MCP SERVER | NO CODE | CLOUD HOSTED

CrowdSec MCP for AI Agents

Analyze real-time firewall decisions and global IP reputation data

CrowdSec connects your network defense to any AI agent, letting you manage threat intelligence directly through conversation. Query active local firewall decisions, monitor real-time security updates, and check global IP reputation data without logging into a command line.

A+ Quality Score 98.33/100

threat-intelligence | firewall-management | ip-reputation | network-security | intrusion-prevention

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

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.

CrowdSec MCP

3 tools available
Cloud-hosted on Vinkius

This MCP gives your AI client full control over your threat intelligence pipeline and network monitoring. You can query the local decision API to see if an IP or range is currently blocked by your firewall. The agent also polls for real-time updates on any new bans or deleted decisions, keeping you instantly aware of changes in your security posture. Need to know if a source is malicious? Use the global Community Threat Intelligence data to fetch reputation scores and behavioral classifications for any IP address. Instead of digging through complex logs, your AI acts like a dedicated security analyst, summarizing suspicious activity right where you're working. Getting this connected via Vinkius means you can access all these tools from Claude, Cursor, or any other MCP-compatible client.

Core Capabilities

01 — Querying Local Decisions

Use the agent to ask about existing network blocks, policy decisions, or ranges configured in your local firewall.

02 — Streaming Decision Updates

The agent polls for real-time notifications on any new block or deletion event, keeping you instantly updated on changes.

03 — Checking Global CTI Reputation

You can fetch external threat data to assess an IP address's reputation and behavioral risk score worldwide.

One Click on Vinkius — From Prompt to Execution

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

01 Subscribe to this MCP and provide your CrowdSec LAPI URL, along with both the Local API Key and the Community Threat Intelligence (CTI) Key.

02 Your AI client uses these credentials to establish a connection to your local firewall system and the global threat intelligence network.

03 You prompt the agent with natural language—for example, 'What's the reputation of this IP?' or 'Are there active blocks for this range?'—and get immediate data back.

The bottom line is that your AI client handles all the complex API calls and log parsing, letting you talk to your firewall like it’s a person.

Built For

This MCP is built for security professionals who deal with constant IP reputation checks and incident response. It's perfect for the network engineer tired of switching between terminal windows and dashboard UIs, or the DevOps team needing immediate context during an active breach.

Security Engineer

Instantly check local decision statuses and global reputation metrics without leaving their primary command interface.

DevOps Team Member

Monitor security streams and verify suspicious IP behaviors during automated deployment or incident response directly from the IDE.

System Administrator

Automate auditing of blocked network ranges and investigate unusual traffic patterns using plain language prompts.

What Changes When You Connect

01 You get instant visibility into your local network state. Use the `get_decisions` tool to query all active blocks or policy decisions for specific IP ranges in plain English.

02 Stay updated on security changes without manually checking logs. The agent polls for new and deleted decisions using `get_decisions_stream`, providing a continuous, real-time context stream.

03 Stop guessing about malicious IPs. Run the `get_cti_smoke` tool to fetch global reputation data and threat classifications from the community network.

04 The MCP streamlines security auditing. Instead of complex CLI commands, your agent handles checking suspicious actors' metadata and classifications instantly.

05 It integrates directly into your existing workflow. You pull threat intelligence straight from your IDE or terminal, eliminating context switching.

Real-World Applications

Investigating a Sudden Traffic Spike

A system administrator suspects an IP is malicious but doesn't know why. They ask the agent to check its global reputation using `get_cti_smoke`. The agent returns that the IP is flagged as a 'Tor Exit Node', allowing the admin to immediately block it.

Reviewing Blocked Ranges After an Incident

A DevOps team member needs to know exactly which IPs were blocked in the last hour. They use `get_decisions` and get a list of all decisions, confirming that the suspicious range was correctly covered by policy.

Monitoring Firewall Changes During Maintenance

A security engineer needs to track if any blocks or policies change while they are working late. They set up a stream query using `get_decisions_stream` and get instant alerts on every single decision made.

Pre-deployment Vulnerability Check

Before deploying new services, the team uses the agent to check known bad IPs. They run `get_cti_smoke` against a list of potential endpoints and flag any that have high noise scores.

Patterns to Avoid

Manual Log Diving for IP Status

✗ AVOID

Spending twenty minutes in the command line, running multiple `grep` commands across firewall logs just to see if a single bad IP address was blocked yesterday.

✓ INSTEAD

Just ask your agent. The MCP handles this complex process; use the `get_decisions` tool to query the local decision status directly via conversation.

Ignoring Real-Time Security Changes

✗ AVOID

Assuming a block remains active forever, leading to missed alerts or delayed response when an old ban expires and needs manual re-verification.

✓ INSTEAD

Set up continuous monitoring. Use the `get_decisions_stream` tool so your agent alerts you immediately when decisions are added or removed.

Relying on Internal Knowledge for Threat Scoring

✗ AVOID

Assuming all attackers come from known sources, and missing a threat actor that uses a new, clean-looking IP address.

✓ INSTEAD

Always check the global context. Use `get_cti_smoke` to pull in external reputation metrics and community threat intelligence for any suspect IP.

The Right Fit

Use this MCP if your job requires constant visibility into network boundaries, local firewall policies, or external threat scoring.

Specifically, you need a single place to query active decisions (`get_decisions`), monitor changes in real time (`get_decisions_stream`), and check global reputation (`get_cti_smoke`). Don't use this if you only need static policy documentation; that requires a separate knowledge base tool. If your goal is just to write firewall rules, those tools belong in a different category entirely.

CrowdSec MCP for AI Agents: Solving Local Firewall Decision Queries

Today, checking the status of an IP requires navigating complex command-line interfaces (CLIs). You run multiple commands to check if a specific address is banned locally, then you might have to cross-reference that with another policy system. It's tedious clicking and copy-pasting just to answer: 'Is this safe?'

With this MCP, asking the agent about an IP's local status is simple. You get immediate confirmation on whether a decision exists using `get_decisions` . The result isn't raw JSON; it's a clear, actionable statement telling you exactly what policies are in place and when they expire.

CrowdSec MCP for AI Agents: Enhancing Global Threat Intelligence Context

The biggest gap is relying only on internal data. You might block an IP because of local policy, but you have no idea if that same IP has a global reputation score or if it’s linked to other threat actors. That context is usually found in separate, manual reports.

Now your agent pulls that external intelligence directly into the conversation using `get_cti_smoke` . You don't just know *if* you blocked something; you understand *why* globally, drastically improving your incident response capabilities.

3 Tools for IP Reputation & Firewall Decision Management

Use these tools to manage everything from checking a single IP's global risk score to querying every active local firewall decision.

#	TOOL	DESCRIPTION
01	`get_cti_smoke`	Retrieves the global threat intelligence reputation data for a specified IP address.
02	`get_decisions_stream`	Polls the local API to fetch real-time updates regarding new and deleted firewall decisions.
03	`get_decisions`	Queries the local decision API to list all current active blocks or policy decisions for an IP range.

See It in Action

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

U Check if there are any active decisions for IP 192.0.2.5 in our local CrowdSec database.

Local Decision Status Check:

- **IP:** 192.0.2.5
- **Status:** Active Ban (Confirmed)
- **Reason:** Port scan attempt (Nmap signature)
- **Applied:** 20 minutes ago
- **Expires:** In 4 hours and 10 minutes

This decision was made by the 'Firewall Policy Group' and is set to automatically lift at the specified time.

U What's the global reputation of IP 203.0.113.99 according to CrowdSec CTI?

CrowdSec Global Threat Intelligence Report

METRIC	VALUE
Classification	Tor Exit Node
Noise Score	High (85/100)
Activity Flagged	Scanning, Botnet Communication
Community Status	Multiple Blocklists Hit

The IP shows clear indicators of compromise and is frequently associated with automated attacks.

U Get the latest stream of decisions from CrowdSec to see recent blocks.

Decision Stream Update (Last 5 Minutes)

- ✅ **NEW BAN:** IP `193.20.1.1` blocked for SSH brute force.
- ⚠️ **UPDATE:** Decision on `172.16.0.5` changed from 'Monitor' to 'Block'.
- ❌ **DELETED:** Ban on `10.0.0.2` lifted automatically (Expired).

Three new blocks occurred in the last interval, primarily targeting web services.

Frequently Asked Questions

01 **How do I check if an IP is banned locally using the CrowdSec MCP for AI Agents?**

You can ask the agent to query your local decision API. It will tell you immediately if a block exists, why it was applied (e.g., 'port scan'), and when that ban automatically expires.

02 **Does the CrowdSec MCP for AI Agents track changes in my firewall policies?**

Yes, the agent polls the decision stream so you get real-time updates on any new blocks or any decisions that are lifted. You never have to manually check if your security context is synchronized.

03 **What kind of reputation data can I get with the CrowdSec MCP for AI Agents?**

You fetch global IP reputation scores and classifications from the Community Threat Intelligence network. This tells you how many other systems globally have flagged that IP as suspicious or malicious.

04 **Is the CrowdSec MCP for AI Agents useful for DevOps teams during an incident?**

Absolutely. During an active breach, you can use the agent to check both local blocks and global reputation scores simultaneously, speeding up containment decisions by hours.

05 **Does this tool require me to be a security expert?**

No. The MCP is designed for natural conversation. You talk to your AI client like you're talking to a colleague; the agent handles all the technical API calls and data parsing.

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 `"crowdsec": { "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

CrowdSec 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 CrowdSec. 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	CrowdSec MCP
Server ID	`019e3881-32c2-7289-b7cb-c9be90d4cf07`
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/crowdsec.