MCP SERVER | NO CODE | CLOUD HOSTED

Levo.ai Security MCP

Audit every API flow and data exposure instantly.

Levo.ai (API Security & Observability) MCP helps you audit and secure your APIs using natural conversation. It maps out every API endpoint—even undocumented ones—and flags vulnerabilities like BOLA or broken authentication. You can monitor for sensitive data (PII/PHI) exposure, generate live OpenAPI specs from actual traffic, and get detailed diagnostic evidence on security flaws.

A+ Quality Score 100/100

api-security | observability | vulnerability-scanning | pii-detection | openapi | threat-detection

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

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.

Levo.ai (API Security & Observability) MCP

10 tools available
Cloud-hosted on Vinkius

You run into a wall when trying to secure your APIs because the documentation is outdated, and the runtime environment is too complex. This MCP lets you hand off that complexity to your AI client. You stop manually sifting through millions of lines of logs or running separate compliance tools. Instead, you ask natural questions about your API structure and security posture.

Your agent can immediately list every single endpoint—whether it was documented years ago or if a developer just spun up a 'shadow' service last week. It checks those endpoints for sensitive data exposure, flagging anything containing PII or PHI. Need to know if an API is vulnerable? Your client runs checks against OWASP standards and gives you specific details on broken authentication instances. You can even get a live OpenAPI specification derived from actual observed traffic patterns; it's precise, not theoretical. This capability makes Levo.ai the ultimate security layer for your APIs, connecting directly to your operational data via Vinkius.

Core Capabilities

01 — Map all API endpoints

List every REST, GraphQL, gRPC, and SOAP endpoint, including any undocumented or unused shadow services.

02 — Audit sensitive data flows

Identify which APIs handle regulated data, like PII (names, emails) or PHI (medical records).

03 — Detect API vulnerabilities

Check for active security flaws against OWASP standards, such as broken object-level authorization.

04 — Generate live OpenAPI specs

Create accurate OpenAPI specifications based on the traffic your APIs are actually receiving right now.

05 — Analyze runtime behavior

Monitor API usage patterns and spot anomalies, like unexpected changes in data structure (schema drift).

06 — Retrieve vulnerability evidence

Get deep diagnostic reports explaining exactly how a specific security flaw was exploited.

One Click on Vinkius — From Prompt to Execution

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

01 First, subscribe to the Levo.ai MCP and input your API token and organization ID.

02 Next, tell your AI client what you need—for instance, 'List all applications that handle PHI.'

03 Your agent runs the query against Levo's live sensors and returns a clean list of endpoints, vulnerabilities, or data flows.

The bottom line is you get real-time security answers for your API stack without writing a single log query.

Built For

This MCP is built for the security engineer who can't afford to wait days for compliance reports, or the backend developer who needs instant proof that their code doesn't introduce new risks. If you deal with regulated data (HIPAA, GDPR), this saves you from massive headaches.

Security Engineer

You use it to hunt for API threats and monitor sensitive data exposure by asking your agent questions instead of manually filtering logs.

Backend Developer

You audit endpoint schemas and verify security build-time results before merging code that hits production.

Compliance Officer

You automate the auditing of regulated data flows, generating reports on your global API security posture across different environments.

What Changes When You Connect

01 Spot undocumented APIs: Use the `list_catalog_endpoints` tool to find 'shadow' or 'zombie' endpoints that nobody knows about, eliminating hidden security risks.

02 Ensure compliance effortlessly: The `list_sensitive_data` tool checks every endpoint for regulated data flows (PII/PHI), giving you instant audit reports.

03 Stop guessing on specs: Instead of writing OpenAPI definitions by hand, use `export_openapi_spec` to generate a specification based on real-time traffic observation. It's always accurate.

04 Deep dive into flaws: When a vulnerability is found, the `get_vulnerability` tool provides diagnostic evidence, telling you exactly what went wrong and how to fix it.

05 Catch behavioral drift: The `list_observations` tool tracks runtime changes in API traffic patterns. This alerts you when an endpoint's structure unexpectedly changes.

06 Understand scope quickly: Use `list_applications` and `list_environments` to map out exactly which services and deployment stages are currently under threat.

Real-World Applications

The compliance officer needs to prove PHI handling across all regions.

Instead of manually pulling reports from five different regional databases, the agent runs `list_sensitive_data` and filters results for 'PHI' exposure. It delivers a consolidated list of endpoints that need immediate policy review.

The developer suspects an old API is leaking data.

The developer asks the agent to check endpoint details using `get_endpoint_details` on a legacy service, confirming it's improperly exposing names and emails, leading to immediate remediation.

The security team needs an instant audit of all APIs.

The engineer runs `list_catalog_endpoints` to get a full inventory, then uses `list_vulnerabilities` to cross-reference the entire set for active OWASP flaws in one go.

A new microservice is deployed and needs immediate schema validation.

The team runs `export_openapi_spec` against the live service. The agent generates a verified, accurate OpenAPI payload that the documentation team can use immediately.

Patterns to Avoid

Manual log analysis for PII.

✗ AVOID

A junior security analyst spends all day filtering Splunk logs across multiple services just to see if a specific endpoint is passing emails, wasting hours and missing context.

✓ INSTEAD

Just ask your AI client to use `list_sensitive_data`. It checks every monitored API automatically and reports exactly where PII flows are happening. No manual log parsing required.

Relying on static documentation.

✗ AVOID

The team assumes a retired 'user profile' endpoint is safe because it's removed from the Wiki, but the code still runs and hasn't been secured or audited.

✓ INSTEAD

Use `list_catalog_endpoints` to find all endpoints, including undocumented shadow APIs. This prevents you from trusting outdated documentation.

Using generic vulnerability scanners.

✗ AVOID

Running a broad scanner that flags hundreds of issues but fails to explain *why* an object access flaw exists or how to fix it.

✓ INSTEAD

Run `get_vulnerability`. This tool provides deep diagnostic exploitation evidence, giving you the root cause and clear remediation steps for every issue.

The Right Fit

Use this MCP if your primary pain point is visibility into your API surface area. You need to know what APIs exist (even forgotten ones), where sensitive data goes, and if they are vulnerable—and you can't afford the time or resources for manual log review.

Don't use it if you only have a single, isolated application that runs in one known environment, and whose code is entirely under version control. In those cases, traditional static analysis tools might be sufficient. However, this MCP shines when your APIs are distributed across multiple applications, environments, or when the runtime behavior itself is the security concern.

The headache of API visibility today

Most companies deal with a sprawling web of microservices. To audit them, teams currently resort to a painful combination: manually checking Wikis for endpoint definitions, running expensive, slow scanners that miss shadow APIs, and then spending days correlating logs from dozens of different services just to find out where PII is flowing.

With this MCP, your agent takes over the detective work. You simply ask it about data exposure. It automatically searches every monitored service—regardless of whether a developer documented it or if it's running in staging or production—and gives you one clean answer.

You get complete API security context with Levo.ai

The tedious steps of compiling endpoint lists, checking schemas, and verifying data flows are gone. You no longer need to jump between `list_applications`, `list_catalog_endpoints`, and then a separate compliance dashboard.

You get one conversational layer that understands the difference between an intended API call and a dangerous security vulnerability. It’s immediate context.

Levo.ai (API Security & Observability) - 10 Tools

These tools give your agent granular control over every aspect of API auditing, from listing services to exporting live OpenAPI specifications.

#	TOOL	DESCRIPTION
01	`list_applications`	Lists all services and applications tracked by Levo.ai for monitoring.
02	`list_catalog_endpoints`	Identifies every API endpoint type—REST, GraphQL, gRPC, SOAP—including undocumented shadow APIs.
03	`get_endpoint_details`	Provides the complete schema structure for a specific discovered API endpoint.
04	`list_environments`	Lists all deployment boundaries that Levo's active monitoring sensors are covering.
05	`export_openapi_spec`	Automatically generates a live OpenAPI payload specification for any given application.
06	`get_observation`	Retrieves detailed information about a specific runtime anomaly detected by the sensors.
07	`get_vulnerability`	Pulls diagnostic exploitation evidence for a single, identified API vulnerability.
08	`list_observations`	Gathers a list of all runtime API behavior observations detected by the Levo sensors.
09	`list_sensitive_data`	Lists which API endpoints handle or expose sensitive, regulated data flows.
10	`list_vulnerabilities`	Gathers a comprehensive list of all active security vulnerabilities found across your applications.

See It in Action

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

U List all discovered API endpoints in our Levo catalog

I've found 24 total endpoints across 3 applications. Highlights include '/v1/users' (REST), '/graphql' (GraphQL), and 5 undocumented shadow APIs discovered in the 'Payments' service. Would you like to inspect the sensitive data status for these?

U Show me the active OWASP vulnerabilities for the 'Checkout' application

Retrieving vulnerabilities for 'Checkout'... I've identified 2 High impact issues: 1) BOLA (Broken Object Level Authorization) on '/api/v1/orders/{id}', 2) Broken User Authentication. I can provide the diagnostic exploitation evidence for both.

U Generate a live OpenAPI spec for the 'User Management' service

Generating OpenAPI spec based on live traffic... Done. I've successfully exported the specification for the User Management service. It includes 12 verified endpoints and mapped request/response schemas. Would you like the JSON payload?

Frequently Asked Questions

01 **How does Levo.ai (API Security & Observability) MCP find shadow APIs?**

The MCP uses the `list_catalog_endpoints` tool to dynamically map all traffic, not just documented routes. This means it finds 'shadow' or undocumented endpoints that are actively being used by your services.

02 **Is this better than traditional API gateway monitoring?**

Yes. While gateways monitor traffic flow, the Levo MCP analyzes *what* is in the traffic—specifically checking for PII/PHI and running deep OWASP vulnerability scans that go beyond simple rate limiting.

03 **What if I only need to check one endpoint's schema?**

You can use `get_endpoint_details` to pull the precise, detailed schema structure for any single API endpoint you discover in your catalog. It provides a deep dive into how that specific resource is built.

04 **Can Levo.ai (API Security & Observability) MCP help with compliance reporting?**

Absolutely. By listing sensitive data flows using `list_sensitive_data`, you automatically gather the evidence needed to prove regulatory adherence, simplifying your audit process.

05 **Does this tool support multiple environments (staging/prod)?**

Yes. You can use `list_environments` and then query specific data or vulnerabilities across those distinct deployment boundaries monitored by the sensors.

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 `"levoai-api-security-observability": { "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

Levo.ai (API Security & Observability) 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 Levo.ai (API Security & Observability). 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	Levo.ai (API Security & Observability) MCP
Server ID	`019d75c6-dc6c-71f7-a5a3-888c65e00720`
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/levoai-api-security-observability.