MCP SERVER | NO CODE | CLOUD HOSTED

Wazuh (SIEM) MCP

Query logs, status, and compliance using natural language.

Wazuh (SIEM) connects security operations and endpoint monitoring directly to any AI agent. Instantly list agents, check compliance reports, and pull manager logs using natural conversation. It lets you run complex security queries—like checking File Integrity Monitoring or mapping MITRE ATT&CK tactics—without ever leaving your chat interface.

A+ Quality Score 100/100

siem | threat-detection | vulnerability-management | endpoint-security | incident-response

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

Wazuh (SIEM) MCP

21 tools available
Cloud-hosted on Vinkius

Managing a Security Information and Event Management (SIEM) system usually means jumping between dashboards, running command-line tools, and filtering massive amounts of data. This MCP changes that process entirely. You connect it to any AI agent through Vinkius, giving your client the ability to speak directly to your Wazuh environment.

Instead of writing complex queries or navigating deep menu structures, you simply ask questions about your infrastructure. Your agent handles everything from checking if cluster nodes are healthy to retrieving security configuration assessment results across all endpoints. This means you get immediate answers on agent status, threat intelligence mappings, and audit data without ever needing to log into the Wazuh UI.

Core Capabilities

01 — Audit System Compliance

Fetch detailed compliance reports from modules like Rootcheck or Security Configuration Assessment (SCA) to confirm endpoint hardening.

02 — Manage Endpoint Agents

List, enroll, restart, or upgrade all agents across the network using simple commands in your AI client.

03 — Analyze Threat Data

Retrieve MITRE ATT&CK mappings and run log decoders to validate threat detection capabilities against specific attack vectors.

04 — Inspect Core Logs & Status

Pull live logs from the manager daemon or check the overall health of the cluster nodes instantly.

05 — Refine Security Rules

List, update, or test security rules and decoders against sample log data to improve detection accuracy.

One Click on Vinkius — From Prompt to Execution

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

01 First, you subscribe to this MCP on Vinkius and provide your specific Wazuh API URL, username, and password.

02 Next, you activate the connection within your preferred AI client (Claude, Cursor, etc.).

03 Finally, tell your agent what you need—like 'Show me all failed SCA checks for agents in the finance department.' The MCP executes the query and returns structured data.

The bottom line is that this connection lets your AI client treat complex security infrastructure like a simple API endpoint, turning manual console work into conversational queries.

Built For

This MCP is for the Security Analyst who hates dashboard clicks and the DevSecOps Engineer who needs to automate agent lifecycle management from their terminal. If your job involves checking compliance or hunting through logs, this tool saves hours of manual work.

Security Analyst

Quickly query agent status, run File Integrity Monitoring (Syscheck) reports, and cross-reference findings with MITRE ATT&CK data during an active investigation.

DevSecOps Engineer

Automate repetitive tasks like upgrading agents or monitoring cluster health directly from a terminal-based AI workflow without manual SSH connections.

Incident Responder

Pull manager logs and check for connection warnings immediately during an active incident, bypassing the need to navigate multiple log viewer tabs.

What Changes When You Connect

01 Stop manual dashboard diving. Instead of clicking through tabs to check agent status, just ask your AI client to `list_agents`. You get the list instantly in plain text.

02 Accelerate incident response. When you need to know if a system was tampered with, use `get_syscheck` to pull File Integrity Monitoring reports immediately, without running console commands.

03 Improve compliance posture checks. Instead of manually checking dozens of policies, ask for the latest Security Configuration Assessment (SCA) results using `get_sca`, and get actionable failure points.

04 Automate maintenance. Need to update a bunch of machines? Run `upgrade_agents` or use `restart_agents`. It's one simple command instead of coordinating multiple SSH sessions.

05 Deepen threat hunting. Use the MCP to pull MITRE ATT&CK mappings via `get_mitre`, which lets you instantly map observed attacker behavior to industry-standard tactics.

Real-World Applications

Investigating a potential breach

An incident hits the network. Instead of logging into three different dashboards, the analyst asks their agent to check `get_syscheck` for file changes and then run `get_mitre` to see if those changes match known attack patterns. The results come back together in one chat window.

Quarterly compliance audit

The auditor needs proof that all agents meet minimum security standards. The DevSecOps engineer simply calls `get_sca` and runs the report through the agent, getting a consolidated list of failures across hundreds of endpoints.

Cluster troubleshooting

Agents start failing randomly. Instead of logging into the cluster manager to check services, the engineer asks the MCP to run `get_manager_status` and then `list_cluster_nodes`. The AI client pinpoints which specific node is offline.

Tuning detection rules

A new log format comes in. Instead of writing a complex decoder, the analyst uses `get_logtest` to feed sample logs and test if existing rules are interpreting the data correctly before deploying changes using `update_rule_file`.

Patterns to Avoid

Using raw API calls

✗ AVOID

Manually constructing complex WQL filters or crafting multiple multi-step API requests to get agent status and compliance data.

✓ INSTEAD

Just ask your AI client directly. Use the natural language interface to run `list_agents` combined with `get_syscheck`. The MCP handles all the underlying complexity for you.

Over-relying on the GUI

✗ AVOID

Opening the Wazuh dashboard, clicking through 'Agents', then switching tabs to 'Compliance' just to gather a list of endpoints and their security roles.

✓ INSTEAD

Use `list_agents` followed by `get_sca`. The MCP pulls all that data into one flow. It's faster, cleaner, and easier to track.

Ignoring cluster health

✗ AVOID

Assuming everything is fine when a service suddenly slows down, without checking the underlying manager processes.

✓ INSTEAD

Always check first by requesting `get_manager_status` or running `restart_cluster` if necessary. This ensures you know the core infrastructure isn't failing.

The Right Fit

Use this MCP if your primary job revolves around querying massive, complex datasets—specifically security event logs, compliance reports, and agent status. You need to ask 'Why?' or 'What changed?' regarding system integrity.

Don't use it if you just need simple ticketing or task management (use a dedicated ITSM tool). Don't use it if your core job is drafting

documents or running basic CRUD operations unrelated to security data. If you only need to list users, `list_security_users` works, but if you need context on *why* those users exist and what they access, this MCP is necessary.

The daily struggle of SIEM dashboards

You know the drill. An alert fires at 2 AM. You log into the Wazuh dashboard. First, you navigate to 'Agents' just to see who's online. Then you have to check a different tab for File Integrity Monitoring results. Next, you might need to run a manual report on configuration assessment failures. It's a clicking nightmare—jumping between tabs and copying data into separate spreadsheets.

With this MCP connection, the process flips. You tell your agent what you're looking for in plain English. Your client runs `list_agents` and then automatically pulls `get_syscheck` results and `get_sca` reports, compiling it all into one clean response right where you are working.

Control agents, logs, and security roles with the Wazuh (SIEM) MCP

The manual steps that vanish include: navigating to the 'System' panel; manually running agent status checks; and then having to jump over to the 'Compliance' section for audit data. You don't need those clicks anymore.

You just ask, and your AI client gives you a structured answer. It's instant access to deep system intelligence that used to take thirty minutes of painful dashboard navigation.

Wazuh (SIEM) MCP – 21 Tools

Use these 21 tools to control agents, audit security policies, test rules, and retrieve deep logs directly from your AI agent.

#	TOOL	DESCRIPTION
01	`list_cluster_nodes`	Retrieves a list of all nodes currently running in your Wazuh cluster.
02	`create_agent`	Enrolls and adds a brand new agent to the monitored network using specified details.
03	`create_security_role`	Defines and creates a specific security role within the Wazuh system for resource management.
04	`list_decoders`	Lists all currently loaded decoders, allowing you to see how log sources are interpreted.
05	`delete_agents`	Removes specified Wazuh agents from the monitoring system using a defined query filter.
06	`list_agents`	Provides a list of all monitored agents, supporting filters to narrow down results.
07	`get_logtest`	Tests specific security rules and decoders against sample log data to see if they trigger correctly.
08	`get_manager_logs`	Retrieves the latest operational logs from the Wazuh manager daemon for review.
09	`get_manager_status`	Checks and reports on the current running status of the main Wazuh manager service.
10	`get_mitre`	Fetches structured data mapping security findings to MITRE ATT&CK framework identifiers.
11	`restart_agents`	Initiates a restart command for selected or all monitored agents.
12	`restart_cluster`	Forces a full restart of the entire Wazuh cluster to resolve deep-seated service issues.
13	`get_rootcheck`	Runs and returns results from Rootcheck, which verifies system file integrity against known baselines.
14	`list_rules`	Lists all active security rules configured in Wazuh, with support for filtering.

#	TOOL	DESCRIPTION
15	get_sca	Runs and returns results from Security Configuration Assessment (SCA), checking system hardening policies.
16	list_security_users	Lists all user accounts that have API access credentials within Wazuh.
17	get_syscheck	Pulls the results from File Integrity Monitoring (Syscheck), showing changes to critical system files.
18	get_syscollector	Retrieves the current inventory of installed software and hardware components across monitored agents.
19	update_rule_file	Allows you to programmatically push updates or changes to a specific Wazuh rule file.
20	update_security_config	Updates the overall security configuration settings for the Wazuh manager system.
21	upgrade_agents	Initiates a controlled upgrade process for monitored agents to ensure they run the latest software version.

See It in Action

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

U List all Wazuh agents that are currently active.

I've retrieved the agent list. You have 12 active agents, including 'web-server-01' (Ubuntu) and 'db-master' (CentOS). Would you like to check the status of a specific one?

U Show me the latest Security Configuration Assessment (SCA) results.

Fetching SCA data... I found several failed checks related to SSH hardening and password policies on 3 agents. Should I list the specific failed policy IDs?

U Check the Wazuh manager logs for any recent errors.

Inspecting manager logs... I found a few 'connection refused' warnings from 10 minutes ago related to agent enrollment. The rest of the daemons are running normally.

Frequently Asked Questions

01 **How do I use the Wazuh (SIEM) MCP to check endpoint compliance?**

To audit security posture, ask your agent to run `get_sca`. This executes the Security Configuration Assessment and provides a list of policies that are failing across your monitored agents.

02 **Can I find out which agents are online using Wazuh (SIEM) MCP?**

Yes, simply ask your agent to run `list_agents`. It uses the system's filtering capabilities to give you a current list of all monitored endpoints.

03 **What is the best way to test new security rules with Wazuh (SIEM) MCP?**

Use the `get_logtest` tool. You can provide sample logs and let the agent run them against your current set of rules and decoders to validate if they trigger correctly.

04 **How do I check the overall health of my Wazuh cluster?**

Check manager status using `get_manager_status`. If that looks okay, you can also run `list_cluster_nodes` to verify every node is communicating correctly.

05 **What if I need to update a rule file after finding an issue with Wazuh (SIEM) MCP?**

You use the `update_rule_file` tool. After troubleshooting, you can push changes directly to your rules without manual API calls.

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 `"wazuh-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

Wazuh (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 Wazuh (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	Wazuh (SIEM) MCP
Server ID	`019e3909-476f-70fb-ad4f-801165cf5846`
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/wazuh-siem.