

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

Auto-Claim Repair Choice MCP

Turn messy repair estimates into a clear booking plan.

Auto-Claim Repair Choice MCP reconciles repair estimates, shop terms, and policy constraints. It helps your AI client build prioritized repair checklists and logical booking sequences by checking vehicle availability and shop authorization against specific policy limits.

A+ Quality Score 100/100

claims

repair

insurance

automotive

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.

Auto-Claim Repair Choice MCP

4 tools available

Cloud-hosted on Vinkius

You can stop manually cross-referencing repair quotes with policy documents. This MCP handles the heavy lifting of automotive insurance claims by acting as a decision engine for repair logistics. When you feed your AI client multiple estimates, it checks them against policy limits, deductible requirements, and parts preferences to find the best fit. It doesn't just look at cost; it verifies if a shop is actually authorized under the current policy and checks if the proposed repair timeline works with when the vehicle is actually available. Once the best option is identified, the MCP generates a step-by-step execution plan so you can move straight to booking. It turns fragmented data from shops and adjusters into a single, actionable sequence.

Core Capabilities

01 — Estimate Comparison

Your agent evaluates multiple quotes against policy limits and parts preferences.

02 — Shop Verification

The AI checks if a repair facility is authorized under specific policy tiers.

03 — Timeline Validation

The MCP confirms repair durations match the vehicle's available window.

04 — Execution Planning

Your agent generates a chronological sequence for completing the repair.

One Click on Vinkius — From Prompt to Execution

Available at vinkius.com/en/ai-agent-connect/auto-claim-repair-choice — connect your AI agent in three steps.

- 01
- Connect your AI client to the Vinkius-hosted MCP.
- 02
- Provide your agent with repair estimates, policy details, or shop information.
- 03
- The agent uses the MCP tools to run comparisons, validations, or sequence generations.
- 04
- Receive a prioritized repair checklist or a complete booking plan.

Connect your preferred AI client to Vinkius to start using these tools immediately.

Built For

This MCP is built for professionals managing automotive insurance workflows who need to bridge the gap between estimates and execution.

Claims Adjusters

Hand off estimate data and policy rules to the AI to find compliant repair options.

Insurance Operations Managers

Use the tool to verify shop networks and standardize repair booking sequences.

Automotive Logistics Coordinators

Provide vehicle availability and shop timelines to ensure repair windows align.

What Changes When You Connect

- 01
- Reduces manual errors by checking estimates against policy limits automatically.
- 02
- Prevents scheduling conflicts by validating repair timelines against vehicle availability.

-
- 03 Ensures policy compliance by verifying shop authorization before booking.
 - 04 Shortens the time from estimate to repair through automated booking sequences.
-

Real-World Applications

Policy Compliance Auditing

Check if a high-cost repair estimate exceeds the policy limit or violates parts preferences.

Repair Scheduling

Match a shop's estimated repair duration with the customer's vehicle availability window.

Network Management

Confirm a specific repair shop is in the authorized tier for a given policy.

Workflow Automation

Convert a selected repair option into a step-by-step chronological plan for execution.

Patterns to Avoid

Ignoring policy constraints

❌ AVOID

You pick the cheapest repair estimate without checking if the shop is actually in your network.

✓ INSTEAD

Always run the estimate through `compare_repair_options` to ensure it meets policy limits and parts preferences first.

Scheduling blind

❌ AVOID

You book a repair based on a shop's estimate but realize later the customer needs their car sooner.

✓ INSTEAD

Use `check_vehicle_window` to confirm the repair timeline actually fits the customer's availability before you commit.

Manual booking steps

❌ AVOID

You manually write out every single step needed to move a claim from estimate to a finished repair.

✓ INSTEAD

Let the MCP handle the logistics by using `generate_booking_sequence` to create a clear, chronological plan.

The Right Fit

Claims adjusters and operations managers should connect this MCP if they handle high volumes of automotive repair claims. If your workflow relies on manual spreadsheet cross-referencing to check policy limits, you need this. Avoid using this MCP if you only handle simple, single-step claims that don't require network or timeline validation.

4 Tools for Automating Repair Decision Logic

These tools handle the heavy lifting of comparing quotes, verifying shop networks, and building repair schedules.

#	TOOL	DESCRIPTION
01	check_vehicle_window	This tool checks if a proposed repair timeline actually fits within the vehicle's availability.
02	compare_repair_options	Use this to evaluate different repair estimates against specific policy constraints and user preferences.
03	generate_booking_sequence	This tool creates a chronological, step-by-step plan to execute the chosen repair.
04	validate_shop_eligibility	This tool verifies whether a specific repair shop is authorized under the current policy terms.

See It in Action

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

U Compare these repair estimates against my policy: estimates=[{"id":"opt1","totalCost":1200}, {"id":"opt2","totalCost":1500}] and policyConditions={"deductible":200,"limit":2000} and partsPreferences={"type":"OEM"}.



The recommended option is opt1 with a total cost of 1200, which is fully compliant with your 2000 limit.

U Is shop_abc123 eligible for my policy? approvedShopTerms=[{"shopId":"shop_abc123","tier":"Gold"}]



Yes, shop_abc123 is an eligible Gold tier shop.

U Can I schedule a 5-day repair if my vehicle is available from 2024-05-01 to 2024-05-10?



Yes, the 5-day repair fits within your availability window.

Frequently Asked Questions

01 What AI clients can I use with this MCP?

You can use this MCP with any 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 your client and start working.

03 Can this MCP help with scheduling?

Yes, it uses the `check_vehicle_window` tool to ensure repair timelines fit within the vehicle's availability.

04 How does it handle different repair shops?

The `validate_shop_eligibility` tool checks if a specific shop is authorized under the current policy terms.

05 Can it compare multiple quotes at once?

Yes, the `compare_repair_options` tool evaluates multiple estimates against policy constraints and preferences.

06 How does the tool evaluate repair options?

The `'compare_repair_options'` tool evaluates estimates by checking if the total cost stays within policy coverage limits and how well they match user parts preferences.

07 Can I verify if a shop is in my network?

Yes, you can use `'validate_shop_eligibility'` to check if a specific shop is authorized under your current policy terms.

08 What is a booking sequence?

A booking sequence is a chronological plan generated by `'generate_booking_sequence'` that covers shop confirmation, part procurement, vehicle drop-off, and final inspection.

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>"auto-claim-repair-choice": { "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

Auto-Claim Repair Choice 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 Auto-Claim Repair Choice. 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	Auto-Claim Repair Choice MCP
Server ID	01a0e15c-0e51-7333-ad83-65ac3333bb1b
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/auto-claim-repair-choice.