MCP SERVER | NO CODE | CLOUD HOSTED

Elastic Security MCP for AI Agents

Audit and Manage Threat Detection Rules in SIEM Environments

Elastic Security connects your AI client directly to your SIEM environment, giving you conversational control over threat detection and SOC auditing. You can search raw security alerts, manage complex custom rulesets, audit MITRE ATT&CK coverage, and handle exceptions—all without leaving the chat window.

A+ Quality Score 100/100

siem | threat-detection | soc-operations | cybersecurity | alert-management | mitre-att-ck

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

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.

Elastic Security MCP

10 tools available
Cloud-hosted on Vinkius

Managing a modern security stack is complicated. Usually, checking rule logic or searching for specific threats means jumping through three dashboards, running CLI commands, and cross-referencing spreadsheets. This MCP lets you skip all that overhead. Instead of navigating complex consoles, you talk to your AI client and tell it exactly what you need done with your threat detection environment.

It's like having a security expert sitting next to you who has instant access to every rule, alert, and log entry in the system. Need to check if a new ransomware variant is covered? Just ask. Found a false positive from a scanner? Whitelist it instantly. This MCP brings that level of detailed control straight into your conversation flow. Vinkius hosts this connection, so you can plug directly into your favorite AI client and start managing your SIEM operations immediately.

Core Capabilities

01 — Search and correlate raw security alerts

Retrieve comprehensive security signals by searching across hostnames, user profiles, IP geolocations, and full process trees.

02 — Create and modify detection logic

Build new custom log detection rules or update existing ones to track malicious activity patterns (TTPs) in real-time.

03 — Audit threat coverage and rule status

Search for specific rules by MITRE tactic, check if official prepackaged rules need updates, or list all configured detection rules for gap analysis.

04 — Manage false positive exceptions

Whitelist hostnames in exception lists or add global exception records to prevent known-good administrative behavior from triggering alerts.

05 — Control rule lifecycles and state

Irreversibly delete custom rules or enable/disable specific detection rules across large organizational units as needed for tuning.

One Click on Vinkius — From Prompt to Execution

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

01 Subscribe to this MCP and provide your Kibana Host, Port, and Elastic API Key.

02 Your AI client connects to the service, authenticating your access rights across the security stack.

03 You interact with the system using natural language prompts to execute complex tasks like searching signals or updating detection rules.

The bottom line is you manage threat hunting and SOC operations entirely through conversation.

Built For

This MCP is for security professionals who spend too much time clicking between dashboards, running manual searches, or writing complex queries just to answer simple questions. It's built for the SOC Analyst tired of context switching and the Security Engineer needing instant rule deployment.

SOC Analyst

Using this MCP, you can monitor live security alerts and audit detection rules without ever leaving your chat interface. You get immediate visibility into suspicious activity.

Security Engineer

You manage the full life cycle of detection logic—creating new rules or adjusting existing ones—all using plain English commands to ensure maximum coverage.

Incident Responder

During an active investigation, you quickly search for signals and verify threat coverage against known CVEs without needing deep knowledge of the underlying index structure.

What Changes When You Connect

01 Stop manually cross-referencing rule logic. Use the `find_detection_rules` tool to search by MITRE tactic or name, instantly showing if your coverage is adequate for new threats.

02 Reduce alert fatigue immediately. If you have false positives from scanners, use `add_exception` or `list_exceptions` to whitelist hosts and keep the noise down without disabling vital rules.

03 Gain full visibility into incidents with `search_signals`. Instead of piecing together data, you get a single view that consolidates hostnames, user profiles, and IP geolocations for every alert.

04 Maintain system health effortlessly. Run `get_prepackaged_rules_status` to verify if the official rules need updating, ensuring you're covered by the latest threat models.

05 Tweak your environment with precision. Use `update_rule` or `delete_rule` to manage detection rule state—disabling noisy triggers without deleting necessary logic.

Real-World Applications

Investigating a new ransomware pattern

The team notices unusual activity. They ask their agent to search for security signals from the last hour, focusing on user 'admin_root' and looking for process trees related to volume shadow copy deletion. The system returns specific alerts with source IPs.

Tuning false positive alert noise

The Security Engineer knows a vulnerability scanner runs weekly but triggers dozens of alerts. They tell the agent to check global exception lists and then use `add_exception` to whitelist the scanning host, clearing up the dashboard.

Auditing threat gaps for compliance

During a compliance review, the CISO needs proof of coverage for 'Lateral Movement'. They ask the agent to search detection rules specifically by the MITRE tactic tag. The system returns all relevant rule names and their current status.

Responding to zero-day reports

A new CVE is reported overnight. An Incident Responder asks for a list of all configured detection rules, filtered by 'CVE' or the affected asset type, allowing them to quickly verify if existing logic tracks the threat.

Patterns to Avoid

Blindly disabling rules

✗ AVOID

A junior analyst sees a flood of alerts and decides to simply disable all detection rules related to 'network traffic' because it’s too noisy.

✓ INSTEAD

Don't disable everything. First, use `list_detection_rules` to see the full inventory. Then, if you confirm a specific rule is causing noise, use `update_rule` to only change its state, keeping all other logic active.

Ignoring system health checks

✗ AVOID

The SOC team assumes their security coverage is up-to-date because they haven't had a major incident lately.

✓ INSTEAD

Always check the official status first. Run `get_prepackaged_rules_status` to confirm that your environment has the latest threat models for Windows, Linux, and Cloud before trusting your current defense posture.

Over-relying on manual investigation

✗ AVOID

An investigator must manually search logs, then check rule definitions, then look up user profiles in a separate database to piece together an attack timeline.

✓ INSTEAD

Instead, ask the agent to `search_signals` for the activity. The system consolidates all necessary data—hostnames, users, process trees—into one report.

The Right Fit

Use this MCP if your security workflow requires rapid, conversational access to deep SIEM functions. You need to manage the full detection rule lifecycle (create, delete, update) and correlate raw alerts with metadata like user profiles or geolocations. It's perfect for SOC Analysts who live in a chat interface.

Don't use it if your primary goal is simple log ingestion—if you just need to view pure, unstructured logs without threat context, an alternative logging tool might be better. Also, this MCP requires deep API knowledge (Kibana Host/API Keys), so ensure your team can manage those credentials before connecting.

Elastic Security MCP: Managing SIEM Detection Rules via AI Agents

Right now, managing threat detection rules is a grind. You spend hours in the console listing rules or writing complex KQL queries to check for coverage gaps against new attack patterns. If you need to verify if your current setup covers a specific MITRE Tactic, it means multiple clicks and switching between rule lists and mapping documents.

With this MCP, you simply ask: 'What detection rules track lateral movement?' The system instantly executes the query and presents the results in natural language, giving you immediate assurance of coverage without touching a single dashboard.

Elastic Security MCP: Auditing False Positives and Alert Signals

The worst part about SIEMs is false positives. Every time a scanner runs, you get dozens of alerts that require manual triage to whitelist the source. You're constantly copying hostnames and

Now, tell your agent: 'Whitelist this specific machine for vulnerability scans.' The MCP handles the `add_exception` logic instantly, resolving false positives and letting you focus on real threats.

pasting them into exception lists just to clear the noise.

Elastic Security MCP: 10 Tools for Alert & Rule Management

These tools allow you to perform deep SIEM operations like listing all rules, searching security signals, and managing exception lists using natural language commands.

#	TOOL	DESCRIPTION
01	add_exception	Adds a hostname to an exception list, telling the SIEM engine to ignore telemetry from that specific host for certain rules.
02	create_rule	Defines and activates a new log detection rule that tracks malicious activity patterns across your environment.
03	delete_rule	Permanently removes custom-written detection rules from the system, which is an irreversible action.
04	list_exceptions	Retrieves a list of global exception lists that manage specific rules and detect known administrative behavior to bypass alerts.
05	find_detection_rules	Searches for existing detection rules using criteria like name, tag, or MITRE tactic to quickly evaluate coverage against new threats.
06	get_rule	Retrieves exact details, including run intervals and query logic, for a single specified detection rule.
07	get_prepackaged_rules_status	Checks if your environment's official prepackaged rules are up to date against the latest threat models (Windows, Linux, Cloud).
08	list_detection_rules	Displays a comprehensive list of every configured detection rule within the SIEM system for full inventory.
09	search_signals	Searches raw, generated security alerts (Signals), enriching them with user profiles and IP geolocation data in one view.
10	update_rule	Enables or disables an existing detection rule to manage noisy triggers or re-activate rules after tuning.

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 active rules related to suspicious process execution.

Active Rules Found:
- **Volume Shadow Copy Deletion via WMIC:** Severity: High | MITRE Tactic: Defense Evasion
- **Suspicious PowerShell Call:** Severity: Medium | MITRE Tactic: Execution
- **Ransomware-linked Process Execution:** Severity: Critical | MITRE Tactic: Impact

You can use the full logic for any of these, or disable 'Suspicious PowerShell Call' if it generates too much noise.

U What happened with the user admin_root yesterday?

Security Signals for admin_root (Yesterday):
- **10:15 AM:** Suspicious LDAP Query detected from 192.168.1.5.
Source IP: 192.168.1.5 | *Process Tree:* winlogbeat → powershell.exe
- **10:45 AM:** Brute Force Attempt recorded on the Domain Controller.
Details: Failed logins from multiple IPs within a five-minute window.

U Check if my current detection rules are up to date.

System Health Check Results:
- **Windows Environment:** Status - Needs Update. The latest official threat model for credential dumping was released two days ago.
- **Linux/Cloud:** Status - Current. All required models are within the last 48 hours.

Action Recommended: Run an update on the Windows prepackaged rules immediately to close this gap.

Frequently Asked Questions

01 **How does the Elastic Security MCP improve my SOC alert management?**

The MCP lets you manage complex SIEM operations entirely through natural conversation. Instead of clicking between dashboards to find threat coverage, you can ask the agent directly if a specific vulnerability is tracked by existing rules.

02 **Can I use the Elastic Security MCP to handle false positive alerts?**

Yes. You can whitelist hostnames or add global exception lists using this MCP. This prevents known-good administrative activity, like scanner checks, from generating unnecessary alerts and cleaning up your dashboard.

03 **What kind of security events can I search for with the Elastic Security MCP?**

You can search raw generated security signals (alerts). The system consolidates all necessary metadata —hostnames, user profiles, and IP geolocations—into one view, making investigations much faster.

04 **Is this MCP good for auditing compliance against MITRE ATT&CK?**

Absolutely. You can find detection rules by specific tags or the MITRE tactic they cover. This lets you prove your coverage status quickly and easily, which is essential during audits.

05 **How do I update or modify an existing security rule using Elastic Security MCP?**

You can enable or disable rules using this MCP via natural language commands. This allows you to manage noisy triggers across large units without manually editing the rule logic in the console.

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 `"elastic-security": { "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

Elastic Security 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 Elastic Security. 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	Elastic Security MCP
Server ID	`019d758e-e1d2-7288-9b2a-4e5027cf644e`
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/elastic-security.