MCP SERVER | NO CODE | CLOUD HOSTED

Snyk MCP for AI Agents

Audit code vulnerabilities and security flaws directly in your development environment.

Snyk MCP connects your security dashboard to your AI agent. Analyze vulnerabilities, check project metadata, and manage organization permissions without leaving your workspace. It turns your AI into a security analyst that knows your codebase's specific risks.

A+ Quality Score 100/100

code-security | cve-scanning | devsecops | container-security | dependency-analysis | vulnerability-management

MAINTAINED BY VINKIUS CLOUD FOR AI AGENTS · JULY 2026
vinkius.com/mcp/snyk

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.

Snyk MCP

9 tools available
Cloud-hosted on Vinkius

You're tired of the constant context switching between your code and the Snyk web console. Every time a build fails because of a dependency issue, you have to leave your IDE, find the right project, and dig through CVE reports to see what actually needs fixing. This MCP changes that. You can just ask your AI agent to tell you which packages are broken and what the specific fix looks like. It pulls the data directly into your workspace. By using this through the Vinkius catalog, you get a direct line to your security data. You can check who has admin access, see how much of your monthly scan quota is left, or list every project in your organization. It moves the security check into the conversation where the work actually happens.

Core Capabilities

01 — Fetch specific CVE remediation steps

Get actionable instructions on how to patch a vulnerability without leaving your editor.

02 — Audit organizational membership roles

See every user and their permissions within your Snyk organization instantly.

03 — Query project configuration details

Retrieve internal configurations and metadata for any project in your account.

04 — Monitor API integration status

Check which external pipelines and tools are currently connected to your security dashboard.

05 — Check billing and usage limits

View your current scan limits and billing status via natural language.

One Click on Vinkius — From Prompt to Execution

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

01 Connect your Snyk account to the Vinkius platform.

02 Provide your personal Snyk API token to authorize the connection.

03 Start asking your agent about vulnerabilities, project status, or organization members.

The bottom line is you get your security data directly in your chat window.

Built For

The DevSecOps engineer who is tired of clicking through dashboards at 2am to find out why a container build failed.

DevSecOps Engineer

Investigates CVEs and generates threat models during PR reviews to block risky code from merging.

App Developer

Finds out which package versions triggered a build error without having to run local dependency scans.

SysAdmin

Audits user permissions and monitors API usage limits via text commands to keep the organization secure.

What Changes When You Connect

01 Stop hunting for flaws by getting a full list of active security issues in one view.

02 Get faster fixes by pulling remediation steps for specific CVEs without leaving your editor.

03 Audit your team permissions quickly by viewing every user and their role in your organization.

04 Manage your costs more effectively by checking your billing status and usage stats during planning.

05 Keep your dev environment organized by pulling project metadata and configurations on demand.

Real-World Applications

Fixing a failed build

A developer gets a build error and asks the agent to find the faulty dependency and its remediation steps.

Pre-merge security check

A DevSecOps engineer needs to know if a new PR introduces a critical flaw and asks the agent to check project details.

Quota monitoring

A SysAdmin needs to see how many seats are left and asks for usage stats to plan for next month.

User access audit

An auditor needs a list of all team members and their roles to ensure the principle of least privilege.

Patterns to Avoid

Manually copying CVE IDs

✗ AVOID

Copying a long CVE ID from a browser and pasting it into a chat to ask for help.

✓ INSTEAD

Ask the agent to pull the full list of active issues for the project first to see all flaws at once.

Guessing project IDs

✗ AVOID

Trying to query a project and getting an error because you don't know the exact ID.

✓ INSTEAD

Ask the agent to list all projects in the organization to identify the correct one first.

Manual billing checks

✗ AVOID

Leaving your terminal to check the Snyk dashboard for usage limits.

✓ INSTEAD

Ask the agent to check your billing info or usage stats directly in your current chat.

The Right Fit

Use this if you need to query your Snyk security data via natural language or want to automate security checks within your IDE. It's perfect for DevSecOps teams who want to see vulnerabilities without opening a browser. Don't use this if you need to perform deep static analysis on raw source code files or if you need to manually edit Snyk policies. Use a dedicated security scanner for those tasks.

Snyk MCP for Eliminating Manual CVE Research

You're stuck in a loop of opening the Snyk dashboard, searching for a project, filtering by severity, and then copying remediation steps into your notes. It's a tedious process that breaks your focus every time a vulnerability is flagged.

This MCP lets you stay in your editor. Just ask your agent to summarize the critical flaws and provide the fix paths. You get the answers you need without the tab switching.

Snyk MCP for Scaling Organization Security Audits

Auditing who has admin access or checking how much of your monthly scan quota is left usually requires a manual walk-through of the Snyk UI. It's easy to miss a permission or overlook a billing spike.

You can now audit your entire organization's membership and check your usage stats with a single prompt. It turns a chore into a quick status check.

Snyk MCP with 9 Tools for Code Security

Query your security dashboard, audit organization members, and check CVE remediation steps using natural language.

#	TOOL	DESCRIPTION
01	`get_project_details`	Get the specific configuration and metadata for a project. This helps you understand how a project is set up in Snyk.
02	`list_integrations`	See which external pipelines are currently connected to your organization. This is useful for auditing your security stack.
03	`list_issues`	Pull a list of all active security flaws for a specific project. You can use this to see exactly what needs fixing.
04	`list_organization_members`	View every user and their role within your Snyk organization. This makes it easy to manage team permissions.
05	`list_organizations`	Browse all different organizations you have access to. Use this to switch contexts between different company entities.
06	`list_projects`	Get a full list of projects within a specific organization. This helps you find the right project to scan.
07	`get_billing_info`	View current billing status and organization costs. This keeps you on top of your security spend.
08	`get_issue_details`	Retrieve the deep details and fix paths for a specific security flaw. It tells you exactly how to patch the code.
09	`get_usage_stats`	Check your current scan limits and usage metrics. This helps you monitor your team's security activity.

See It in Action

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

U What are the biggest security risks in our main app right now?

I've scanned your 'Core Backend' project and found **3 Critical Issues**:

- **SQL Injection** - Found in `db-driver@1.2.0`
- **Cross-Site Scripting** - Found in `auth-lib@0.9.4`
- **Insecure Deserialization** - Found in `json-parser@2.1.1`

Would you like me to pull the specific remediation steps for the SQL Injection flaw?

U Who has admin rights in our security org?

I've pulled the member list for your organization. Here are the users with **Admin** roles:

NAME	EMAIL	ROLE
Alex Mercer	a.mercer@company.co	`Admin`
Sarah Chen	s.chen@company.co	`Admin`

There are 2 other members with `Collaborator` or `Viewer` roles. Would you like to see the full list?

U How much of our Snyk budget is left?

Here is your current usage summary:

- **Monthly Scan Limit**: 50,000
- **Used This Month**: 12,450
- **Remaining**: 37,550

Your billing status is currently **Active**. You are well within your limits for the current cycle.

Frequently Asked Questions

01 **Can Snyk MCP help me find CVEs faster?**

Yes. It allows your AI agent to pull the full list of active security issues and CVEs for your projects directly into your chat window.

02 **How do I check my Snyk billing with an AI agent?**

You can ask your agent to check your billing status or usage stats. It will pull that data from your Snyk account and summarize it for you.

03 **Can I see who has access to our Snyk organization?**

Yes. Your agent can list all members of your Snyk organization and show you their specific roles, making it easy to audit permissions.

04 **Does Snyk MCP work with Cursor or Claude?**

It works with any MCP-compatible client, including Claude, Cursor, Windsurf, and VS Code.

05 **Can I see which projects are in my Snyk account?**

Yes. You can ask your agent to list all projects within a specific organization to help you find the right one to analyze.

06 **How do I get remediation steps for a specific flaw?**

Once the agent identifies a security issue, you can ask it for the specific remediation steps. It will pull the fix paths directly from Snyk.

07 **Can the AI give me the code fix for a Snyk security vulnerability?**

Yes! The bot uses `get_issue_details` to read Snyk's extensive remediation context natively. Because it operates inside your IDE (like Cursor), it seamlessly merges Snyk's advisory with your actual local file context to write a highly secure patch immediately.

08 **How do I find my organization ID if I only know my project name?**

You don't need to manually hunt for it. Simply tell your AI agent: 'Find my React Frontend project and list its issues'. The AI will autonomously query `list_organizations`, isolate the correct ID, run `list_projects` under it, find the matching name, and then execute the issue retrieval.

09 Is it safe to expose my project vulnerabilities to an AI?

Yes. Vinkius operates transparently—your Snyk API Token is securely isolated and requests route directly from your local MCP client to Snyk endpoint APIs. No underlying CVE issue is retained or spied upon on cloud databases you don't control.

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 `"snyk": { "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

Snyk 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 Snyk. 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	July 2026
MCP Server	Snyk MCP
Server ID	`019d760a-c2eb-712b-bade-fa9f33027e7a`
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/snyk.