

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

Fleet Vehicle Assignment Optimizer MCP

Match service jobs to the right vehicles and drivers instantly.

Fleet Vehicle Assignment Optimizer MCP automates the logic of matching service jobs to your fleet. It calculates payload capacity, checks driver certifications, and verifies vehicle availability to ensure every job is assigned to a capable vehicle and an authorized driver.

A+ Quality Score 100/100

logistics

fleet

scheduling

optimization

transportation



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.

Fleet Vehicle Assignment Optimizer MCP

4 tools available

Cloud-hosted on Vinkius

You can stop manually checking driver logs and vehicle specs every time a new job comes in. This MCP handles the heavy lifting of logistical planning by looking at your fleet's constraints in real time. When you give your AI client a list of jobs, it cross-references payload limits, travel range, and driver credentials to build an optimized schedule. It prevents scheduling conflicts by checking if a vehicle is actually available and ensures you don't assign a heavy load to a vehicle that can't carry it. If you need to know how much room is left in a truck after a run, the MCP tracks remaining capacity for both weight and distance. It's a direct way to turn your fleet data into actionable assignments without the manual math.

Core Capabilities

01 — Automated Job Assignment

Your agent uses this to distribute multiple service tasks across your fleet at once.

02 — Driver Compliance Checks

The AI checks driver certifications to prevent unauthorized vehicle operation.

03 — Capacity Monitoring

Your agent calculates remaining payload and range to prevent overloading.

04 — Availability Verification

The AI checks vehicle schedules to avoid double-booking assets.

One Click on Vinkius — From Prompt to Execution

Available at vinkius.com/en/ai-agent-connect/fleet-vehicle-assignment-optimizer — connect your AI agent in three steps.

- 01 Connect your AI client to the Vinkius-hosted MCP.
- 02 Provide your job, vehicle, and driver data to your AI agent.
- 03 Ask the agent to assign jobs or check specific constraints.
- 04 The agent executes the corresponding tool to process the request.
- 05 Receive the finalized assignments or status checks immediately.

You connect the MCP to your AI client and start managing your fleet through natural language.

Built For

Logistics professionals and fleet managers use this to automate the tedious parts of dispatching and scheduling.

Fleet Manager

Hands off job lists and vehicle data to the AI to generate optimized assignments.

Logistics Coordinator

Uses the AI to verify driver eligibility and vehicle availability for specific routes.

Dispatch Supervisor

Relies on the AI to monitor vehicle capacity and prevent overages during high-volume shifts.

What Changes When You Connect

- 01 Reduces manual errors in payload and range calculations.

-
- 02 Prevents unauthorized driver assignments through automated credential checks.
 - 03 Eliminates scheduling conflicts by verifying vehicle availability.
 - 04 Speeds up the dispatch process by batching job assignments.
-

Real-World Applications

Daily Dispatching

Assign a morning batch of deliveries to a fleet of trucks while respecting weight limits.

Resource Planning

Check how much remaining capacity a vehicle has after completing its first three stops.

Driver Compliance

Verify that a driver is legally allowed to operate a heavy-duty vehicle before assigning a route.

Conflict Resolution

Confirm a vehicle is not already booked for maintenance or another job before assigning a new task.

Patterns to Avoid

Manual capacity math

❌ AVOID

You try to guess if a truck can handle a heavy load by looking at a spreadsheet. You end up overworking a vehicle or risking an overload.

✓ INSTEAD

Let your agent handle the math. Use ``calculate_remaining_capacity`` to get the exact weight and range left after every assignment.

Ignoring driver credentials

❌ AVOID

You assign a driver to a heavy-duty truck without checking if their license covers that specific vehicle class.

✓ INSTEAD

Run a check before you dispatch. Use ``check_driver_eligibility`` to ensure the driver is authorized for the assigned vehicle.

Double booking assets

❌ AVOID

You send a driver to a new job while the vehicle is already scheduled for maintenance or a different route.

✓ INSTEAD

Verify the schedule first. Use ``validate_vehicle_availability`` to confirm the vehicle is actually free for the requested window.

The Right Fit

Fleet managers and logistics coordinators should connect this MCP if they manage multiple vehicles and drivers. If you only run a single vehicle or don't need to track driver certifications, you don't need this. The decision comes down to whether you want to stop manually cross-referencing driver logs and vehicle specs.

4 Tools for Fleet Assignment and Capacity Management

These tools handle the math and logic behind driver eligibility, vehicle availability, and payload limits.

#	TOOL	DESCRIPTION
01	calculate_remaining_capacity	This tool monitors leftover resources like payload weight and driving range. It tells you exactly how much more a vehicle can carry or travel after a job is assigned.
02	check_driver_eligibility	This tool verifies driver credentials against specific vehicle requirements. It ensures your drivers are authorized for the class of vehicle they are assigned to.
03	validate_vehicle_availability	This tool checks for scheduling conflicts. It confirms whether a specific vehicle is free to take on a job during a requested time window.
04	assign_fleet_jobs	This tool distributes a batch of service jobs across your fleet. It matches jobs to vehicles based on capacity, range, and driver authorization.

See It in Action

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

U Assign these jobs to my fleet: `jobs=[{'id': 'j1', 'weight': 500, 'distance': 50}, {'id': 'j2', 'weight': 200, 'distance': 30}], vehicles=[{'id': 'v1', 'capacity': 1000, 'range': 200, 'class': 'Medium'}], drivers=[{'id': 'd1', 'classes': ['Medium']}]`



Job j1 has been assigned to vehicle v1 with driver d1. Job j2 has also been assigned to vehicle v1 with driver d1. Remaining capacity for v1 is 300kg payload and 120km range.

U Is driver d1 eligible to drive a Heavy class vehicle?



No, driver d1 is only authorized for Medium and Light vehicle classes.

U Check if vehicle v1 is available from 2024-05-01T10:00:00Z to 2024-05-01T14:00:00Z.



Vehicle v1 is available for the requested time window.

Frequently Asked Questions

01 How does this MCP handle vehicle capacity?

The MCP uses the `calculate_remaining_capacity` tool to track payload and range, ensuring you don't exceed vehicle limits.

02 Can I use this with Claude or Cursor?

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

03 Does it check if drivers are allowed to drive certain trucks?

Yes, the `check_driver_eligibility` tool verifies driver credentials against vehicle classes.

04 How do I avoid scheduling two jobs for the same truck?

The `validate_vehicle_availability` tool checks for scheduling conflicts before you finalize an assignment.

05 Do I need to host the MCP myself?

No, Vinkius hosts and manages the MCP for you. You just connect your client and start using the tools.

06 How does the assignment logic handle driver qualifications?

The system uses `'check_driver_eligibility'` to ensure that a driver's specific vehicle class authorization matches the vehicle they are being assigned to.

07 Can I check if a vehicle is available for a specific time window?

Yes, you can use the `'validate_vehicle_availability'` tool to check if a vehicle is free or if it conflicts with scheduled maintenance blackouts.

08 How do I know if a vehicle has enough space for a heavy load?

You can use `'calculate_remaining_capacity'` to determine the leftover payload and range of a vehicle after previous assignments.

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>"fleet-vehicle-assignment-optimizer": { "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

Fleet Vehicle Assignment Optimizer 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 Fleet Vehicle Assignment Optimizer. 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	October 2026
MCP Server	Fleet Vehicle Assignment Optimizer MCP
Server ID	01a0ed98-5857-7234-af93-f188d95f8cb2
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/fleet-vehicle-assignment-optimizer.