MCP SERVER | NO CODE | CLOUD HOSTED

Ethylene Plant Yield Model MCP

Predict and optimize petrochemical yields through steam cracking simulations.

Ethylene Plant Yield Model MCP gives your AI client the ability to simulate chemical transformations in steam cracking processes. It models how feedstock composition, cracking severity, and furnace architecture dictate product output. You can run sensitivity analyses, compare furnace designs, and find the exact cracking severity needed to hit specific yield targets.

A+ Quality Score 100/100

ethylene | steam-cracking | yield-optimization | petrochemical | industrial-modeling

MAINTAINED BY VINKIUS CLOUD FOR AI AGENTS · SEPTEMBER 2026
vinkius.com/en/ai-agent-connect/ethylene-plant-yield-model

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.

Ethylene Plant Yield Model MCP

4 tools available
Cloud-hosted on Vinkius

You can use this MCP to bring chemical engineering simulation directly into your AI-driven workflows. Instead of manually calculating how feedstock changes affect your output, you can ask your agent to run the numbers for you. It handles the heavy lifting of modeling hydrocarbon transformations by looking at your specific feedstock, the severity of the cracking, and the furnace design you are using.

Whether you are trying to figure out how a shift in paraffin content will change your ethylene production or you need to decide which furnace architecture is more efficient for a new feedstock, this MCP provides the mathematical backbone. You can test different scenarios in seconds, finding the sweet spot for cracking severity to maximize specific chemical yields. It turns your AI client into a specialized tool for petrochemical process optimization and feedstock sensitivity testing.

Core Capabilities

01 — Yield Prediction

Your agent calculates the exact product distribution based on feedstock and severity.

02 — Severity Optimization

The AI finds the specific cracking intensity required to maximize target chemical production.

03 — Feedstock Sensitivity Analysis

Your agent determines how much your output changes when feedstock composition shifts.

04 — Furnace Comparison

The tool compares different furnace architectures to evaluate efficiency differences.

One Click on Vinkius — From Prompt to Execution

Available at **vinkius.com/en/ai-agent-connect/ethylene-plant-yield-model** — connect your AI agent in three steps.

01 Connect your preferred MCP-compatible client to Vinkius.

02 Provide your feedstock composition and furnace details to your AI agent.

03 The agent calls the specific tool needed for your calculation.

04 The MCP runs the simulation and returns the precise yield or optimization data.

05 Your agent presents the results in a clear, actionable format.

You connect the MCP to your AI client and start running simulations immediately.

Built For

This MCP is built for professionals managing petrochemical production and process design.

Chemical Engineers

They hand off complex yield calculations and sensitivity modeling to the AI.

Process Optimization Specialists

They use the tool to find optimal cracking severities for specific feedstocks.

Plant Managers

They use the simulations to understand the impact of feedstock changes on plant output.

What Changes When You Connect

01 Reduces manual calculation time for complex yield profiles.

02 Identifies optimal operating points for specific chemical targets.

03	Quantifies the impact of feedstock variability on production.
04	Provides a way to compare furnace designs through simulation.

Real-World Applications

Feedstock Switching

Determine how a new batch of feedstock will change your ethylene and propylene yields before you run it through the furnace.

Targeted Yield Maximization

Find the exact cracking severity needed to push ethylene production to its highest possible level.

Equipment Design Evaluation

Compare the efficiency of different furnace architectures using simulated process data.

Operational Sensitivity Testing

Run simulations to see how sensitive your plant is to small changes in paraffin or olefin concentrations.

4 Tools Available

#	TOOL	DESCRIPTION
01	`calculate_yield_profile`	Predicts the distribution of products from a specific cracking process. It outputs the expected percentages for chemicals like ethylene, propylene, and butadiene.
02	`compare_furnace_efficiency`	Evaluates different furnace designs. It helps you see how changes in furnace architecture impact overall process performance.
03	`optimize_severity`	Finds the ideal cracking severity for your target chemicals. It identifies the exact setting needed to hit peak yield for a given feedstock.
04	`analyze_feedstock_sensitivity`	Measures how changes in feedstock composition impact your final output. Use this to see how shifts in paraffins or olefins change your product mix.

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 expected yield profile for a feedstock with 70% paraffins and 30% olefins at a cracking severity of 50 using a short residence time furnace?

The expected yield profile is: Ethylene 35.2%, Propylene 18.5%, Butadiene 4.1%, BTX 6.2%, and Byproducts 36.0%.

U What cracking severity should I use to maximize ethylene yield with this feedstock: {'paraffins': 60, 'olefins': 20, 'naphthenes': 10, 'aromatics': 10} in a standard furnace?

The optimal severity to maximize ethylene yield is 75, which results in an expected ethylene yield of 42.5%.

U How sensitive is the ethylene yield to changes in the paraffin content?

Increasing the paraffin content by 5% results in a 3.2% increase in ethylene yield.

Frequently Asked Questions

01 **What kind of simulations can this MCP perform?**

It performs steam cracking simulations, including yield profiling, severity optimization, feedstock sensitivity, and furnace efficiency comparisons.

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 **How does the MCP handle feedstock data?**

You provide the feedstock composition, such as paraffin or olefin percentages, and the MCP uses that data to calculate the resulting product distribution.

04 **Can I use this to optimize for specific chemicals?**

Yes, the optimize_severity tool is specifically designed to find the peak yield for specific chemicals based on your feedstock.

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

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

06 **How can I predict the yield of ethylene?**

You can use the `calculate_yield_profile` tool by providing the feedstock composition, cracking severity, and the furnace type.

07 **Can I optimize for propylene instead of ethylene?**

Yes, use the `optimize_severity` tool and set the target product to 'propylene'.

08 **How do I compare different furnace types?**

Use the `compare_furnace_efficiency` tool to evaluate how different furnace architectures perform for your specific feedstock and severity settings.

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 `"ethylene-plant-yield-model": { "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

Ethylene Plant Yield Model 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 Ethylene Plant Yield Model. 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	Ethylene Plant Yield Model MCP
Server ID	`01a094ea-84c0-7130-afb4-d3708285ed3a`
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/ethylene-plant-yield-model.