

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

Fired Heater Design Engineer MCP

Perform precise thermal modeling and equipment sizing for industrial heaters.

Fired Heater Design Engineer MCP gives your AI client the specialized math needed for industrial heater engineering. It handles the heavy lifting for radiant section sizing, tube coil configurations, stack temperature predictions, and fuel consumption estimates based on your specific heating duty and fluid properties.

A+ Quality Score 100/100

thermal

heat-transfer

industrial-design

engineering

combustion



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.

Fired Heater Design Engineer MCP

4 tools available

Cloud-hosted on Vinkius

You can now bring industrial thermal engineering directly into your AI workflow. This MCP acts as a specialized calculator for fired heater design, allowing your agent to move past general reasoning and into precise thermal modeling. Instead of manually checking heat transfer equations, you can provide your heating duty, fluid properties, and fuel composition to get immediate, actionable sizing data.

Your agent uses this tool to determine the physical dimensions required for radiant sections and the specific configurations for tube coils. It also handles the efficiency side of the design, predicting stack exit temperatures and calculating the exact fuel mass or volume needed to meet your targets. It's built to handle the complex relationship between heating duty and environmental losses, making it a practical tool for equipment sizing and efficiency optimization in industrial settings.

Core Capabilities

01 — Radiant Section Sizing

Your agent calculates the necessary volume and dimensions for heat transfer.

03 — Thermal Loss Prediction

Your agent predicts stack exit temperatures and environmental energy losses.

02 — Tube Coil Configuration

The AI determines the number of tubes and their layout for heating coils.

04 — Fuel Demand Estimation

The tool calculates the fuel mass or volume required for a specific duty.

One Click on Vinkius — From Prompt to Execution

Available at vinkius.com/en/ai-agent-connect/fired-heater-design-engineer — connect your AI agent in three steps.

- 01 Connect your MCP-compatible client to Vinkius.
- 02 Provide your heating duty, fluid properties, and fuel data to your agent.
- 03 The agent calls the specific design tool needed for the task.
- 04 Receive precise sizing, temperature, or consumption results.

Connect your client to Vinkius and start running engineering calculations immediately.

Built For

This MCP is built for engineers working in thermal design and industrial combustion.

Thermal Design Engineer

Hands off heating duty and fluid data to the AI to generate equipment sizing.

Process Engineer

Uses the tool to check fuel efficiency and stack temperature predictions.

Mechanical Engineer

Provides duty requirements to get tube coil configurations and physical dimensions.

What Changes When You Connect

- 01 Automates complex thermal modeling calculations.
- 02 Provides precise tube and radiant section dimensions.
- 03 Calculates fuel requirements based on specific compositions.

04 Predicts stack exit temperatures to manage energy losses.

Real-World Applications

Equipment Sizing

Determine the physical dimensions of a radiant section for a new heater design.

Fuel Planning

Estimate the total fuel mass needed to satisfy a specific heating duty.

Efficiency Optimization

Predict stack temperatures and energy losses to improve fuel efficiency.

Coil Design

Calculate the number and configuration of tubes for heating coils.

4 Tools Available

#	TOOL	DESCRIPTION
01	get_fuel_consumption	This tool estimates the total mass or volume of fuel you need to meet a specific heating demand.
02	get_radiant_section_sizing	Use this to find the physical dimensions required for a radiant section to hit your heat transfer targets.
03	get_stack_temperature_and_losses	This tool predicts the flue gas temperature at the stack and calculates energy lost to the environment.
04	get_tube_coil_design	This tool calculates the number of tubes and their physical configuration for heating coils.

See It in Action

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

- U** Calculate the radiant section size for a heating duty of 50 MW with a specific fluid and fuel profile.



The required radiant section volume is 1250 m³, with a tube surface area of 450 m² and an estimated height of 12 meters.

- U** What will be the stack temperature if I use a 15% excess air ratio?



The predicted stack exit temperature is 185C with a flue gas volume flow of 4500 m³/h.

- U** Estimate the fuel consumption for a 25 MW duty at 85% efficiency.



The estimated fuel mass flow rate is 420 kg/h based on the provided fuel composition.

Frequently Asked Questions

01 What kind of engineering calculations does this MCP perform?

It performs thermal modeling for radiant section sizing, tube coil design, stack temperature prediction, and fuel consumption estimation.

02 Which AI clients can I use with this MCP?

You can use it with any MCP-compatible client like Claude, Cursor, Windsurf, or 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 working.

04 Can I use this for fuel efficiency analysis?

Yes, the MCP can predict stack exit temperatures and calculate energy losses to help with efficiency targets.

05 Does it require any special setup?

No setup is required. Once you subscribe on Vinkius, it is ready to use with one click.

06 What can I calculate with this tool?

You can calculate radiant section dimensions, tube coil geometry, stack exit temperatures, heat losses, and total fuel consumption.

07 How do I provide fluid properties?

Fluid properties should be provided as a JSON string containing heatCapacity, density, and viscosity.

08 Does this tool account for efficiency targets?

Yes, the efficiency target is a required input for sizing and fuel consumption calculations to ensure the design meets specific performance goals.

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>"fired-heater-design-engineer": { "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

Fired Heater Design Engineer 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 Fired Heater Design Engineer. 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	Fired Heater Design Engineer MCP
Server ID	01a0914f-eb48-7026-bd8a-3aec72b5a944
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/fired-heater-design-engineer.