

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

Flare Radiation Modeling MCP.

Calculate safety zones and stack heights using API 521 standards.

Flare Radiation Modeling MCP provides technical simulation tools for flare systems based on API 521 standards. Your AI client can model thermal radiation intensity, identify safety exclusion zones for personnel and equipment, and determine required stack heights by accounting for gas composition, mass flow rates, wind tilt, and solar radiation.

A+ Quality Score 100/100

api-521

thermal-radiation

flare-systems

safety-modeling

industrial-safety



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.

Flare Radiation Modeling MCP

4 tools available

Cloud-hosted on Vinkius

You can now run complex flare system simulations directly through your AI client. This MCP applies API 521 standards to your technical workflows, allowing you to model thermal radiation intensity and safety requirements without manual calculation. It handles the heavy lifting of accounting for gas composition, mass flow rates, and wind-induced flame tilt. You can use it to identify necessary exclusion zones to keep personnel and equipment safe from heat. It also calculates the minimum stack height needed to meet specific safety criteria. Instead of jumping between spreadsheets and manual formulas, you simply provide the parameters to your agent and get precise, standard-compliant results.

Core Capabilities

01 — Thermal Radiation Modeling

Your agent calculates radiation flux at specific coordinates.

02 — Safety Zone Identification

The AI determines necessary exclusion radii for personnel and equipment.

03 — Stack Height Optimization

Your client finds the minimum height needed to meet safety standards.

04 — Gas Property Analysis

The MCP retrieves heating values for specific gas mixtures.

05 — Environmental Factor Accounting

The tools incorporate wind tilt and solar radiation into the models.

One Click on Vinkius — From Prompt to Execution

Available at vinkius.com/en/ai-agent-connect/flare-radiation-modeling — connect your AI agent in three steps.

- 01 Connect your AI client to the Vinkius hosted MCP.
- 02 Provide gas composition, flow rates, or wind speeds to your agent.
- 03 The agent uses the specific API 521 tools to run the math.
- 04 Receive precise radiation, height, or zone data immediately.

You connect the MCP to your AI client and start running simulations through natural language.

Built For

This MCP is built for engineers working with flare systems and industrial safety standards.

Process Engineers

They hand off gas composition and flow rate data to the AI to get rapid safety calculations.

Safety Engineers

They use the tool to define exclusion zones and ensure personnel protection.

Flare System Designers

They rely on the stack height calculations to meet API 521 requirements.

What Changes When You Connect

- 01 Calculations follow API 521 standards directly.
- 02 Accounts for wind-induced flame tilt and solar radiation.
- 03 Reduces manual math for thermal radiation flux.
- 04 Provides specific exclusion radii for equipment and people.

Real-World Applications

Flare Stack Design

Determine the minimum height required for a new flare stack to meet safety limits.

Gas Composition Analysis

Get the heating value of a specific gas mixture to feed into radiation models.

Site Safety Planning

Map out exclusion zones around existing flares to protect workers.

Radiation Impact Assessment

Check the thermal radiation intensity at specific points on a facility layout.

4 Tools Available

#	TOOL	DESCRIPTION
01	calculate_radiation_intensity	This tool determines the thermal radiation flux at a specific coordinate relative to the flare.
02	calculate_required_stack_height	Use this to find the minimum height a flare stack must reach to satisfy safety standards.
03	determine_exclusion_zones	This tool identifies the safety radii needed to protect personnel and equipment from heat.
04	get_gas_energy_content	This tool retrieves the specific energy or heating value for a given gas mixture.

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 radiation intensity at 50 meters for an elevated flare with a methane-heavy gas composition and a flow rate of 10 kg/s?



The calculated radiation intensity at 50 meters is 4.2 kW/m².

- U** Calculate the required stack height for an elevated flare to keep ground radiation below 1.5 kW/m² given a 10 kg/s flow rate and 5 m/s wind.



The minimum required stack height is 24.5 meters.

- U** What is the personnel exclusion zone for a ground flare with a flow rate of 5 kg/s and a wind speed of 10 m/s?



The personnel exclusion radius is 32.0 meters.

Frequently Asked Questions

01 Does this MCP follow specific industry standards?

Yes, all calculations and models are based on the API 521 standard.

02 What kind of data do I need to provide for radiation modeling?

You typically need to provide flare type, gas composition, mass flow rates, and wind conditions.

03 Can I use this with Claude or Cursor?

Yes, you can use this MCP with any MCP-compatible client like Claude, Cursor, or Windsurf.

04 Does it account for wind speed?

Yes, the tools account for wind-induced flame tilt in the modeling process.

05 How does it handle different gas mixtures?

The MCP includes a tool to retrieve the specific energy or heating value for various gas mixtures.

06 How does this server handle wind effects?

The ``determine_exclusion_zones`` and ``calculate_required_stack_height`` tools incorporate wind speed to model flame tilt, which shifts the radiation profile and affects safety distances.

07 Can I calculate radiation for specific gas mixtures?

Yes, you can use ``calculate_radiation_intensity`` by providing a JSON string of the gas composition to determine the heat release rate and resulting flux.

08 What standards does this model follow?

The modeling logic is built upon the API 521 standard for flare systems.

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>"flare-radiation-modeling": { "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

Flare Radiation Modeling 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 Flare Radiation Modeling. 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	Flare Radiation Modeling MCP
Server ID	01a09156-98b1-7233-94b6-c4ff855aad73
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/flare-radiation-modeling.