

MCP SERVER

NO CODE

CLOUD HOSTED

Cyclic Steam Stimulation (CSS) Designer MCP

Optimize heavy oil recovery through precise thermal reservoir modeling.

Cyclic Steam Stimulation (CSS) Designer MCP gives reservoir engineers the computational tools needed to manage thermal EOR operations. Use it to calculate injection volumes, determine soak times, and predict oil recovery across multiple cycles. It turns your AI client into a specialized thermal engineering assistant that handles the heavy math for reservoir optimization.

A+ Quality Score 100/100

thermal-eor

steam-injection

reservoir-engineering

heavy-oil

optimization



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.

Cyclic Steam Stimulation (CSS) Designer MCP

4 tools available

Cloud-hosted on Vinkius

You can now run complex thermal EOR calculations directly within your AI client. This MCP provides the specific mathematical frameworks required to manage Cyclic Steam Stimulation (CSS) workflows. Instead of manually calculating heat transfer or injection requirements, you provide your reservoir parameters and let your agent handle the heavy lifting.

You can use it to determine exactly how much steam you need to hit thermal targets in a specific reservoir or find the sweet spot for soak periods to maximize heat transfer. It also helps you look ahead by predicting how much oil you'll actually pull from the well during a cycle. For long-term management, you can track how well your well is performing against theoretical maximums to see if your heat utilization is dropping over time. It's built to turn raw reservoir data into actionable injection and soak schedules.

Core Capabilities

01 — Steam Volume Planning

Your agent calculates the exact steam needed to hit specific thermal targets.

03 — Production Prediction

Your client estimates oil recovery volumes and cycle durations based on reservoir data.

02 — Soak Time Optimization

The AI determines the best idle period to maximize heat transfer into the heavy oil.

04 — Efficiency Monitoring

The MCP compares real-world performance against theoretical maximums to track well health.

One Click on Vinkius — From Prompt to Execution

Available at vinkius.com/en/ai-agent-connect/cyclic-steam-stimulation-css-designer — connect your AI agent in three steps.

- 01

Connect your AI client to the Vinkius hosted MCP.
- 02

Provide reservoir data like porosity, permeability, and viscosity to your agent.
- 03

Ask your agent to run specific calculations like steam volume or soak duration.
- 04

Receive precise engineering outputs directly in your chat interface.

Get your thermal engineering tools running in your AI client in minutes.

Built For

This MCP is built for professionals managing thermal enhanced oil recovery in heavy oil assets.

Reservoir Engineer

Hands off reservoir parameters like porosity and permeability to get rapid injection and soak schedules.

Production Engineer

Uses the tool to predict cycle outcomes and monitor well efficiency over time.

Field Operations Manager

Relies on the calculated steam volumes and soak durations to plan daily field activities.

What Changes When You Connect

- 01

Reduces manual calculation errors in steam volume and soak time planning.
- 02

Provides rapid estimates for oil recovery across multiple CSS cycles.

-
- 03** Identifies declining well performance by comparing actual results to theoretical maximums.
-
- 04** Speeds up the transition from reservoir data to actionable injection plans.
-

Real-World Applications

Initial Cycle Planning

Calculate the first steam injection volume and soak time needed for a new heavy oil well.

Recovery Forecasting

Predict the oil yield for upcoming cycles based on current reservoir conditions.

Multi-Cycle Optimization

Track how heat utilization changes over several cycles to adjust future injection volumes.

Well Health Audits

Compare current cycle performance against theoretical limits to decide when to move to a different EOR method.

4 Tools Available

#	TOOL	DESCRIPTION
01	calculate_soak_duration	This tool finds the ideal amount of time to let a well sit idle to ensure maximum heat transfer to the reservoir.
02	estimate_recovery_cycle	Use this to predict the total volume of oil produced and the overall length of a specific production cycle.
03	evaluate_cycle_efficiency	This tool assesses well health by comparing current performance metrics against theoretical maximums.
04	plan_steam_injection	This tool calculates the specific steam volume required to reach your thermal targets in a given reservoir.

See It in Action

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

- U** Calculate the steam injection volume needed for a reservoir with 25% porosity, 150 mD permeability, an initial viscosity of 5000 cP, and a target viscosity of 50 cP.



The required steam injection volume is 4500 m3 with an estimated heat loss of 120 m3.

- U** What is the optimal soak time for an injection volume of 3000 m3 in a reservoir with 100 mD permeability and a thermal diffusivity of 0.05?



The optimal soak time is 12 days.

- U** Estimate the recovery for the 2nd cycle where 4000 m3 of steam was injected, the soak time was 10 days, and porosity is 20%.



The expected oil recovery is 850 m3 and the total cycle duration is 25 days.

Frequently Asked Questions

01 What kind of reservoirs is this MCP for?

This MCP is specifically designed for heavy oil reservoirs requiring thermal EOR via Cyclic Steam Stimulation (CSS).

02 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.

03 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.

04 Can I use this to track performance over time?

Yes, you can use the `evaluate_cycle_efficiency` tool to compare current performance against theoretical maximums to monitor well health.

05 What data do I need to provide for steam injection planning?

You typically need to provide reservoir parameters such as porosity, permeability, and viscosity levels to get accurate steam volume calculations.

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>"cyclic-steam-stimulation-css-designer": { "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

Cyclic Steam Stimulation (CSS) Designer 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 Cyclic Steam Stimulation (CSS) Designer. 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	Cyclic Steam Stimulation (CSS) Designer MCP
Server ID	01a09ae5-e4b5-704f-8b8f-f97bffc72eee
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/cyclic-steam-stimulation-css-designer.