MCP SERVER | NO CODE | CLOUD HOSTED

AI ESG Scoring Engine MCP.

Quantify the ethical and social impact of your AI systems.

AI ESG Scoring Engine provides your AI client with the tools to measure Environmental, Social, and Governance impacts. It calculates composite responsibility scores, identifies compliance gaps, and simulates the capital needed to hit specific ESG milestones. You can use it to audit bias, check organizational oversight, and plan for future regulatory requirements.

A+ Quality Score 100/100

esg | ai-ethics | compliance | governance | risk-management

MAINTAINED BY VINKIUS CLOUD FOR AI AGENTS · SEPTEMBER 2026
vinkius.com/en/ai-agent-connect/ai-esg-scoring-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 ESG Scoring Engine MCP

4 tools available
Cloud-hosted on Vinkius

You can now give your AI agent the ability to audit its own ethical footprint. This MCP provides a specialized scoring engine designed to evaluate how AI systems impact Environmental, Social, and Governance (ESG) metrics. Instead of guessing at compliance, you use this tool to get hard numbers on responsibility scores and specific gaps in your current framework.

You can dig into the social side of things by looking at bias risks or check if your company's structure is actually ready to handle AI-related risks. It also handles the financial side of ethics. If you have a target ESG score in mind, you can run simulations to see exactly how much investment is required to get there. It turns abstract ethical concepts into measurable data points that your agent can use to guide your development or compliance strategy.

Core Capabilities

01 — ESG Scoring

Your agent calculates a composite responsibility score based on your specific AI metrics.

02 — Bias Assessment

The AI analyzes social risk and demographic parity to find hidden biases.

03 — Governance Auditing

Your agent evaluates if your organizational oversight is strong enough to manage AI risks.

04 — Financial Projection

The tool simulates the capital needed to reach your target ESG milestones.

05 — Compliance Gap Analysis

Your agent identifies specific areas where your AI fails to meet ESG standards.

One Click on Vinkius — From Prompt to Execution

Available at **vinkius.com/en/ai-agent-connect/ai-esg-scoring-engine** — connect your AI agent in three steps.

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

02 Provide your AI agent with your current AI system metrics.

03 The agent calls the specific tool needed, such as calculate_ai_esg_score.

04 Your agent receives the score, gaps, or investment projections.

05 Use the results to adjust your AI development or compliance strategy.

Connecting this MCP to your AI client gives your agent immediate access to specialized ESG calculation tools.

Built For

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

AI Compliance Officers

They use the tool to identify regulatory gaps and document ESG scores for audits.

ESG Analysts

They use the scoring engine to quantify the social and governance impact of new technologies.

Risk Managers

They use the governance and bias tools to assess potential social and organizational liabilities.

What Changes When You Connect

01 Turns qualitative ethical concerns into quantitative scores.

02 Provides a clear financial roadmap for reaching ESG targets.

03 Identifies specific social risks like demographic parity gaps.

04 Automates the assessment of organizational governance readiness.

Real-World Applications

Pre-deployment Audits

Run a full ESG score calculation before launching a new AI model to ensure it meets company standards.

Budget Planning

Simulate the investment needed to improve a low ESG score to a target level.

Bias Mitigation

Use the social risk tool to find and fix demographic disparities in your training data.

Governance Reviews

Check if your current management structure is prepared for the risks of your AI deployment.

Patterns to Avoid

Treating ethics as a guess

AVOID

You ask your agent to tell you if your new model is 'mostly ethical' without any data. This leads to vague answers that won't pass a real audit.

✓ INSTEAD

Use calculate_ai_esg_score to get a specific score from 0 to 100 along with the exact financial and compliance metrics tied to it.

Ignoring the cost of compliance

AVOID

You decide to improve your ESG score but have no idea how much capital to set aside for the necessary changes.

✓ INSTEAD

Run simulations with project_investment_needs to estimate the specific capital required to hit your target ESG milestones.

Overlooking organizational gaps

AVOID

You assume your current management team can handle the risks of a new AI deployment without checking their readiness.

✓ INSTEAD

Use assess_governance_readiness to verify if your current organizational structure can actually manage the risks of the technology you are deploying.

The Right Fit

Connect this MCP if you are an AI compliance officer, ESG analyst, or risk manager who needs to quantify ethical impact. Do not connect it if you are looking for general sentiment analysis or qualitative feedback. The deciding factor is whether you need hard, measurable data to back up your corporate responsibility strategy.

AI ESG Scoring Engine: 4 Specialized Audit Tools

These tools turn abstract ethical concerns into hard numbers. Your agent uses them to calculate scores, find bias, and plan the budget needed for compliance.

#	TOOL	DESCRIPTION
01	`assess_governance_readiness`	This tool checks if your current organizational structure can actually manage AI risks. It evaluates whether your oversight is sufficient for the technology you are deploying.
02	`calculate_ai_esg_score`	This tool generates a primary ESG score from 0 to 100. It also outputs the specific financial and compliance metrics tied to that score.
03	`evaluate_bias_risk`	This tool performs a deep dive into the social risk components of your AI. It focuses on identifying potential social impacts and demographic disparities.
04	`project_investment_needs`	This tool runs simulations to estimate the capital required to reach specific ESG milestones. It helps you plan the budget needed to improve your scores.

See It in Action

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

U Calculate the AI ESG score for a system with a bias mitigation level of 0.8, transparency rating of 0.7, governance strength of 0.9, and stakeholder impact of 0.8.

The calculated AI ESG score is 82. The system shows high governance stability but has minor compliance gaps in transparency documentation.

U What is the social risk if my bias mitigation level is 0.5 and the demographic parity gap is 0.3?

The social risk score is 45. Mitigation urgency is Scheduled. It is recommended to increase training data diversity to reduce the demographic parity gap.

U How much investment is needed to move my current score of 65 to a target of 85?

The estimated required investment is €45,000 to achieve a projected score increase of 20 points.

Frequently Asked Questions

01 **What can this MCP do for AI ethics?**

It allows your AI agent to calculate ESG scores, evaluate bias risks, and check if your organization is ready to manage AI governance.

02 **How is the ESG score calculated?**

The score is a composite value from 0 to 100 based on metrics like bias mitigation, transparency, governance, and stakeholder impact.

03 **Can I use this to plan my ESG budget?**

Yes. The project_investment_needs tool simulates the capital required to reach specific ESG target milestones.

04 **Which AI clients work with this MCP?**

Any MCP-compatible client, including Claude, Cursor, Windsurf, and VS Code, can use these tools.

05 **Does this MCP help with bias detection?**

Yes. The evaluate_bias_risk tool provides a deep dive into the social risk components of your AI system.

06 **What is the AI ESG score?**

The AI ESG score is a composite metric from 0 to 100 that represents the holistic responsibility and ethical standing of an AI system based on its governance, social impact, and transparency.

07 **How can I estimate the cost of improving my AI's ethical standing?**

You can use the `project_investment_needs` tool to simulate the required investment in Euros to reach a specific target ESG score.

08 **Does this tool account for regulatory requirements like the EU AI Act?**

Yes, the scoring model incorporates compliance gaps that reflect the distance between current system states and requirements set by regulatory bodies and institutional investors.

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-esg-scoring-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 ESG Scoring 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 ESG Scoring 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 ESG Scoring Engine MCP
Server ID	`01a09db1-4288-7355-aed7-25e380139a02`
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-esg-scoring-engine.