

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

Backcountry Rescue Logistics MCP

Model emergency response and evacuation plans for wilderness rescues.

Backcountry Rescue Logistics MCP gives your AI client the ability to model critical wilderness emergency response data. It calculates total response times, checks if helicopters can fly in specific conditions, and helps you decide on the best evacuation strategy based on terrain and injury severity.

rescue

wilderness

logistics

emergency

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.

Backcountry Rescue Logistics MCP

4 tools available

Cloud-hosted on Vinkius

When you are managing a wilderness emergency, you need data that accounts for the reality of the terrain. This MCP lets your AI agent act as a logistics coordinator by processing variables like injury severity, communication delays, and weather constraints. You can use it to run quick simulations of how long a rescue will actually take or to determine if an aerial extraction is even an option given the current weather and topography. Instead of guessing, you get specific estimates for response times and evacuation priorities. It helps you move from a chaotic situation to a structured plan by identifying which resources to send where and how to prioritize different victims based on their condition and the difficulty of the terrain.

Core Capabilities

01 — Response Time Modeling

Your agent uses this to predict how long it takes to reach a victim.

02 — Aerial Feasibility Checks

The AI checks weather and terrain to see if helicopters can fly.

03 — Evacuation Strategy Selection

Your agent picks the best way to move a patient based on the situation.

04 — Urgency Categorization

The AI assigns priority levels to help you manage limited resources.

05 — Communication Delay Analysis

Your agent calculates the lag between a distress signal and responder awareness.

One Click on Vinkius — From Prompt to Execution

Available at vinkius.com/en/ai-agent-connect/backcountry-rescue-logistics — connect your AI agent in three steps.

- 01 Connect your client like Claude or Cursor to the Vinkius platform.
- 02 Provide the AI with incident details like terrain, injury, and weather.
- 03 The AI invokes the necessary tools to process the data.
- 04 You receive specific estimates for time, method, and priority.

Connect your AI client to the Vinkius hosted MCP to start running logistics models.

Built For

This MCP is built for professionals managing emergency response in remote areas.

Search and Rescue Coordinators

They hand off terrain and injury data to the AI to generate logistics plans.

Wilderness First Responders

They use the tool to estimate how much time they have before a situation changes.

Emergency Dispatchers

They use the tool to prioritize incoming calls from remote locations.

What Changes When You Connect

- 01 Reduces guesswork in time-sensitive rescue planning.
- 02 Provides objective data for helicopter deployment decisions.
- 03 Organizes resource allocation through clear priority levels.

Real-World Applications

Rapid Response Planning

A coordinator uses the tool to quickly estimate arrival times for a ground team in steep terrain.

Triage Support

An agent categorizes the urgency of multiple incidents to help prioritize limited teams.

Aerial Resource Management

A dispatcher checks if fog makes a helicopter rescue impossible before launching a flight.

Communication Assessment

Responders calculate signal delays to understand the true timeline of an incident.

4 Tools Available

#	TOOL	DESCRIPTION
01	evaluate_communication_latency	This tool calculates the time gap between an incident occurring and rescuers actually receiving the signal.
02	assess_aerial_feasibility	Use this to find out if a helicopter can actually be used for the rescue based on current conditions.
03	calculate_rescue_estimate	This provides your primary estimate for total response time and the most viable evacuation strategy.
04	determine_evacuation_priority	This tool categorizes the urgency of the situation to help you decide how to allocate resources.

See It in Action

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

U We are 5km from the roadhead in steep terrain. One person has a critical injury and we have a satellite messenger. How long will rescue take?



The estimated rescue time is 4.5 hours, and the suggested evacuation method is Air.

U Is it possible to use a helicopter for a rescue if the weather is foggy and the terrain is a flat meadow?



No, aerial extraction is not feasible due to the primary constraint of Weather.

U What is the priority for a group of 4 with a stable injury in clear weather?



The priority level is Stable, and no specialized resources are required beyond standard ground response.

Frequently Asked Questions

01 What AI clients can I use with this MCP?

You can use any MCP-compatible client, including Claude, Cursor, Windsurf, and VS Code.

02 Do I need to host the MCP myself?

No, Vinkius hosts and manages the MCP for you. You just connect and use it.

03 Can it tell me if a helicopter can fly?

Yes, the `assess_aerial_feasibility` tool determines if a helicopter is a viable option based on the conditions you provide.

04 How does it determine rescue priority?

It uses the `determine_evacuation_priority` tool to categorize urgency based on injury and terrain.

05 Does this work for ground rescues too?

Yes, it models various evacuation strategies, including ground-based responses.

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>"backcountry-rescue-logistics": { "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

Backcountry Rescue Logistics 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 Backcountry Rescue Logistics. 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	Backcountry Rescue Logistics MCP
Server ID	01a0a250-60c7-7282-b549-c126e64f4264
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/backcountry-rescue-logistics.