

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

Care Transport Booking Plan MCP

Manage medical transit and mobility logistics through your AI client.

Care Transport Booking Plan MCP handles the heavy lifting of medical logistics. It syncs medical appointments with specialized transit providers while respecting mobility requirements, driver availability, and specific budget constraints. Your AI agent uses this to build precise trip schedules and driver summaries without manual coordination.

A+ Quality Score 100/100

medical

transport

mobility

scheduling

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.

Care Transport Booking Plan MCP

4 tools available

Cloud-hosted on Vinkius

You can use this MCP to manage the complex movement of patients between medical appointments and their homes. Instead of manually checking if a vehicle can accommodate a wheelchair or if a driver is available, you let your AI client handle the math. It looks at mobility constraints and budget limits to pick the right provider for the job. You can generate full logistical plans, create specific timelines for pickups, and even pull up driver summaries to ensure everyone is on the same page. If a primary plan falls through, the MCP helps you find backup contacts immediately. It turns messy scheduling data into organized, actionable transport plans.

Core Capabilities

01 — Mobility-Aware Planning

Your agent uses this to ensure transport providers can actually handle specific rider needs like wheelchairs.

02 — Budget Management

The AI checks provider costs against your set limits during the planning phase.

03 — Timeline Calculation

Your client uses this to build precise departure and arrival sequences including necessary buffers.

04 — Contingency Identification

The agent pulls secondary contact info when a primary transport plan is no longer viable.

05 — Driver Documentation

Your agent creates quick summaries for drivers to review before they pick up a rider.

One Click on Vinkius — From Prompt to Execution

Available at vinkius.com/en/ai-agent-connect/care-transport-booking-plan — connect your AI agent in three steps.

- 01 Connect your AI client to the Vinkius hosted MCP.
- 02 Provide your agent with appointment details, mobility needs, and budget.
- 03 The agent runs the planning tools to generate schedules and provider matches.
- 04 Review the generated timelines or driver info sheets directly in your chat interface.

Connecting this MCP to your AI client gives you immediate access to medical logistics tools.

Built For

This MCP is built for professionals managing patient movement and medical logistics.

Medical Coordinators

They hand off appointment times and patient mobility needs to the AI to build transport plans.

Logistics Managers

They use the tool to manage provider availability and budget constraints for large-scale transit.

Healthcare Administrators

They rely on the tool to generate driver summaries and backup contact lists for patient safety.

What Changes When You Connect

- 01 Matches specific mobility needs with capable transit providers.
- 02 Calculates trip timelines that include necessary buffers for limited mobility.

-
- 03 Keeps transport costs within defined budgetary limits.
-
- 04 Provides driver-ready summaries to reduce communication errors.
-

Real-World Applications

Wheelchair Accessible Transport

You can coordinate specific pickups for patients requiring specialized vehicles and extra time.

Appointment Scheduling

The AI builds a full logistics plan for a patient's medical visit based on their location and budget.

Emergency Contingency

When a driver is delayed, you quickly identify backup contacts through the MCP.

Driver Briefing

You generate quick info sheets so drivers know exactly what mobility assistance a rider needs.

Patterns to Avoid

Manual timeline estimation

❌ AVOID

You try to guess how long a wheelchair pickup takes without accounting for extra loading time. This leads to missed appointments and frustrated patients.

✓ INSTEAD

Use ``generate_pickup_timeline`` to calculate precise sequences that include necessary buffers for mobility needs.

Ignoring budget constraints

❌ AVOID

You book the first available vehicle without checking if the cost exceeds your patient's coverage or department limits.

✓ INSTEAD

Let your agent use the planning tools to check provider costs against your set limits during the initial phase.

Relying on a single contact

❌ AVOID

You set up a transport plan but don't have a way to reach anyone if the driver gets stuck in traffic.

✓ INSTEAD

Always use ``get_backup_contacts`` to identify secondary contacts before the trip begins.

The Right Fit

Medical coordinators and logistics managers should connect this MCP if they manage high-volume patient movement with specific mobility requirements. If you only handle standard vehicle transport without budget or wheelchair constraints, you don't need this. The deciding factor is whether your transit planning requires matching specific rider needs to provider capabilities.

Care Transport Booking Plan MCP Tools

These tools handle the math and logistics for patient transit. They turn mobility constraints and budget limits into actionable schedules and driver briefings.

#	TOOL	DESCRIPTION
01	create_rider_info_sheet	This tool produces a concise summary for drivers and medical staff. It captures essential rider details and mobility needs in one place.
02	generate_pickup_timeline	This tool calculates the precise sequence of events for a single trip. It accounts for travel time and necessary buffers.
03	get_backup_contacts	This tool identifies secondary contacts to use if the primary transport plan fails. It provides a safety net for changing situations.
04	plan_transport_logistics	This tool generates a complete logistical plan. It includes timelines and priorities based on your specific constraints.

See It in Action

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

- U** Plan a transport for a wheelchair user for an appointment at 2024-10-15T14:00:00Z. Pickup is at 123 Maple St. Budget is \$100. Providers: [{ 'id': 'p1', 'name': 'MedTrans', 'cost': 80, 'capabilities': ['wheelchair'] }].



The transport plan is ready. Priority 1: MedTrans (\$80). Pickup is scheduled for 13:15:00Z to account for mobility lead time and travel buffers.

- U** Generate a rider info sheet for John Doe, who needs assistance with walking, at 456 Oak Ave.



Rider Information Sheet: Name: John Doe. Location: 456 Oak Ave. Mobility Alert: Requires assistance with walking.

- U** Calculate the pickup timeline for an appointment at 2024-11-01T09:00:00Z with a 30 minute travel duration and 15 minute buffer for limited mobility.



Departure time is 08:15:00Z. Arrival time is 08:45:00Z. A 15-minute buffer and mobility lead time have been applied.

Frequently Asked Questions

01 How does this MCP handle mobility constraints?

The MCP factors in specific needs, such as wheelchair access, when selecting transport providers and calculating pickup timelines.

02 Can I set a budget for transport?

Yes, you can provide a budget limit, and the tool will prioritize providers that fit within that amount.

03 What happens if a transport plan fails?

You can use the `get_backup_contacts` tool to identify secondary support contacts immediately.

04 Which AI clients can I use with this MCP?

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

05 Does it create documents for drivers?

Yes, the `create_rider_info_sheet` tool generates concise summaries for drivers and medical staff.

06 How does the tool handle mobility constraints?

The `'plan_transport_logistics'` tool evaluates provider capabilities against the specified mobility needs to ensure the vehicle and driver are appropriate for the rider.

07 Can I get a schedule for a specific trip?

Yes, you can use `'generate_pickup_timeline'` to calculate exact departure and arrival times, including travel buffers and lead times.

08 What happens if the primary transport provider is unavailable?

You can use `'get_backup_contacts'` to retrieve secondary or emergency contact information from the approved provider list.

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>"care-transport-booking-plan": { "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

Care Transport Booking Plan 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 Care Transport Booking Plan. 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	Care Transport Booking Plan MCP
Server ID	01a0d436-4cd3-73f6-a956-e4e8204a8e93
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/care-transport-booking-plan.