buckets.
05	<code>search_alerts</code>	Fetches security events related to anomalous Kubernetes activity, AWS IAM brute force attempts, or massive data transfers within a time frame.
06	<code>search_cve_exposure</code>	Filters the entire cloud infrastructure to show precisely which machines are currently vulnerable to a given CVE identifier.
07	<code>list_security_policies</code>	Lists all global security policies enforced by Lacework, confirming what structural norms must be maintained.
08	<code>list_lql_queries</code>	Retrieves the available structure of custom queries so you know how to check for specific compliance requirements.
09	<code>list_resource_groups</code>	Lists logical groups (like 'Production' or 'Staging') that help organize and evaluate security policies across your cloud architecture.
10	<code>execute_query</code>	Runs a custom Lacework Query Language request to analyze large datasets for specific patterns, such as anomalous login sequences.

See It in Action

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

U Search for all Critical alerts from the last 24 hours



Searching alerts... I've identified 3 Critical alerts in your Lacework account: 1) Anomalous Kubernetes Process, 2) AWS Root Account Access from new IP, 3) Massive Data Exfiltration detected. Which one would you like to investigate first?

U List all host vulnerabilities for our Production resource group



Retrieving host vulnerabilities... Found 12 active CVEs across 4 instances in the Production group. Highlights: 2 Critical (Log4j related) and 5 High impact. I can provide the Machine IDs and remediation steps for the critical ones.

U Are there any unrestricted S3 buckets currently visible in our inventory?



Auditing cloud inventory... I've discovered 2 unrestricted S3 buckets (AWS:S3:Bucket): 'public-assets-prod' and 'temp-data-dump'. Both allow world-readable access. Would you like to see the associated security policies for these assets?

Frequently Asked Questions

01 How does Lacework (Cloud Security & CNAPP) MCP find unrestricted S3 buckets?

It uses the `search_cloud_inventory` tool to query the real-time cloud control plane. This finds any bucket that is publicly readable or writable, regardless of where it appears in your account structure.

02 What if I want to check for a specific vulnerability like Log4j?

You use `search_cve_exposure`. You provide the CVE ID, and this MCP filters all integrated machines across your cloud estate to tell you exactly which nodes are impacted.

03 Can I find evidence of a brute force attempt using Lacework (Cloud Security & CNAPP) MCP?

Yes. Running `search_alerts` will fetch events related to AWS IAM brute-forcing attempts, giving you the specific time window and accounts involved in the attack.

04 Does this MCP only check my live VMs?

No. It checks both running hosts using `list_host_vulnerabilities` AND it scans container images in registries like ECR/DockerHub using `list_container_vulnerabilities`.

05 What is the best way to use Lacework (Cloud Security & CNAPP) MCP for compliance?







First, run `list_security_policies` to understand your ruleset. Then, use a custom query via `execute_query` to test specific compliance checks against your actual data.

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 <code>"lacework-cloud-security-cnapp": { "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

Lacework (Cloud Security & CNAPP) 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 Lacework (Cloud Security & CNAPP). 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	Lacework (Cloud Security & CNAPP) MCP
Server ID	019d75c3-aef6-7074-9995-43120b6aae55
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/lacework-cloud-security-cnapp.