MCP SERVER | NO CODE | CLOUD HOSTED

AI R&D Efficiency Analytics MCP

Turn research spending into measurable innovation metrics.

AI R&D Efficiency Analytics MCP connects your AI client to a specialized engine that tracks the ROI of artificial intelligence research. It correlates your financial R&D spend against tangible outputs like model releases and patent filings to generate a normalized efficiency score. Use it to monitor productivity, compare performance against industry benchmarks, and track innovation trends over time.

A+ Quality Score 100/100

ai | rd | productivity | benchmarking | innovation

MAINTAINED BY VINKIUS CLOUD FOR AI AGENTS · SEPTEMBER 2026
vinkius.com/en/ai-agent-connect/ai-rd-efficiency-analytics

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 R&D Efficiency Analytics MCP

4 tools available
Cloud-hosted on Vinkius

You can finally stop guessing if your AI research budget is actually producing results. This MCP gives your AI client the ability to analyze the relationship between what you spend on R&D and what you actually ship. Instead of looking at raw spend, you can look at normalized efficiency scores that account for both intellectual property gains and product velocity.

Your agent can break down your innovation profile to see if you are leaning too heavily into patent filings at the expense of model releases, or vice versa. It provides a clear view of your momentum by tracking how efficiency shifts across different project cycles. Whether you are evaluating a single research sprint or comparing your entire department's output against industry standards, this tool provides the mathematical backbone to justify your R&D roadmap.

Core Capabilities

01 — ROI Correlation

Your agent links financial R&D spend directly to tangible innovation outputs.

02 — Productivity Scoring

The MCP generates a normalized score to quantify research effectiveness.

03 — Benchmarking

Your AI client compares your internal metrics against industry standards.

04 — Trend Analysis

The tool tracks efficiency changes across multiple time periods to show momentum.

05 — Composition Breakdown

Your agent evaluates the balance between IP creation and product velocity.

One Click on Vinkius — From Prompt to Execution

Available at **vinkius.com/en/ai-agent-connect/ai-rd-efficiency-analytics** — connect your AI agent in three steps.

01 Connect your AI client to the Vinkius platform.

02 Provide your R&D spend and output data to your agent.

03 The MCP calculates efficiency scores and composition ratios.

04 Review the resulting metrics and benchmarks directly in your chat interface.

Get your analytics running in minutes by connecting your client to the Vinkius hosted MCP.

Built For

This MCP is built for leadership teams managing high-cost AI research budgets who need data-driven proof of progress.

AI Research Leads

They use the tool to justify budget allocations and track team velocity.

Venture Capitalists

They use it to evaluate the R&D efficiency of portfolio companies.

Finance Directors

They use the metrics to correlate R&D spend with actual IP and product growth.

What Changes When You Connect

01 Quantifies the relationship between R&D spend and innovation output.

02 Identifies if research is too focused on patents or product releases.

03 Provides industry context for internal productivity scores.

04 Tracks efficiency momentum over time to catch slowing productivity early.

Real-World Applications

Budget Justification

Show leadership how much innovation is generated per Euro spent on AI research.

Strategic Balancing

Determine if you need to shift focus from patent filings to faster model releases.

Competitive Benchmarking

Check if your R&D efficiency is keeping pace with industry standards.

Performance Monitoring

Watch how efficiency scores fluctuate as you scale your research teams.

Patterns to Avoid

Measuring spend without output

AVOID

You report that the R&D budget increased by 20% last quarter to show progress.

✓ INSTEAD

Connect your spend and output data to `calculate_efficiency_score` to show the actual normalized efficiency gain per dollar spent.

Ignoring the IP vs product balance

AVOID

You assume high patent filings mean the research team is highly productive.

✓ INSTEAD

Use `analyze_innovation_composition` to see if you are over-indexing on patents while neglecting product release speed.

Relying on internal intuition

AVOID

You tell leadership that your team is performing well because the engineers feel busy.

✓ INSTEAD

Run `compare_to_benchmarks` to get a mathematical comparison of your efficiency against current industry standards.

The Right Fit

Connect this MCP if you manage high-cost AI research budgets and need to prove the ROI of innovation. Do not connect it if you are a solo developer or a small team without significant R&D spend. The decision rests on whether you need to justify research roadmaps to leadership using hard data.

AI R&D Efficiency Analytics: 4 Specialized Tools

These tools convert raw R&D spending and output data into actionable efficiency metrics and industry benchmarks.

#	TOOL	DESCRIPTION
01	`analyze_innovation_composition`	This tool determines if your efficiency comes from building intellectual property or increasing product release speed. It breaks down the ratio between patent filings and model releases.
02	`calculate_efficiency_score`	This tool generates a core R&D efficiency metric for a specific project or timeframe. It uses your spend and output data to create a normalized score.
03	`compare_to_benchmarks`	This tool measures your specific efficiency score against established industry standards. It shows you exactly where you sit within the current market landscape.
04	`get_efficiency_trends`	This tool tracks how your productivity changes over multiple periods. It helps you see if your R&D momentum is increasing or slowing down.

See It in Action

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

U Calculate my R&D efficiency. I spent 500000 Euros, released 5 models, filed 2 patents, and my time to market was 8 months.

Your R&D efficiency score is 72. Your productivity status is Optimal, with improvement opportunities in increasing patent filings to balance your product velocity.

U Compare an efficiency score of 65 against an industry intensity of 0.15.

Your score is 5 points below the industry average. You are currently in the Average performance tier.

U Analyze my innovation composition for 10 model releases and 2 patent filings with a 1000000 Euro spend.

Your innovation is Product-heavy, with a release ratio of 83% and a patent ratio of 17%.

Frequently Asked Questions

01 **What data does this MCP need to calculate efficiency?**

You need to provide R&D spend, the number of model releases, and the number of patent filings.

02 **How does it define innovation composition?**

It compares the ratio of intellectual property (patents) against product velocity (model releases).

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

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

04 **Does this tool provide industry benchmarks?**

Yes, the compare_to_benchmarks tool allows you to measure your efficiency against established industry standards.

05 **How is the efficiency score calculated?**

The score is a normalized metric that correlates your financial R&D investment with your specific innovation outputs.

06 **What is the R&D efficiency score?**

It is a normalized metric from 0 to 100 that measures how effectively capital is converted into innovation, such as through `calculate_efficiency_score`.

07 **How can I compare my performance to the industry?**

You can use the `compare_to_benchmarks` tool to evaluate your score against established industry standards and performance tiers.

08 **Can I track efficiency over time?**

Yes, the `get_efficiency_trends` tool allows you to analyze momentum, direction, and volatility across multiple historical periods.

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-rd-efficiency-analytics": { "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 R&D Efficiency Analytics 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 R&D Efficiency Analytics. 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 R&D Efficiency Analytics MCP
Server ID	`01a09dde-a1b9-7243-8fce-c0e7cf196835`
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-rd-efficiency-analytics.