MCP SERVER | NO CODE | CLOUD HOSTED

Acunetix 360 MCP for AI Agents

Automate web application vulnerability scanning and compliance auditing

Acunetix 360 connects your AI agent to automated web vulnerability scanning, giving you full control over application security directly from chat. Launch scans for your APIs and web apps, track detailed vulnerabilities by severity level, and audit scan progress without logging into a console. It handles everything from initial testing to compliance reporting.

F Quality Score 10.14/100

web-security | penetration-testing | automated-scanning | app-security | vulnerability-management | cybersecurity

MAINTAINED BY VINKIUS CLOUD FOR AI AGENTS · JUNE 2026
vinkius.com/mcp/acunetix-360

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.

Acunetix 360 MCP

3 tools available
Cloud-hosted on Vinkius

Running secure web applications involves constant monitoring and deep auditing—it's tedious manual work. This MCP lets you automate the entire security workflow using natural conversation. Instead of jumping between dashboards and writing complex API calls, your agent manages vulnerability scans for all your web apps and APIs on demand. You can start a new scan right in chat, check progress across multiple systems, or instantly pull up lists of identified vulnerabilities, including suggested fixes. The system supports auditing past builds to keep security checks running throughout the entire development lifecycle. By connecting this MCP through Vinkius, you give your AI client immediate access to robust, industry-standard web application testing tools.

Core Capabilities

01 — Launch new vulnerability scans

Start a full security scan on a specific web application or API endpoint directly from your chat interface.

02 — List and track active scans

Get the current status, progress, and historical record of all running vulnerability assessments across your infrastructure.

03 — Retrieve identified security issues

Pull a detailed list of every discovered flaw, including its severity level (e.g., Critical, High) and basic remediation guidance.

04 — Audit specific vulnerability reports

Quickly gather all recorded vulnerabilities to support compliance checks or risk assessments for management reporting.

One Click on Vinkius — From Prompt to Execution

Available at **vinkius.com/mcp/acunetix-360** — connect your AI agent in three steps.

01 Subscribe to this MCP and provide your Acunetix 360 User ID and API Token.

02 Connect it to your preferred AI client (like Cursor or Claude).

03 Ask your agent a natural language question, such as 'List all high-severity vulnerabilities found in my last scan.' The agent handles the rest.

The bottom line is that you talk to your agent like you're talking to a security analyst; it talks to Acunetix 360 and gives you the answer.

Built For

This MCP is for Security Engineers, DevSecOps teams, and Compliance Officers who are tired of manually switching between scanning tools, dashboard APIs, and ticketing systems. If your job involves verifying that a web app build meets strict security standards before launch, this tool saves hours of tedious clicking.

Security Engineers

You use this MCP to automate vulnerability triage, launching large-scale scans and monitoring the results across multiple applications simultaneously.

DevSecOps Teams

You integrate automated security checks into CI/CD pipelines and audit historical scan results using your AI client instead of writing complex scripts.

Compliance Officers

You retrieve comprehensive security reports and detailed vulnerability logs to easily pass risk assessments or meet regulatory requirements.

What Changes When You Connect

01 Start scans instantly. Instead of navigating the Acunetix console to launch a job, your agent handles it with one simple request.

02 Track progress easily. Use the `list_scans` tool to see the status of all ongoing security assessments without needing dashboard access.

03 Pinpoint flaws quickly. The `list_vulnerabilities` tool lets you pull up all identified issues—like SQL Injection or XSS—to focus remediation efforts immediately.

04 Audit builds efficiently. You can check scan results from recent development builds to ensure security standards never slip, even with fast releases.

05 Manage complexity. Your agent organizes the output, giving you actionable reports that go beyond just listing findings.

Real-World Applications

Pre-release vulnerability check

A developer needs to know if a new API endpoint is safe before going live. They ask their agent to launch_scan on the specific URL, and minutes later, they get a summary of critical flaws back in chat.

Compliance audit preparation

The Compliance Officer needs proof that all services were scanned last quarter. They use list_scans to retrieve a complete history report needed for the regulatory body.

Triage after an incident

A security engineer suspects a recent breach and wants to know exactly what flaws exist. They use list_vulnerabilities to get a categorized, prioritized rundown of all potential entry points.

Patterns to Avoid

Manual dashboard navigation

✗ AVOID

The user logs into the Acunetix portal, finds the 'Scan Management' tab, clicks 'New Scan,' and fills out multiple forms—a process that takes five minutes of clicking.

✓ INSTEAD

Instead, simply tell your agent: 'Launch a scan for my main checkout page.' The tool handles all the form filling and API calls instantly.

Copying vulnerability lists

✗ AVOID

The user runs a report and then has to manually copy hundreds of vulnerability findings into a separate spreadsheet for the team, risking data loss.

✓ INSTEAD

Ask your agent to list_vulnerabilities. It compiles the actionable data directly in chat format, making it ready for immediate review or pasting.

Ignoring scan status

✗ AVOID

The user starts a large scan and then forgets if it's running, stalled, or finished until they try to access the results hours later.

✓ INSTEAD

Check the status immediately by asking your agent to list_scans. You get real-time feedback on where all your security efforts stand.

The Right Fit

Use this MCP if you need automated, chat-based control over the entire web vulnerability scanning lifecycle. Specifically, if your workflow requires launching new scans (`launch_scan`), monitoring ongoing tests, and pulling detailed reports of flaws (`list_vulnerabilities`) without touching a UI, this is for you. Don't use it if all you need is to view an existing PDF report; in that case, a simple document reader tool suffices. However, if your team relies on proprietary ticketing systems for remediation tracking, you might want a dedicated issue tracker MCP instead.

Acunetix 360: Automating Web Vulnerability Scanning Workflows

Right now, managing web security is a mess of clicks. You have to log into the scanning platform, manually define scope rules for every new endpoint, hit 'Start Scan,' and then keep checking the status tab until it’s done. If you want to audit historical results or check progress across multiple applications, you end up copy-pasting data between different dashboards.

With this MCP, you just talk to your agent. You ask for a new scan, and it starts it. You need to know the status? Ask. The tool returns the necessary information immediately in chat. It takes massive manual overhead and turns it into simple conversation.

Acunetix 360: Streamlining AppSec Compliance Reporting

Before, generating a compliance report meant running the scan, waiting for completion, downloading a massive file, and then wading through pages of data just to extract the 'critical' findings. It was slow and prone to human error.

Now, you ask your agent to list_vulnerabilities. It filters out the noise and gives you only what matters—the actionable flaws needed for compliance sign-off. Security reporting is now instant.

Acunetix 360: 3 Tools for Web Vulnerability Scanning

Use these tools to start scans, list existing scan results, or pull comprehensive lists of all discovered security vulnerabilities.

#	TOOL	DESCRIPTION
01	`list_scans`	Retrieves a list of all past, present, and completed vulnerability scans within your Acunetix 360 account.
02	`launch_scan`	Initiates a brand new security scan on specified web applications or APIs.
03	`list_vulnerabilities`	Gets a comprehensive list of all identified security issues, including severity ratings and remediation advice.

See It in Action

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

U Can you show me all high-severity vulnerabilities found in my main login flow?

Vulnerability Report: Login Module

Cross-Site Scripting (XSS): *Found 2 instances on* `/login/profile` *.* ***Severity:*** *High.* Remediation:* Sanitize input fields.
SQL Injection: *Found 1 instance on* `user_id` *parameter.* ***Severity:*** *Critical.* Remediation:* Use prepared statements for all database queries.

(Total Flaws: 3. Two require immediate attention.)

U What's the status of my scan for the API gateway?

Scan Status Update:
The scan on `api/v1/gateway` started successfully at 2024-06-12 10:30 AM.
- **Current Status:** In Progress (78% complete)
- **Estimated Completion:** Within the next 25 minutes.

Please wait a moment, and I'll give you an alert when it finishes.

U List all scans we ran last month.

Scan History (Last 30 Days):

SCAN NAME	START DATE	STATUS	FLAWS FOUND
Main App v2.1	2024-05-28	Completed	12
API Beta Test	2024-06-01	Completed	5
Marketing Landing Page	2024-06-10	Canceled	N/A

Note: The 'Main App v2.1' scan contained the highest number of critical findings.

Frequently Asked Questions

01 **How do I use Acunetix 360 with my AI agent to check for security flaws?**

You simply ask your agent what you need. Instead of logging into the tool, just tell it, 'Check my application for critical vulnerabilities.' It uses its tools to run the necessary scans and provides the results directly in chat.

02 **Can Acunetix 360 help me audit security reports for compliance?**

Yes. You can retrieve comprehensive lists of identified issues using your agent. This makes generating documentation for audits much faster because you get structured data, not just a raw file download.

03 **If I launch a scan through the AI MCP, will it work on APIs?**

Yes, this tool is designed to scan more than just standard web pages. You can specify API endpoints and run full vulnerability scans against them, which is crucial for modern microservices architecture.

04 **What if I need a history of all my past security checks?**

You can ask the agent to list all previous scans. It pulls up a summary table showing when everything was run and how many flaws were found, helping you track long-term risk.

05 **Is Acunetix 360 MCP better than just using built-in IDE security tools?**

While IDE tools are great for code review, this MCP gives you a full platform view. It launches deep, automated scans against deployed applications and APIs, finding vulnerabilities that simple code analysis misses.

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 `"acunetix-360": { "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

Acunetix 360 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 Acunetix 360. 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	Acunetix 360 MCP
Server ID	`019d7546-591f-70da-9efd-bfb30235277e`
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/acunetix-360.