

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

Ferry Capacity Checker MCP

Keep your maritime boarding operations safe and organized.

Ferry Capacity Checker MCP gives your AI agent direct access to real-time maritime occupancy data. It handles the heavy lifting of tracking vessel limits, checking available passenger spots, and validating boarding requests to prevent overcapacity issues. You can connect it to your preferred AI client and start managing fleet logistics immediately.

A+ Quality Score 100/100

ferry

capacity

occupancy

maritime

logistics



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.

Ferry Capacity Checker MCP

4 tools available

Cloud-hosted on Vinkius

You can use this MCP to manage vessel occupancy without manually checking spreadsheets or logs. It gives your AI agent the ability to act as a real-time maritime coordinator. Instead of guessing how many people can fit on a boat, your agent can pull exact capacity numbers and check remaining space instantly. This is particularly useful for preventing safety violations by validating boarding requests before they happen. If you need a high-level view of how your fleet is performing, the agent can pull utilization status to see which vessels are nearing their limits. It turns your AI client into a logistics coordinator that understands the physical constraints of your maritime operations.

Core Capabilities

01 — Real-time occupancy monitoring

Your agent checks current passenger counts against vessel limits.

02 — Automated boarding validation

The agent confirms if a group can safely board a vessel.

03 — Fleet status reporting

Your agent identifies which ships in the fleet are reaching capacity.

04 — Capacity limit retrieval

The agent pulls the maximum passenger count for any specific ferry.

One Click on Vinkius — From Prompt to Execution

Available at vinkius.com/en/ai-agent-connect/ferry-capacity-checker — connect your AI agent in three steps.

- 01 Connect the MCP to your AI client via Vinkius.
- 02 Ask your agent to check a specific vessel's capacity or space.
- 03 The agent calls the necessary tool to pull real-time maritime data.
- 04 Your agent provides you with the exact passenger count or boarding status.

Connecting this MCP to your AI client gives your agent immediate maritime data access.

Built For

This MCP is built for maritime logistics professionals who need to automate passenger flow and safety checks.

Logistics Coordinators

They hand off fleet occupancy monitoring to the agent.

Port Operators

They use the agent to validate boarding requests for arriving groups.

Fleet Managers

They rely on the agent to report on overall fleet utilization.

What Changes When You Connect

- 01 Prevents overcapacity by validating passenger groups before boarding.
- 02 Provides instant access to vessel limits and remaining space.
- 03 Identifies fleet-wide utilization trends through a single request.

- 04 Reduces manual math for passenger counts and boarding approvals.
-

Real-World Applications

Safety Compliance

The agent blocks boarding requests that would exceed a vessel's legal passenger limit.

Dockside Management

Staff use the agent to quickly check how many people can still fit on an incoming ferry.

Fleet Optimization

Managers use the agent to see which vessels are underutilized or overbooked.

Passenger Flow Control

The agent manages the timing of boarding groups based on real-time space availability.

Patterns to Avoid

Manual capacity math

❌ AVOID

You ask your agent to guess if a group of 50 people can fit on a boat based on general knowledge.

✓ INSTEAD

Use ``check_remaining_space`` to get the exact number of available spots before approving any group.

Ignoring fleet trends

❌ AVOID

You focus only on one boat at a time and miss the fact that your entire fleet is hitting capacity limits.

✓ INSTEAD

Run ``get_fleet_utilization_status`` to see which vessels are nearing their limits across the whole operation.

Overlooking legal limits

❌ AVOID

You approve a boarding request because the deck looks empty, ignoring the official passenger cap.

✓ INSTEAD

Always use ``validate_boarding_request`` to ensure every group stays within legal safety margins.

The Right Fit

Maritime logistics teams and port operators should connect this MCP to automate safety checks. If you manage passenger flow through manual spreadsheets or verbal confirmations, you need this. Avoid this MCP if your operations do not involve real-time vessel occupancy or strict passenger limits.

Ferry Capacity Checker: 4 Operational Tools

These tools give your agent direct control over vessel occupancy, passenger limits, and fleet-wide utilization data.

#	TOOL	DESCRIPTION
01	check_remaining_space	This tool calculates exactly how many more passengers can board a specific ferry.
02	get_fleet_utilization_status	Use this to get a high-level overview of how many vessels in your fleet are currently at or near capacity.
03	get_vessel_capacity	This tool retrieves the maximum passenger limit for a specific ferry.
04	validate_boarding_request	This tool determines if a specific group of passengers is allowed to board based on the current vessel status.

See It in Action

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

U How many more passengers can fit on ferry V-102 if there are currently 45 people on board?



There are 55 spaces remaining on ferry V-102.

U Can a group of 10 people board ferry V-500 if it already has 95 passengers and a capacity of 100?



No, the group cannot board because the total occupancy would exceed the vessel's capacity.

U What is the maximum capacity for vessel B-77?



The maximum capacity for vessel B-77 is 200 passengers.

Frequently Asked Questions

01 How does this MCP prevent overcapacity?

The agent uses the `validate_boarding_request` tool to check if a group can safely board based on current occupancy and vessel limits.

02 Can I check the status of my entire fleet at once?

Yes, you can use the `get_fleet_utilization_status` tool to see how many vessels are near or at capacity.

03 Which AI clients can use this MCP?

You can use this with any MCP-compatible client like Claude, Cursor, Windsurf, or VS Code.

04 Does this MCP provide real-time data?

Yes, it is designed for real-time maritime capacity management and occupancy monitoring.

05 How do I find the limit for a specific boat?

You can ask your agent to use the `get_vessel_capacity` tool for any specific ferry.

06 How can I check if a ferry is full?

You can use the `check_remaining_space` tool to see how many spots are left or if the vessel is currently full.

07 Can I validate a group of passengers before they board?

Yes, the `validate_boarding_request` tool determines if a specific number of passengers can safely board based on current occupancy.

08 How do I see the overall fleet status?

Use `get_fleet_utilization_status` to get a high-level overview of how many vessels are near capacity or overloaded.

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>"ferry-capacity-checker": { "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

Ferry Capacity Checker 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 Ferry Capacity Checker. 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	Ferry Capacity Checker MCP
Server ID	01a0d18a-3227-73a8-94c4-ff3b5de75490
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/ferry-capacity-checker.