MCP SERVER | NO CODE | CLOUD HOSTED

Wallarm MCP

Turn API Security Audits into Natural Conversation

Wallarm MCP connects your AI agent to an enterprise API security platform. Monitor live traffic for attacks like SQLi and XSS, identify vulnerabilities in exposed endpoints, and manage IP allow/denylists—all through natural conversation. This lets you skip the security dashboard deep-dive and get immediate threat intel.

A+ Quality Score 100/100

api-security | waf | threat-detection | forensics | sql-injection | xss-protection

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

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.

Wallarm MCP

10 tools available
Cloud-hosted on Vinkius

Running a modern API means constantly worrying about who's hitting your endpoints and if they're safe. Instead of manually logging into complex security consoles, you just talk to your AI agent. This MCP turns that massive security headache into simple chat commands. You can ask the agent what attacks were detected recently, grouping threats by type like XSS or SQLi. Need to dig deeper? You can search through individual malicious requests, looking at full headers and payloads for forensic details. It also helps you find vulnerabilities—the agent lists them up so you know exactly what needs fixing. Plus, you can check the health of your WAF nodes or instantly block bad actors by managing IP rules. All this deep security data is available in one place via Vinkius, letting your AI client act like a full-time SOC analyst.

Core Capabilities

01 — Find detected attacks

Search for recent security threats and group them by the attack type (like XSS or SQLi).

02 — Forensically analyze payloads

Deeply search intercepted traffic to view full headers and payloads from malicious HTTP requests.

03 — List security vulnerabilities

Get a list of all open vulnerabilities found in the live API traffic, including diagnostic data for remediation.

04 — Manage IP access rules

Add or remove specific IPs or CIDR ranges to your global allowlist or denylist.

05 — View API endpoint map

Automatically pull a list of every exposed API endpoint and method found in the traffic.

One Click on Vinkius — From Prompt to Execution

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

01 Subscribe to this MCP, then enter your Wallarm API Token and Client ID into your AI client.

02 Your agent connects using those credentials, granting it read/write access across your security dashboard tools.

03 You simply ask a question—like 'What's the status of our filtering nodes?'—and the agent executes the necessary action.

The bottom line is you get instant, actionable API threat intelligence without ever leaving your chat window.

Built For

This MCP is built for security and platform teams. It's for the DevSecOps engineer who needs to triage critical vulnerabilities in minutes instead of hours. It’s for SOC analysts who need rapid incident forensics on demand, and API developers who want to verify their exposed endpoints are secure before deployment.

Security Engineer (DevSecOps)

Monitoring live traffic for zero-day threats or listing vulnerabilities during the CI/CD pipeline.

SOC Analyst

Responding to an alert by searching security hits and immediately blocking malicious IPs via chat commands.

API Developer

Verifying the entire list of exposed API endpoints to ensure all methods are properly secured.

What Changes When You Connect

01 Stop manually digging through security dashboards. With this MCP, you simply ask your agent to 'List all open vulnerabilities,' and it pulls the exact report data instantly.

02 Manage access rules without logging into a separate console. You can use the `create_ip_acl_rule` tool to add or remove IPs globally via chat.

03 Drill down on threats immediately. Instead of wading through logs, you run `search_security_hits` to see full payloads for any malicious request.

04 Maintain visibility into your entire attack surface by running `get_discovered_api_inventory` and getting a complete map of exposed endpoints.

05 Accelerate incident response time. You can use the agent to search attacks via `search_security_attacks`, which groups threats by vector, saving critical minutes.

Real-World Applications

Immediate Threat Triage

A SOC analyst notices unusual traffic spikes. Instead of jumping between the WAF logs and the vulnerability tracker, they ask their agent to 'Search for security attacks.' The MCP responds with grouped threats (e.g., 5 XSS attempts), allowing them to immediately focus on remediation.

Patching Vulnerabilities

A DevSecOps engineer needs to assess the risk of a recently found vulnerability. They run 'Search for vulnerabilities' and find an IDOR issue. Using `get_vulnerability_details`, they get the full diagnostic data needed to write a fix.

Onboarding New Services

An API developer launches a new microservice endpoint. They use 'Get discovered API inventory' through their agent, verifying that the MCP has successfully cataloged all exposed methods and endpoints for security review.

Blocking Malicious Users

During an active breach attempt, the team identifies a bad IP address. They use `create_ip_acl_rule` via chat to instantly add the IP to the global denylist, blocking further access without manual rule deployment.

Patterns to Avoid

Copying and pasting logs

✗ AVOID

The analyst has to navigate six different tabs—WAF status, attack reports, payload details, IP lists—and manually copy the relevant findings into a ticket or spreadsheet.

✓ INSTEAD

Let your agent do the heavy lifting. Use `search_security_hits` for forensic payloads and then use `list_ip_acl_rules` to document the required block action, all in one conversation.

Relying on dashboards alone

✗ AVOID

The team views a dashboard that says 'Vulnerabilities Found: 12.' They then have to click into 12 different records to understand the actual impact and fix status.

✓ INSTEAD

Instead, ask your agent to run `search_vulnerabilities`. It lists all open issues upfront, and you can use `get_vulnerability_details` for deep context on any specific ID.

Manual rule deployment

✗ AVOID

An attacker is identified. A human must log into the firewall console, locate the IP management section, and manually add a new block rule.

✓ INSTEAD

Use `create_ip_acl_rule` via your agent to instantly enforce an IP ban by specifying 'black' list type. It's faster than any UI click.

The Right Fit

You should use this MCP if you need a single, conversational interface for managing high-stakes API security operations. This is perfect when your workflow involves checking multiple systems—for example, confirming an attack occurred (using `search_security_attacks`), finding the details of the exploit (`get_vulnerability_details`), and then immediately blocking the source IP (`create_ip_acl_rule`). However, don't use this if you just need simple logging or metrics viewing. If your only goal is to track monthly usage statistics, a dedicated billing API will be better suited. This tool is for *actionable* security intelligence, not passive reporting.

API Security Audits Are Too Hard to Manually Track

Today, managing an API's security posture means living in a nightmare of consoles. You jump from the WAF dashboard to check for attacks; then you open another tab to list vulnerabilities; after that, you dive into payload logs just to find one bad IP address. It’s constant switching, copy-pasting data between Jira and three different monitoring dashboards.

With this MCP, your agent handles the clicks. You ask a natural question—like 'What's wrong with our access controls?'—and it gathers all the necessary information: listing open vulnerabilities via `search_vulnerabilities`, checking node health with `list_filtering_nodes`, and even pulling the API inventory using `get_discovered_api_inventory`. The result is a single, comprehensive answer.

Wallarm MCP Gives You Real-Time Threat Command

The manual process of checking threat status involves finding an attack vector, searching for the specific hit payload, and then manually creating a rule to block it. That’s three separate workflows across multiple interfaces.

Now, you can coordinate these actions conversationally. Ask your agent to find attacks using `search_security_attacks`, review the payloads with `search_security_hits`, and immediately execute `create_ip_acl_rule` on the offending IP—all in one flow. You control the entire threat response cycle from chat.

Wallarm: 10 Tools for API Security Management

These tools let you run specific security operations—from listing all known vulnerabilities to instantly blocking malicious IP addresses—all through your AI chat client.

#	TOOL	DESCRIPTION
01	`create_ip_acl_rule`	Adds an IP or CIDR range to either the global allowlist or denylist.
02	`get_discovered_api_inventory`	Retrieves a comprehensive list of all API endpoints and methods automatically found in traffic.
03	`get_client_info`	Pulls details about your Wallarm account, subscription level, and current feature status.
04	`get_vulnerability_details`	Retrieves full diagnostic data and exploit evidence for a specific vulnerability ID.
05	`list_ip_acl_rules`	Displays all currently configured IP allowlist and denylist rules.
06	`list_filtering_nodes`	Shows the deployed status and health of your WAF/API gateway filtering nodes.
07	`search_security_attacks`	Searches for security attack clusters, grouping them by vector type like SQLi or XSS.
08	`search_security_hits`	Shows full request headers and payloads for individual malicious HTTP requests intercepted by WAF nodes.
09	`search_vulnerabilities`	Lists all open security vulnerabilities discovered from analyzing live API traffic.
10	`update_vulnerability_status`	Changes the lifecycle status of a vulnerability, marking it as closed or false positive.

See It in Action

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

U List all security attacks detected in the last hour.

I've searched the recent security events in Wallarm. I found 12 attacks clusters, including 3 SQL Injections targeting the /api/login endpoint and 5 XSS attempts on /search. Would you like to see the individual hits and payloads for any of these?

U Block the malicious IP address 1.2.3.4 immediately.

Understood. I have successfully added 1.2.3.4 to your global denylist (black list) in Wallarm. All traffic from this IP will now be blocked by your filtering nodes. Would you like to see if there are other IPs with similar behavior?

U What vulnerabilities are currently open in our production API?

I found 3 open vulnerabilities: 1. Broken Object Level Authorization (BOLA) on /api/user/{id} (Critical), 2. Information Disclosure via verbose error messages (Medium), and 3. Insecure Direct Object Reference (IDOR) on /api/orders (High). I can provide the remediation guidance for any of these.

Frequently Asked Questions

01 **How does Wallarm MCP help with finding vulnerabilities?**

The MCP lets you run `search_vulnerabilities` to list all open flaws found in live API traffic. You can then use `get_vulnerability_details` to get full diagnostic data and understand exactly how to fix it.

02 **Can Wallarm MCP help me block a bad IP?**

Yes, you use the `create_ip_acl_rule` tool. You simply ask your agent to add an IP to the global denylist or allowlist, and it executes the rule change for you.

03 **What is the purpose of get_discovered_api_inventory?**

This tool automatically gathers a map of every exposed API endpoint and method. It's crucial for auditing your entire attack surface to ensure nothing was accidentally left open.

04 **Does Wallarm MCP support finding XSS attacks?**

Yes, you can use `search_security_attacks` which groups detected threats by vector. This allows you to specifically find and review XSS or SQLi attempts that were intercepted.

05 **What if I need to change a vulnerability status?**

You use the `update_vulnerability_status` tool. You can mark vulnerabilities as 'closed' or 'falsepositive' directly through your agent, keeping your security records accurate.

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 `"wallarm": { "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

Wallarm 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 Wallarm. 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	Wallarm MCP
Server ID	`019d761e-1214-714d-83fe-00370e8b59dc`
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/wallarm.