MCP SERVER | NO CODE | CLOUD HOSTED

AI Cyber Risk Assessment Engine.

Turn technical AI vulnerabilities into clear financial risk data.

AI Cyber Risk Assessment Engine gives your AI client the ability to measure the security posture of your machine learning models. It calculates risk scores, estimates potential financial losses in Euros, and identifies specific attack vectors like prompt injection. You can use it to plan security budgets by getting direct recommendations on where to invest your mitigation funds.

A+ Quality Score 100/100

ai-security | cyber-risk | risk-quantification | model-vulnerability | security-investment

MAINTAINED BY VINKIUS CLOUD FOR AI AGENTS · SEPTEMBER 2026
vinkius.com/en/ai-agent-connect/ai-cyber-risk-assessment-engine

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.

AI Cyber Risk Assessment Engine MCP

4 tools available
Cloud-hosted on Vinkius

You can now give your AI agent the ability to audit the security of your AI systems. This MCP acts as a specialized risk engine that translates technical vulnerabilities into numbers you can actually use for business decisions. Instead of guessing how vulnerable a model is to prompt injection or model inversion, you get a quantified risk score.

If you are managing a deployment, you can use this to see how much a specific vulnerability might cost the company in Euros. It moves beyond vague security warnings by providing a concrete estimate of loss exposure. When it comes to budgeting, the MCP helps you decide where to put your money by recommending specific security investments based on the identified threats. It's a direct way to bridge the gap between technical model flaws and the financial reality of running AI in production.

Core Capabilities

01 — Risk Quantification

Your agent uses this to turn model vulnerabilities into a standardized risk score.

02 — Financial Impact Modeling

Your agent calculates the potential Euro-based loss exposure for a given risk level.

03 — Vulnerability Mapping

Your agent identifies specific attack paths like model inversion or prompt injection.

04 — Budget Planning

Your agent suggests specific investment amounts to lower your risk score.

05 — Security Control Evaluation

Your agent assesses how effective current controls are against specific AI threats.

One Click on Vinkius — From Prompt to Execution

Available at **vinkius.com/en/ai-agent-connect/ai-cyber-risk-assessment-engine** — connect your AI agent in three steps.

01 Connect the MCP to your client like Claude or Cursor via Vinkius.

02 Provide your agent with details about your AI system and its current controls.

03 The agent calls the risk engine to calculate scores and exposure.

04 Review the financial impact and investment recommendations generated by the agent.

Connecting this MCP to your AI client gives your agent immediate access to specialized security math.

Built For

This MCP is built for professionals managing the intersection of AI deployment and corporate security.

AI Security Engineers

They use the tool to identify specific attack vectors and validate mitigation strategies.

CISO and Risk Officers

They use the financial exposure data to report risk to the board in Euro terms.

Product Managers

They use the investment recommendations to justify security spend in their product roadmaps.

What Changes When You Connect

01 Converts technical AI flaws into Euro-based financial estimates.

02 Identifies specific attack paths like prompt injection.

03	Provides data-driven budget recommendations for security mitigation.
04	Generates a single risk score for quick high-level assessments.

Real-World Applications

Pre-deployment Audits

Run a risk assessment on a new model before it goes live to find critical vulnerabilities.

Security Budgeting

Calculate exactly how much money is needed to move a risk score from a critical level to a manageable one.

Financial Risk Reporting

Translate a high risk score into a potential loss amount for executive reporting.

Vulnerability Deep Dives

Use attack vector analysis to understand if your model is specifically prone to inversion or injection.

Patterns to Avoid

Treating risk as a vague concept

AVOID

Asking your agent to tell you if your model is 'generally safe' without any specific data points.

✓ INSTEAD

Use `calculate_risk_score` to get a concrete number that represents your actual threat level.

Ignoring the financial cost of flaws

AVOID

Focusing only on the technical details of a prompt injection without knowing the business impact.

✓ INSTEAD

Run `estimate_loss_exposure` to see exactly how much that specific vulnerability could cost the company in Euros.

Guessing security budget requirements

AVOID

Setting a flat security budget for the next quarter without looking at specific model threats.

✓ INSTEAD

Use `recommend_security_investment` to find out the specific funding needed to lower your risk score.

The Right Fit

Security engineers and CISOs should connect this MCP to bridge the gap between technical flaws and financial reporting. Product managers should use it to justify security spend in their roadmaps. If you only need to check basic model uptime or general performance, this isn't for you.

4 Tools for Quantifying AI Cyber Risk

These tools turn technical model vulnerabilities into actionable financial data and budget plans.

#	TOOL	DESCRIPTION
01	`analyze_attack_vectors`	This tool identifies which specific AI vulnerabilities are driving your current risk levels. It pinpoints critical paths like prompt injection or model inversion based on your model type.
02	`calculate_risk_score`	This tool provides a high level assessment of your current cyber risk. It gives your agent a single number to represent the overall threat level.
03	`estimate_loss_exposure`	This tool translates technical security risks into potential financial impact. It calculates expected losses in Euros to help you understand the stakes.
04	`recommend_security_investment`	This tool determines the specific funding needed to mitigate your identified risks. It helps you plan your security budget around actual needs.

See It in Action

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

U What is the current cyber risk score for an LLM with a high attack surface and low control effectiveness?

The calculated risk score is 8.5, which is classified as Critical, primarily driven by the high attack surface and insufficient security controls.

U Estimate the potential loss for an AI system worth €500,000 with a risk score of 7.

The estimated potential loss is €350,000, with a loss range of €250,000 to €450,000 and High confidence.

U How much should I invest to reduce my risk score from 8 to 3 for a system valued at €1,000,000?

The recommended budget is €120,000, with an Immediate priority level, focusing on Data Privacy Enhancement and Input Filtering.

Frequently Asked Questions

01 **What kind of AI risks does this MCP cover?**

It focuses on AI-specific threats, including prompt injection, model inversion, and data privacy risks.

02 **How does it calculate financial loss?**

The tool translates technical risk scores and system value into a potential loss exposure expressed in Euros.

03 **Can I use this with Claude or Cursor?**

Yes, you can connect this MCP to any MCP-compatible client including Claude, Cursor, and Windsurf.

04 **Does this MCP provide specific budget numbers?**

Yes, the recommend_security_investment tool provides specific funding amounts needed to mitigate identified risks.

05 **How is the risk score determined?**

The score is a high level assessment based on the model's attack surface and the effectiveness of your security controls.

06 **How is the cyber risk score calculated?**

The `calculate_risk_score` tool determines the score by analyzing the intersection of the attack surface, model vulnerabilities, and data sensitivity, then adjusting for the effectiveness of existing security controls.

07 **Can I estimate the financial impact of an AI breach?**

Yes, you can use `estimate_loss_exposure` to calculate the potential financial loss in Euros by providing the current risk score, the asset value, and the predicted incident frequency.

08 **What kind of AI vulnerabilities can be analyzed?**

By using `analyze_attack_vectors`, you can identify critical paths such as prompt injection, adversarial evasion, or model inversion for various model types like LLMs or Diffusion models.

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 `"ai-cyber-risk-assessment-engine": { "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

AI Cyber Risk Assessment Engine 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 AI Cyber Risk Assessment Engine. 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	September 2026
MCP Server	AI Cyber Risk Assessment Engine MCP
Server ID	`01a09d56-a36c-7181-9a98-767475d39483`
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/en/ai-agent-connect/ai-cyber-risk-assessment-engine.