

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

Catalyst Lifetime Prediction MCP

Model deactivation kinetics and plan refinery maintenance.

Catalyst Lifetime Prediction MCP gives your AI client the ability to model catalyst deactivation and operational lifespans. Use it to calculate remaining life, determine the best time for regeneration, and simulate how changes in feed or temperature impact your catalyst's health through coking and metal poisoning analysis.

A+ Quality Score 100/100

catalyst

refinery

kinetics

maintenance

prediction



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.

Catalyst Lifetime Prediction MCP

4 tools available

Cloud-hosted on Vinkius

You can use this MCP to bring chemical kinetics and refinery operational data directly into your AI workflow. Instead of manually calculating deactivation rates, you let your agent handle the heavy lifting. It analyzes coking, metal poisoning, and specific operating conditions to give you a clear picture of your catalyst's status.

You can run simulations to see how a change in feed temperature or process conditions will accelerate deactivation before you actually make the change in the field. It helps you move away from guesswork by providing specific windows for regeneration and clear urgency levels for replacement. Whether you are trying to optimize a maintenance schedule or assess the risk of a new operating parameter, this MCP provides the mathematical backbone for your decisions.

Core Capabilities

01 — Deactivation Modeling

Your agent uses this to project how quickly a catalyst loses activity over time.

02 — Maintenance Planning

The AI identifies the best time for regeneration to minimize production impact.

03 — Risk Assessment

Your client evaluates the urgency of replacement based on coking and metal poisoning.

04 — Process Simulation

The agent tests how new operating conditions will change the catalyst's lifespan.

One Click on Vinkius — From Prompt to Execution

Available at vinkius.com/en/ai-agent-connect/catalyst-lifetime-prediction — connect your AI agent in three steps.

- 01 Connect your preferred MCP-compatible client like Claude or Cursor to Vinkius.
- 02 Provide your current catalyst operating data to your agent.
- 03 Ask your agent to run a simulation or a life prediction.
- 04 Receive specific, actionable data on remaining life or replacement urgency.

Connect your AI client to Vinkius to start running refinery analytics immediately.

Built For

This MCP is built for professionals managing refinery operations and chemical processing units.

Refinery Engineers

They use the agent to run deactivation simulations and check replacement urgency.

Maintenance Planners

They rely on the agent to find optimal regeneration windows for scheduled shutdowns.

Process Technicians

They use the agent to understand how feed changes affect catalyst health.

What Changes When You Connect

- 01 Reduces guesswork in maintenance scheduling by providing specific regeneration windows.
- 02 Prevents unplanned downtime by predicting replacement urgency before failure.

- 03** Quantifies the impact of process changes on catalyst lifespan through simulation.
-

Real-World Applications

Optimizing Regeneration

Find the exact window to regenerate a catalyst to maximize uptime.

Shutdown Planning

Determine if a catalyst needs replacement during the next scheduled turnaround.

Feed Change Impact Analysis

Simulate how a new feed composition will accelerate coking or poisoning.

Operational Risk Management

Monitor real-time deactivation kinetics to assess the risk of running a catalyst too long.

4 Tools Available

#	TOOL	DESCRIPTION
01	get_replacement_urgency	This tool determines how soon you need to replace a catalyst. It analyzes current operating conditions and poisoning levels to set an urgency level.
02	predict_remaining_life	This tool calculates the expected operational lifespan left for your catalyst. It uses current deactivation data to project future performance.
03	simulate_deactivation_scenario	This tool models the impact of changing process or feed conditions. It shows how shifts in temperature or feed composition affect deactivation rates.
04	calculate_regeneration_window	This tool identifies the best time to perform catalyst regeneration. It finds the optimal window to restore activity without unnecessary downtime.

See It in Action

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

U How many days of operational life are left for our catalyst?



There are 45 days of operational life remaining before the catalyst must be replaced.

U What happens if we increase the feed temperature by 10 degrees?



Increasing the temperature will accelerate coking, reducing the projected remaining life by 12 days.

U Is it urgent to replace the catalyst now?



The urgency level is High. You have 5 days until the activity reaches the critical threshold.

Frequently Asked Questions

01 How does this MCP predict catalyst life?

It uses deactivation kinetics, including coking and metal poisoning data, to model the remaining operational lifespan.

02 Can I use this with Claude or Cursor?

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

03 What happens if I change my feed temperature?

You can use the simulation tool to see exactly how much that temperature change will accelerate deactivation and reduce remaining life.

04 How do I know when to replace a catalyst?

The tool provides a replacement urgency level based on current activity and operating conditions.

05 Do I need to host the MCP myself?

No, Vinkius hosts and manages the MCP for you, so it is ready to use as soon as you connect your client.

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>"catalyst-lifetime-prediction": { "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

Catalyst Lifetime Prediction 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 Catalyst Lifetime Prediction. 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	Catalyst Lifetime Prediction MCP
Server ID	01a097f8-48ca-70a5-8172-f2ee76bd4598
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/catalyst-lifetime-prediction.