

M C P S E R V E R

N O C O D E

C L O U D H O S T E D

AI Use Case Diversification Engine.

Calculate the ROI and risk of your next AI expansion.

AI Use Case Diversification Engine provides the financial modeling you need to scale your AI portfolio. It calculates expansion revenue, required investment, and ROI by weighing market demand against technical feasibility. Use it to move from gut feelings to data-backed growth strategies for your AI product roadmap.

A+ Quality Score 100/100

roi

revenue

investment

market-analysis

ai-strategy



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.

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.

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.

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

Encryption at rest

Ed25519

PKI vault signatures

24h TTL

Ephemeral session keys

V8 Isolate

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.

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.

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.

06 — Honeypot Trap System

Phantom credentials are injected into isolated environments. If a honeypot 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.

FinOps

Cost tracking per tool, payload compression savings, budget optimization signals, and consumption trends.

Firewall & DLP

PII redaction activity, sensitive data protection counters, and security event timeline.

Agent Activity

Which AI clients are connecting, how often, and what they're doing — real-time session tracking.

Tool Health

Slowest and most error-prone tools, with actionable root-cause insights and performance baselines.

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 Use Case Diversification Engine MCP

4 tools available

Cloud-hosted on Vinkius

You shouldn't guess which AI features to build next. This MCP gives your agent the math it needs to evaluate new opportunities. Instead of debating whether a new model or tool is worth the cost, you can run specific scenarios to see the projected revenue and investment requirements. It looks at market demand and technical feasibility to give you a clear picture of what's actually viable. You can look at a single idea to see if it makes sense, or compare multiple strategies to see which one offers the best return for your specific budget. It also helps you understand the technical risk involved, so you don't overextend your engineering team on projects that are too complex to deliver.

Core Capabilities

01 — Revenue Modeling

Your agent calculates projected revenue based on market demand and cross-sell potential.

02 — Investment Analysis

The MCP determines the total investment required for new AI initiatives.

03 — Risk Assessment

Your agent evaluates technical complexity to assign a risk level to your strategy.

04 — Scenario Comparison

The tool compares two different growth paths to see which yields better results.

05 — Portfolio Summaries

Your agent generates high-level financial overviews for multiple opportunities at once.

One Click on Vinkius — From Prompt to Execution

Available at vinkius.com/en/ai-agent-connect/ai-use-case-diversification-engine — connect your AI agent in three steps.

- 01 Connect the MCP to Claude, Cursor, or Windsurf via Vinkius.
- 02 Provide your agent with data on demand, cost, and feasibility.
- 03 Ask the agent to run specific calculations or comparisons.
- 04 Receive immediate financial models and risk assessments.

You connect the MCP to your client and start asking questions about your AI roadmap.

Built For

This MCP is built for leaders who need to justify AI investments with hard numbers.
It turns technical ideas into financial models.

Product Managers

They use it to prioritize the feature roadmap based on ROI.

Finance Leads

They use it to model the impact of AI expansion on the company budget.

AI Strategists

They use it to weigh technical feasibility against market demand.

What Changes When You Connect

- 01 Replaces guesswork with calculated ROI and revenue projections.
- 02 Identifies high-risk technical paths before you commit resources.

-
- | | |
|----|---|
| 03 | Provides a standardized way to compare different growth strategies. |
| 04 | Quantifies the financial potential of a group of AI use cases. |
-

Real-World Applications

Roadmap Prioritization

Decide which AI features to build first by comparing their ROI and technical risk.

Budget Planning

Determine exactly how much capital is needed to fund a new AI expansion.

Risk Management

Check if your planned AI portfolio is too technically complex to execute safely.

Strategic Comparison

Compare two different ways to grow your AI offerings to see which is more profitable.

Patterns to Avoid

Relying on gut feeling

❌ AVOID

You decide to build a complex new feature because it sounds cool, without checking if the revenue justifies the engineering cost.

✓ INSTEAD

Use `analyze_opportunity_viability` to see if the potential payoff actually justifies the effort before you assign developers to it.

Ignoring technical complexity

❌ AVOID

You plan a massive rollout of three new models at once, assuming your team can handle the integration load.

✓ INSTEAD

Run `calculate_portfolio_risk` to see how much danger your roadmap faces due to technical complexity.

Comparing unrelated growth paths

❌ AVOID

You try to decide between two different expansion strategies by looking at them in isolation.

✓ INSTEAD

Use `compare_scenarios` to weigh one specific growth path against another to see which yields better results.

The Right Fit

Product managers and finance leads should connect this MCP if they need to justify AI spending with hard numbers. Engineering teams who prefer to work without clear ROI metrics will find it unnecessary. The deciding factor is whether you need to turn technical ideas into repeatable financial models.

4 Tools for AI Use Case Diversification

These tools provide the mathematical foundation for evaluating technical feasibility and financial returns on new AI initiatives.

#	TOOL	DESCRIPTION
01	analyze_opportunity_viability	This tool evaluates a single specific opportunity to see if it is worth pursuing. It checks if the potential payoff justifies the effort.
02	calculate_portfolio_risk	This tool assesses the risk level of your expansion strategy. It uses technical complexity to determine how much danger your roadmap faces.
03	compare_scenarios	This tool compares two different expansion strategies. Use it to weigh one growth path against another.
04	get_expansion_summary	This tool provides a high-level overview of the financial potential for a group of expansion opportunities. It's your quick way to see the big picture.

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 financial potential for these opportunities: a new predictive maintenance tool with demand 8, cross-sell €500,000, and feasibility 7, and a customer churn predictor with demand 6, cross-sell €200,000, and feasibility 9? Base cost is €100,000.



The total expansion revenue is €700,000 and the total investment required is €100,000, resulting in an overall ROI of 7.0.

U Is a new AI-driven supply chain optimizer worth pursuing if it has a market demand of 9, cross-sell potential of €300,000, technical feasibility of 4, and a base cost of €50,000?



The projected revenue is €2,700,000 and the projected cost is €125,000, resulting in an ROI of 21.6. This is a High Priority opportunity.

U What is the risk level for a portfolio with opportunities having feasibility scores of 3, 5, and 2?



The average feasibility is 3.33, which results in a High Risk level.

Frequently Asked Questions

01 How does this MCP calculate ROI?

It uses the projected revenue and the required investment to determine the return on investment for a specific use case or portfolio.

02 What data do I need to provide for a viability check?

You should provide market demand, cross-sell potential, technical feasibility, and the base cost.

03 Can I compare two different AI strategies?

Yes, you can use the `compare_scenarios` tool to weigh two different expansion paths against each other.

04 How is technical risk determined?

The MCP assesses risk by looking at the technical feasibility scores of the opportunities in your portfolio.

05 Which AI clients can use this MCP?

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

06 How is the ROI calculated?

ROI is calculated by dividing the total expansion revenue potential by the total investment required, which includes base development costs adjusted for technical feasibility.

07 Can I compare different expansion strategies?

Yes, you can use the `compare_scenarios` tool to compare the total ROI of two different sets of AI opportunities.

08 What factors influence the investment required?

The investment required is based on the base development cost and is adjusted by the technical feasibility score; lower feasibility increases the projected cost.

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 <code>"ai-use-case-diversification-engine": { "url": "..."} </code>
 Windsurf	MCP Settings → <code>mcp_settings.json</code> → 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 Use Case Diversification 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 Use Case Diversification 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 Use Case Diversification Engine MCP
Server ID	01a09d4f-5f99-735d-817e-d9316d36bb8e
Platform	Vinkius Cloud for AI Agents
Endpoint	<code>https://edge.vinkius.com/{token}/mcp</code>

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-use-case-diversification-engine.