MCP SERVER | NO CODE | CLOUD HOSTED

UpGuard MCP

Assess Vendor Risk & Attack Surface Visibility

UpGuard monitors your entire attack surface and assesses third-party vendor risks through natural conversation. It lets you check security scores, track identity breaches affecting employees, and audit digital assets—all without jumping between dashboards. Connect this MCP to see exactly where vulnerabilities exist before attackers do.

A+ Quality Score 100/100

attack-surface | vendor-risk | cybersecurity | compliance | security-scanning | risk-assessment

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

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.

UpGuard MCP

9 tools available
Cloud-hosted on Vinkius

Monitoring an organization's digital perimeter is a full-time job that used to require running reports across five different consoles. This MCP connects your AI agent directly to UpGuard data, letting you talk through complex security questions like talking to a seasoned analyst. You can ask about specific vendors or track down every instance of identity theft affecting your staff. If you're using Vinkius, this connector pulls together vendor risk profiles, monitored domains, and active account risks into one chat window. It’s simple: instead of building complex queries, you just tell the AI what you need to know about who you trust or where your data might be exposed.

Core Capabilities

01 — Audit Vendor Security Scores

Retrieve security scores and detailed risk metrics for any third-party vendor you work with.

02 — Map Digital Assets

List all monitored domains, IP ranges, and SaaS applications to understand your full digital footprint.

03 — Track Identity Theft Incidents

View recent identity breaches affecting your workforce and get reports on affected email addresses.

04 — Assess User Risk Profiles

Audit user-specific risk data, checking for signs of compromised accounts or behavioral issues.

05 — Identify Active Infrastructure Risks

List all active security risks found across your own infrastructure and your vendor network.

One Click on Vinkius — From Prompt to Execution

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

01 Subscribe to this MCP and enter your UpGuard API Key into the Vinkius catalog.

02 Your AI client authenticates the connection, giving your agent access to all monitored security data.

03 You prompt the agent with a natural language query—like 'Show me the top three vendors with high-risk scores'—and get immediate, summarized results.

The bottom line is you ask a question in plain English and get actionable, data-backed security answers instantly.

Built For

This MCP is for the Security Analyst who spends half their day cross-referencing spreadsheets; for the Compliance Officer needing proof of due diligence; and for the CIO who needs a single source of truth on vendor exposure. It cuts out hours of manual report generation.

Security Analyst

Uses this MCP to check active security risks across both internal infrastructure and external vendors, focusing on specific tools like list_user_risks.

Compliance Officer

Verifies vendor risk profiles against compliance standards and tracks identity breaches to satisfy regulatory audits.

IT Operations Manager

Maintains a complete, real-time inventory of all monitored domains and IPs by calling list_monitored_domains or list_monitored_ips.

What Changes When You Connect

01 Stop manually checking security reports. You can run list_vendors and immediately get a full overview of every monitored vendor, including their current score.

02 Audit user activity risk instantly. Using list_user_risks lets you pinpoint which employees are exposed to identity theft or suspicious behavior without needing the HR team's help.

03 Keep track of your digital footprint by listing all assets with list_monitored_domains and list_monitored_ips in one query, eliminating spreadsheet sprawl.

04 Determine exactly what's wrong with a partner. You can check active issues for any third party using list_vendor_risks, then narrow it down with get_vendor.

05 Respond to breaches faster. Calling list_identity_breaches gives you immediate access to breach reports and affected employee lists when an incident happens.

Real-World Applications

Vendor Due Diligence Check

A compliance officer needs to prove that a new partner, 'Acme Corp,' meets security standards. They ask the agent: 'List active risks for Acme Corp and list all monitored vendors.' The AI responds with specific findings from list_vendor_risks and then gives a list of competitors using list_vendors.

Identifying Internal Exposure

An IT Ops Manager suspects an employee's credentials were stolen. They prompt: 'What is the current risk status for user Jane Doe?' The agent runs list_user_risks, providing a clear report on identity theft exposure and recommending immediate action.

Mapping Forgotten Assets

The security team needs to know every public asset they manage. They run 'List all domains and IPs.' The agent uses list_monitored_domains and list_monitored_ips, giving them a definitive inventory of the organization's digital perimeter.

Post-Incident Review

After an incident, the security team needs to know how many employees were affected. They ask: 'Were there any identity breaches in Q3?' The agent uses list_identity_breaches and provides a detailed report on the scale and source of the breach.

Patterns to Avoid

Trying to query historical data manually

✗ AVOID

Downloading raw CSV reports from UpGuard's website for vendor scores, then trying to compare those scores with a separate list of monitored domains in Excel.

✓ INSTEAD

Instead, use the MCP. You can ask the agent to check multiple sources at once: 'Show me vendors whose security score dropped and list their associated domains.' This combines data points from list_vendors, get_vendor, and list_monitored_domains into one chat response.

Running risk checks in a siloed dashboard

✗ AVOID

Checking user risks on the Identity page, then checking vendor risks on the Vendor portal. You end up missing the connection between them.

✓ INSTEAD

Use list_user_risks and list_vendor_risks together. Ask: 'Are there any high-risk users associated with vendors who show active security risks?' The MCP correlates these two data sets for a holistic view.

Forgetting the full scope of assets

✗ AVOID

Only looking at public domains and ignoring internal IP ranges or SaaS apps that could be compromised.

✓ INSTEAD

Always include asset visibility checks. Use list_monitored_ips, list_monitored_domains, and list_saas_apps in a single prompt to guarantee you've covered your entire attack surface.

The Right Fit

Use this MCP if your core problem is connecting disparate security data points: vendor risk, domain inventory, user behavior, and identity breaches. You need an agent that can correlate findings from list_vendors, get_vendor, and list_user_risks in a single conversation thread. Don't use it if you are building a real-time SIEM feed or needing to ingest massive amounts of raw log data; for that, stick with dedicated logging platforms. However, if your workflow involves asking 'What is the security status of X?' across multiple domains (vendors, users, assets), this MCP cuts out all the manual clicking and report generation.

Security review used to be a nightmare of tabs and PDFs.

Today, checking your company's security posture means opening five different portals: one for vendors, one for user logins, one for IP ranges, and another for SaaS apps. You spend hours copy-pasting scores into spreadsheets just to figure out which partner is actually the biggest risk.

With this MCP, you talk to your agent. You ask about a vendor's security score or check for identity breaches affecting employees, and it pulls all that information together instantly. You get actionable intelligence without ever leaving your chat window.

UpGuard: Security Visibility in Conversation

The manual steps of querying list_vendors, checking the specific details with get_vendor, and then cross-referencing those findings against your asset inventory using list_monitored_domains are now one prompt. You don't click; you ask.

Your security team moves from data gathering to risk remediation. The MCP turns weeks of report aggregation into minutes of conversation.

UpGuard With 9 Tools

These tools let you query specific data points, from listing all vendors to checking active user risk profiles, giving you complete visibility into your attack surface.

#	TOOL	DESCRIPTION
01	`get_vendor`	Retrieves detailed security information for a single, specified vendor.
02	`list_saas_apps`	Provides a list of all SaaS applications monitored by UpGuard.
03	`list_identity_breaches`	Generates a full report detailing identity breaches found across the organization's workforce.
04	`list_monitored_domains`	Lists all domain names that are currently under continuous monitoring.
05	`list_monitored_ips`	Provides a comprehensive list of monitored IP addresses for asset visibility.
06	`list_account_risks`	Lists all currently detected security risks affecting the entire account or organization.
07	`list_user_risks`	Retrieves specific risk data and threat indicators for individual users.
08	`list_vendor_risks`	Lists all active security risks found specifically within a target vendor's profile.
09	`list_vendors`	Retrieves an overview list of every third-party vendor currently being monitored.

See It in Action

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

U List all monitored vendors and their security scores.

I've retrieved your vendor list. You are monitoring 15 vendors. Top scorers include 'Cloud Analytics' (950/950) and 'Payment Gateway X' (920/950). 'Old Legacy Supplier' has a lower score of 450. Which one would you like a detailed risk report for?

U Show me the active risks for the vendor 'Microsoft'.

I've fetched the risks for Microsoft. There are 2 active findings: 'Unencrypted backup found' and 'TLS 1.0 support detected'. Shall I retrieve the remediation steps for these?

U Are there any recent identity breaches affecting our domain?

Inspecting breach data... I found one recent identity breach affecting 3 employees. It occurred on a third-party site 'TrainingPortal.com'. Would you like the list of affected email addresses?

Frequently Asked Questions

01 **How do I check a vendor's score using UpGuard MCP?**

You can use list_vendors to see an overview, or get_vendor to pull deep security details for a single partner. The agent presents this data in plain English so you don't have to read technical reports.

02 **Can UpGuard MCP track my employees' personal breaches?**

Yes. You can use list_identity_breaches to monitor identity theft affecting your workforce, giving you immediate alerts on compromised credentials or domains.

03 **Does this MCP show me all my assets?**

It provides comprehensive visibility by letting you list_monitored_domains and list_monitored_ips. This ensures your entire digital footprint is accounted for in one place.

04 **How do I check if a user account is risky with UpGuard MCP?**

Use the list_user_risks tool. It aggregates behavioral data to show you specific risks associated with individual users, helping you preemptively address compromised accounts.

05 **Is this better than just looking at vendor reports?**

Yes. While vendor reports are useful, the MCP allows you to run list_vendor_risks and compare those findings against your own monitored IP ranges (list_monitored_ips) in a single, cross-referenced view.

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 `"upguard": { "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

UpGuard 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 UpGuard. 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	UpGuard MCP
Server ID	`019dd17d-cb81-706e-a890-a1594e5fb815`
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/upguard.