

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

Claim Follow-up Scheduler MCP

Keep insurance claims on track with proactive contact timelines.

Claim Follow-up Scheduler MCP builds intelligent communication schedules for insurance adjusters and claims teams. It analyzes claim status and contact history to predict upcoming deadlines and identify the most urgent tasks. You can use it within your AI client to ensure you meet SLAs and respect claimant communication preferences without manual tracking.

A+ Quality Score 100/100

claims

sla

automation

insurance-tech

scheduling-engine



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.

Claim Follow-up Scheduler MCP

4 tools available

Cloud-hosted on Vinkius

Managing insurance claims means juggling constant updates, strict deadlines, and specific claimant preferences. This MCP acts as a specialized scheduling engine that sits inside your AI client. Instead of you manually checking status logs to see who needs a call, your agent uses this MCP to build proactive timelines based on real claim data. It looks at current momentum and historical contact patterns to tell you exactly when to reach out and what your goal should be for each interaction. You can use it to spot risks before they become SLA violations or to verify that a proposed contact plan doesn't accidentally violate a claimant's request for email-only communication. It turns your AI agent into a proactive claims coordinator that knows what to do next.

Core Capabilities

01 — Timeline Generation

Your agent builds full contact plans with specific objectives for every touchpoint.

02 — Urgency Identification

The AI identifies the single most important task to prevent claim stagnation.

03 — SLA Risk Assessment

The tool predicts the probability of missing deadlines based on claim progress.

04 — Compliance Checking

Your agent verifies that contact methods align with claimant preferences and business rules.

One Click on Vinkius — From Prompt to Execution

Available at vinkius.com/en/ai-agent-connect/claim-follow-up-scheduler — connect your AI agent in three steps.

- 01
- Connect the MCP to Claude, Cursor, or Windsurf via Vinkius.
- 02
- Provide your agent with claim details like status, deadlines, and contact history.
- 03
- Ask your agent to generate schedules or check for urgent tasks.
- 04
- The agent executes the specific tools to return actionable timelines or risk alerts.

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

Built For

This MCP is built for professionals managing high volumes of insurance claims who need to maintain strict communication standards.

Claims Adjusters

They hand off the task of tracking contact dates and deadline monitoring to the AI.

Claims Managers

They use the tool to oversee SLA compliance and team workload distribution.

Insurance Operations Teams

They use it to ensure standardized communication across all active investigations.

What Changes When You Connect

- 01
- Reduces missed deadlines by predicting SLA risks before they happen.
- 02
- Prevents communication errors by validating contact methods against claimant preferences.

-
- | | |
|-----------|---|
| 03 | Eliminates manual schedule building by generating complete contact timelines automatically. |
| <hr/> | |
| 04 | Focuses attention on high-priority tasks through immediate action identification. |
| <hr/> | |
-

Real-World Applications

Preventing SLA Breaches

An adjuster uses the risk assessment tool to see which claims are falling behind schedule before a deadline is missed.

Preference Enforcement

The AI prevents an agent from sending an SMS to a claimant who has requested email-only contact.

Automated Contact Planning

A team creates a month-long communication roadmap for a new investigation in seconds.

Daily Task Prioritization

An agent starts their morning by asking the AI for the single most urgent task across their entire claim load.

Patterns to Avoid

Manual Schedule Creation

❌ AVOID

An adjuster spends two hours every Monday morning manually mapping out when to call every claimant on their desk.

✓ INSTEAD

Let the agent build the entire roadmap at once using ``generate_followup_schedule`` to save time and ensure nothing is missed.

Ignoring Claimant Preferences

❌ AVOID

A team sends a text message to a claimant who specifically requested all communication happen via email.

✓ INSTEAD

Run your proposed contact sequence through ``validate_communication_plan`` to catch preference violations before you hit send.

Reactive Deadline Management

❌ AVOID

A manager realizes a claim has breached its SLA only after the deadline has already passed.

✓ INSTEAD

Use ``calculate_sla_compliance_risk`` to identify which claims are trending toward a breach while there is still time to act.

The Right Fit

Claims professionals managing high volumes and strict communication standards should connect this MCP. If you are a solo adjuster with a handful of low-complexity claims, you likely don't need this. The deciding factor is whether your current workload makes manual deadline tracking and preference enforcement a constant risk.

Claim Follow-up Scheduler Toolset

These tools handle everything from predicting deadline risks to building full contact timelines for your claims.

#	TOOL	DESCRIPTION
01	calculate_sla_compliance_risk	This tool evaluates how likely you are to miss a critical deadline by analyzing current claim momentum and history.
02	generate_followup_schedule	This tool creates a complete timeline of future contact points and specific goals for each one.
03	get_next_immediate_action	This tool identifies the single most urgent follow-up task you need to handle right now.
04	validate_communication_plan	This tool checks if your proposed contact sequence breaks any business rules or specific claimant preferences.

See It in Action

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

- U** Generate a follow-up schedule for a claim in 'Active Investigation' status with a promised date of 2024-11-01 and a deadline of 2024-12-15.



The schedule includes contact points on 2024-10-28 via Email for information gathering, 2024-11-01 via Phone for a status update, and 2024-11-15 via Email for further information gathering.

- U** What is the most urgent action for a claim that has a deadline of 2024-10-25 and is currently in 'Pending Documentation' status?



The most urgent action is to call the claimant on 2024-10-24 to request the missing documentation.

- U** Check if this schedule is valid: contact on 2024-11-05 via SMS, where the claimant only allows Email.



The plan is invalid because the claimant has opted out of SMS communication.

Frequently Asked Questions

01 How does this MCP help with insurance SLAs?

It uses the `calculate_sla_compliance_risk` tool to evaluate the likelihood of missing deadlines based on current claim momentum.

02 Can I use this with my existing AI tools?

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

03 Does it handle claimant communication preferences?

Yes, the `validate_communication_plan` tool checks your proposed contact sequence against specific claimant preferences and business constraints.

04 How does it decide what to do next?

The `get_next_immediate_action` tool identifies the single most urgent follow-up task required for a specific claim.

05 Do I need to host this myself?

No, Vinkius hosts and manages the MCP, so it is ready to use immediately after you connect your client.

06 How does the scheduler determine contact frequency?

Frequency is determined by the claim status and the proximity to the deadline. For example, an Active Investigation requires more frequent contact than a Review/Approval stage.

07 Can I ensure the schedule respects claimant preferences?

Yes. You can use `'validate_communication_plan'` to check if a proposed sequence of contacts violates any specific channel or frequency preferences set by the claimant.

08 How is SLA risk calculated?

The `'calculate_sla_compliance_risk'` tool evaluates risk based on the current claim status, the approaching deadline, and the success rates found in the contact history.

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>"claim-follow-up-scheduler": { "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

Claim Follow-up Scheduler 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 Claim Follow-up Scheduler. 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	Claim Follow-up Scheduler MCP
Server ID	01a0e104-b6e8-7274-aceb-08dd32ca2c3c
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/claim-follow-up-scheduler.