MCP SERVER | NO CODE | CLOUD HOSTED

Bow-Tie Risk Analysis Engine.

Model complex risk paths and barrier effectiveness with your agent.

Bow-Tie Risk Analysis Engine gives your AI client a mathematical framework to map the relationship between hazards, threats, and consequences. Instead of guessing, your agent uses quantitative data to evaluate how preventive and mitigative barriers actually impact risk likelihood and severity. It handles everything from barrier degradation to total risk reduction profiles.

A+ Quality Score 100/100

risk | safety | bow-tie | quantitative | assessment

MAINTAINED BY VINKIUS CLOUD FOR AI AGENTS · SEPTEMBER 2026
vinkius.com/en/ai-agent-connect/bow-tie-risk-analysis-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.

Bow-Tie Risk Analysis Engine MCP

4 tools available
Cloud-hosted on Vinkius

You can now give your AI agent a rigorous way to handle safety and risk modeling. This MCP provides the mathematical backbone needed to run bow-tie analyses, moving beyond simple descriptions into actual quantitative assessment. You use it to map out how a specific threat leads to a top event and how subsequent consequences unfold.

Your agent can look at the specific barriers you have in place to see if they actually work. It doesn't just assume a barrier is perfect; it accounts for degradation to give you a realistic view of your safety posture. By running these models, you can see exactly how much risk you are actually reducing and where your gaps live. It turns your AI client into a specialized risk modeling partner that understands the mechanics of prevention and mitigation.

Core Capabilities

01 — Threat Path Modeling

Your agent calculates how preventive barriers reduce the likelihood of a threat reaching a top event.

02 — Consequence Assessment

Your agent evaluates how mitigative barriers limit the severity of an event once it occurs.

03 — Barrier Health Monitoring

Your agent checks for barrier degradation to adjust risk calculations based on real capability.

04 — Risk Reduction Summaries

Your agent produces a high level view of the total risk reduction achieved by your current safety measures.

One Click on Vinkius — From Prompt to Execution

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

01 Connect your AI client to the Vinkius catalog with one click.

02 Provide your agent with the specific threat, top event, or barrier data.

03 The agent invokes the necessary tools to run the mathematical calculations.

04 Receive a quantitative breakdown of threat likelihood, consequence severity, or barrier health.

Get your risk modeling running in minutes by connecting your preferred AI client to the Vinkius hosted MCP.

Built For

This MCP is built for professionals managing high-stakes safety and operational risks who need mathematical rigor in their modeling.

Risk Engineers

You hand off complex threat and consequence data to the agent to run quantitative bow-tie models.

Safety Managers

You use the agent to monitor barrier health and verify that mitigation strategies are meeting targets.

Process Safety Specialists

You task the agent with calculating risk reduction profiles for specific industrial hazards.

What Changes When You Connect

01 Quantifies risk likelihood and severity through mathematical modeling.

02 Accounts for real world barrier degradation rather than assuming perfect performance.

03 Provides clear summaries of total risk reduction across entire event paths.

04 Connects directly to your existing AI workflow in Claude or Cursor.

Real-World Applications

Safety Barrier Verification

Use the agent to check if your current preventive measures are enough to block a specific threat path.

Degradation Impact Analysis

Run health checks on specific barriers to see how much they increase your overall risk profile.

Risk Mitigation Planning

Model different sets of mitigative barriers to see which ones provide the best risk reduction profile.

Audit and Compliance

Generate quantitative summaries of risk reduction to support safety documentation and reviews.

Patterns to Avoid

Assuming perfect barrier performance

AVOID

You tell your agent that a safety valve is 100% effective and use that to justify a low risk score.

✓ INSTEAD

Use `evaluate_barrier_health` to account for degradation and get a realistic view of how much a barrier has actually weakened.

Vague risk descriptions

AVOID

You ask your agent to describe why a chemical leak is dangerous without any quantitative data.

✓ INSTEAD

Use `analyze_consequence_path` to factor in mitigative barriers and determine the actual severity of the outcome.

Ignoring the threat-to-event link

AVOID

You focus only on the consequences of an accident while ignoring how likely the initial threat is to occur.

✓ INSTEAD

Run `analyze_threat_path` to see how preventive barriers influence the likelihood of a threat reaching a top event.

The Right Fit

Connect this MCP if you manage high-stakes industrial or operational risks that require mathematical modeling rather than just descriptive lists. If your safety process relies on qualitative guesswork rather than quantitative assessment, this tool isn't for you. The deciding factor is whether you need to account for real-world barrier degradation in your risk calculations.

Bow-Tie Risk Analysis Engine 4 Tools

These tools provide the mathematical framework to model threat paths, assess consequences, and monitor the real-world health of your safety barriers.

#	TOOL	DESCRIPTION
01	`analyze_threat_path`	This tool evaluates the link between a specific threat and a top event. It calculates how preventive barriers influence the likelihood of that threat occurring.
02	`calculate_risk_reduction_profile`	This tool generates a high level summary of your risk posture. It shows exactly how much total risk is removed by your current set of barriers.
03	`analyze_consequence_path`	This tool examines the relationship between a top event and its resulting consequences. It factors in mitigative barriers to determine the severity of the outcome.
04	`evaluate_barrier_health`	This tool checks the actual capability of a specific barrier. It factors in degradation to show if a barrier is still performing at its nominal level.

See It in Action

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

U Analyze the threat path for threat 'T-101' leading to top event 'TE-202' using preventive barriers 'B-01' and 'B-02'.

The threat likelihood for T-101 is 0.8. After applying barriers B-01 and B-02, the residual threat likelihood is reduced to 0.15, and the path is successfully blocked.

U What is the current health status of barrier 'B-500'?

Barrier B-500 has a nominal effectiveness of 0.95, but due to degradation, its current effectiveness is 0.72. The status is currently 'Degraded'.

U Calculate the risk reduction profile for top event 'TE-99'.

For top event TE-99, the unmitigated risk score is 100.0 and the mitigated risk score is 12.5, resulting in a total reduction percentage of 87.5%. No critical gaps were identified.

Frequently Asked Questions

01 **What is bow-tie analysis?**

It is a risk assessment method that models the relationship between a hazard, the threats that lead to an event, and the consequences that follow, using barriers to manage the risk.

02 **How does this MCP handle barrier degradation?**

The evaluate_barrier_health tool specifically looks at degradation factors to adjust the nominal effectiveness of a barrier to its current, real world capability.

03 **Which AI clients can I use with this MCP?**

You can use this MCP with any compatible client, including Claude, Cursor, Windsurf, and VS Code.

04 **Do I need to host the MCP myself?**

No, Vinkius hosts and manages the MCP for you. You just connect your client and start using the tools.

05 **Can the agent tell me if my risk reduction is sufficient?**

Yes, by using the calculate_risk_reduction_profile tool, your agent can provide a high level summary of how much risk your current barriers are actually reducing.

06 **How does the engine handle barrier degradation?**

The engine uses `evaluate_barrier_health` to assess how degradation factors reduce a barrier's nominal effectiveness, providing a real-time view of current capability.

07 **Can I see the total risk reduction for a specific event?**

Yes, you can use `calculate_risk_reduction_profile` to get a summary of unmitigated vs mitigated risk scores and identify critical gaps.

08 **What is the difference between preventive and mitigative barriers?**

Preventive barriers act before a top event to stop threats, while mitigative barriers act after a top event to reduce the impact of consequences.

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 `"bow-tie-risk-analysis-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

Bow-Tie Risk Analysis 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 Bow-Tie Risk Analysis 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	Bow-Tie Risk Analysis Engine MCP
Server ID	`01a092f1-70ec-7370-b5cb-d53a177ea565`
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/bow-tie-risk-analysis-engine.