MCP SERVER | NO CODE | CLOUD HOSTED

Halo Security MCP

Manage your entire attack surface conversationally.

Halo Security MCP automates attack surface management for your organization's digital perimeter. Monitor assets, scan vulnerabilities, and track risk scores—all through natural conversation with your AI agent. Add targets, list open ports, inspect certificates, and trigger new scans without leaving your chat interface.

A+ Quality Score 100/100

attack-surface-management | asset-discovery | security-posture | threat-monitoring | vulnerability-assessment

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

SCAN TO VIEW

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.

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.

Halo Security MCP

11 tools available
Cloud-hosted on Vinkius

Managing an entire attack surface used to mean clicking into a dozen different dashboards, downloading massive CSV files, and piecing together fragmented reports. This MCP changes that. It lets you manage security posture conversationally. You connect this tool through Vinkius, giving your AI agent immediate access to critical data about your network assets. Instead of manually exporting vulnerability findings or tracking risk scores across separate tools, you simply ask your client—Claude, Cursor, or any compatible agent—to find out what's wrong. Your agent acts like a dedicated Security Analyst: it can list all discovered issues, check the health of SSL certificates, and even kick off new security assessments on demand. It keeps you focused on risk mitigation, not report generation.

Core Capabilities

01 — Inventory Assets

List all monitored domains, IPs, and applications to map your entire digital footprint.

02 — Review Vulnerabilities

Access a list of discovered security issues, getting detailed information on severity and remediation status for any specific finding.

03 — Check Network Components

Examine network infrastructure by listing open ports, detecting technologies used, or reviewing SSL/TLS certificates.

04 — Assess Risk Over Time

Retrieve overall security risk scores and trends to measure how your organization's posture is changing month over month.

05 — Run Scans On Demand

Trigger new, immediate security assessments for any target you need to re-validate.

One Click on Vinkius — From Prompt to Execution

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

01 Subscribe to this MCP and enter your Halo Security API Key via Vinkius.

02 Connect your preferred AI client (Claude, Cursor, etc.) to the catalog.

03 Ask your agent a question—like 'What are the high-severity vulnerabilities on example.com?'—and get instant answers based on live security data.

The bottom line is you manage complex security tasks using simple conversation rather than complicated manual workflows.

Built For

This MCP is essential for Security Engineers and DevSecOps teams who spend too much time clicking through dashboards at 2 a.m. It gives them the power to query vast amounts of security data instantly, turning overwhelming reports into actionable dialogue.

Security Engineer

Needs to quickly pull asset lists and check vulnerability details during incident triage without opening multiple consoles.

DevSecOps Lead

Must automate the triggering of security scans and monitor findings in real-time as code gets deployed or infrastructure changes.

CISO

Requires a high-level, conversational overview of organizational risk scores and overall attack surface health without deep technical diving.

What Changes When You Connect

01 Stop digging through manual reports. Instead of exporting vulnerability data, you can ask your agent directly for 'all high-severity issues' using the list_issues tool, getting a clean summary immediately.

02 Map your full risk exposure by checking security trends with get_security_risk. You instantly see if recent changes have lowered or raised your organization's overall score.

03 Eliminate guesswork about network health. Use list_open_ports and list_certificates to check for exposed ports or expiring SSL/TLS credentials without running separate scripts.

04 Stay ahead of threats by adding new assets using add_target, ensuring that any newly acquired domain or IP address is immediately part of your security perimeter.

05 Get current status on all monitored systems. List targets and list technologies gives you a complete inventory—knowing what's running where before it becomes a vulnerability.

Real-World Applications

Post-Merger Asset Discovery

The M&A team needs to know if the newly acquired company's systems are secure. They ask their agent to list targets and then check for open ports, immediately identifying any unexpectedly exposed services or unmonitored IPs.

Pre-Deployment Security Check

A DevSecOps engineer is about to deploy a new API service. They use the agent to trigger_scan on the target and then list_issues, ensuring all known vulnerabilities are patched before going live.

Compliance Audit Preparation

The CISO needs quick proof of certificate compliance. They ask the agent to list_certificates, instantly verifying if every critical service has a current, non-expiring SSL/TLS credential across the board.

Investigating a Breach Alert

The security team gets an alert about unusual traffic. They ask the agent to list_dns_records and then list_technologies for the affected asset, narrowing down the potential entry vector faster than manual research.

Patterns to Avoid

Manual Report Chasing

✗ AVOID

Downloading weekly vulnerability reports from Dashboard A, cross-referencing them with IP lists from Console B, and then manually checking expiration dates in a third spreadsheet.

✓ INSTEAD

Instead, ask your agent to list_issues for the affected assets. Then run list_certificates if you need to check credentials. This centralizes all findings into one conversational flow.

Forgetting New Targets

✗ AVOID

A new department launches a website (new domain/IP) that isn't added to the monitoring tool, leaving it completely blind to attack.

✓ INSTEAD

Use add_target immediately when any new resource comes online. This ensures your agent includes the asset in all future list_issues and trigger_scan operations.

Scope Creep Confusion

✗ AVOID

Trying to figure out if a vulnerability is due to outdated software or poor network configuration by checking multiple separate tools.

✓ INSTEAD

Start with list_technologies to identify the specific software. Then use get_issue to understand exactly why that technology version poses a risk.

The Right Fit

Use this MCP if your primary pain point is synthesizing massive amounts of disparate security data—vulnerabilities, network topology, and compliance status—into actionable insights conversationally. You need the ability to ask natural language questions like 'What's wrong with my web assets?' and get a structured list of issues, port statuses, and risk scores in return.

Don't use this if your goal is pure, raw data ingestion for an internal database (you might just need a standard API call). Also, don't use it if you only care about one specific domain; this MCP excels when managing broad attack surfaces across multiple targets. If you only need to check DNS records, consider using a specialized network mapping tool instead, but remember that list_dns_records is useful for context.

The Dashboard Overload

Today, managing your security perimeter feels like juggling ten different dashboards. You're clicking into the vulnerability scanner to get a raw list of issues; switching tabs to check open ports; then opening a third tool just for SSL certificate status. You spend more time copying and pasting data than actually analyzing risk.

With this MCP, your AI agent handles the navigation. Instead of manually aggregating reports, you tell your client what you need—like listing all detected technologies or getting security trend scores—and it pulls everything together in one chat window. It's instant context.

Halo Security MCP: Get Clear Asset Visibility

Manually identifying every asset, port, and technology used across a growing infrastructure is slow and error-prone. You have to remember to check the domain list, then cross-reference it with the open ports list, then verify certificates separately.

This MCP lets you manage that entire lifecycle conversationally. By listing targets and triggering_scan on demand, your agent acts as an always-on security consultant, giving you a single, accurate view of what's vulnerable right now.

Halo Security: 11 Tools for Security Posture Management

These tools let you examine asset status, list open ports, check certificates, trigger scans, and retrieve detailed security findings across your monitored infrastructure.

#	TOOL	DESCRIPTION
01	add_target	Adds a specific domain or IP address to the list of assets being monitored for security issues.
02	get_issue	Retrieves full details about one particular security finding, including its severity and impact.
03	get_security_risk	Pulls the organization's overall risk score and historical trend data for quick comparison.
04	list_certificates	Displays a list of all SSL/TLS certificates found on monitored targets and their current expiration status.
05	list_dns_records	Shows the DNS records that were discovered during the initial reconnaissance phase for your assets.
06	list_issues	Lists all security vulnerabilities and issues found across all monitored targets at once.
07	list_open_ports	Displays every open network port discovered on the assets, helping map potential entry points.
08	list_scans	Provides a history log of all security scans that have been run against your targets.
09	list_targets	Retrieves the full list of assets currently included in the security monitoring perimeter.
10	list_technologies	Identifies and lists all specific software technologies or libraries running on your monitored targets.
11	trigger_scan	Initiates a brand new, immediate security assessment scan for a specified target asset.

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 issues with 'High' severity.

I've retrieved your issues. You currently have 3 High severity findings: 'Outdated jQuery Library' on example.com, 'Missing Security Headers', and 'SSL Certificate Expiring Soon'. Would you like the details for the SSL issue?

U Trigger a security scan for target ID 5592.

Scan triggered! I've successfully requested a new security assessment for target 5592 (api.example.com). You can check the progress by listing scans in a few minutes.

U Show me the risk score summary for our organization.

Accessing risk scores... Your organization's overall security score is 78/100. The score has improved by 5% since last month. Your primary risk drivers are outdated software on 2 targets. Would you like to see those targets?

Frequently Asked Questions

01 **How does the Halo Security MCP help with vulnerability reports?**

It moves beyond static reports. Instead of downloading CSV files, you ask your agent to list_issues or get_issue details directly in the chat. This gives you immediate context on severity and remediation steps.

02 **Can I find out if my domain is secure using Halo Security MCP?**

Yes, you can check multiple angles. You use list_certificates to verify SSL/TLS status and then run get_security_risk for an overall health score.

03 **What kind of assets can the add_target tool handle?**

The add_target function accepts domains, IP addresses, or application identifiers. You just need to tell your agent what needs monitoring.

04 **Does Halo Security MCP show me past scan results?**

Absolutely. Use list_scans to review the history of assessments run against your targets, and you can also manually trigger a new one using trigger_scan if needed.

05 **Is this MCP suitable for DevSecOps teams?**

Yes. It allows automated workflows where agents can monitor issues in real-time, ensuring that every change or deployment is checked against known vulnerabilities via list_issues.

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 `"halo-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

Halo 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 Halo 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	Halo Security MCP
Server ID	`019d75ad-b288-72e7-b83e-a949d090cd2e`
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/halo-security.