MCP SERVER | NO CODE | CLOUD HOSTED

Float MCP

Manage team capacity in plain conversation.

Float connects your team's resource data to any AI agent. Manage scheduling, track time off, and allocate project hours by simply talking to your client. This MCP acts as a centralized source of truth for who's working on what, letting you instantly check capacity and update timelines without jumping between systems.

A+ Quality Score 100/100

resource-planning | capacity-management | team-scheduling | time-tracking | project-allocation | workload-management

MAINTAINED BY VINKIUS CLOUD FOR AI AGENTS · JUNE 2026
vinkius.com/mcp/float

SCAN TO VIEW

The infrastructure that powers AI agents in the real world.

Vinkius connects AI to the world's software through secure, enterprise-grade infrastructure — enabling real-world execution at scale, built on the Model Context Protocol (MCP).

Your AI Connections Run Through Vinkius Cloud

The world's largest managed MCP catalog

Vinkius is the cloud infrastructure where AI agents connect 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.

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.

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.

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	Ed25519	24h TTL	V8 Isolate
Encryption at rest	PKI vault signatures	Ephemeral session keys	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.

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.

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.

06 — Honeytoken Trap System

Phantom credentials are injected into isolated environments. If a honeytoken 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.

Agent Activity

Which AI clients are connecting, how often, and what they're doing — real-time session tracking.

FinOps

Cost tracking per tool, payload compression savings, budget optimization signals, and consumption trends.

Tool Health

Slowest and most error-prone tools, with actionable root-cause insights and performance baselines.

Firewall & DLP

PII redaction activity, sensitive data protection counters, and security event timeline.

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.

Float MCP

12 tools available
Cloud-hosted on Vinkius

Float lets you handle resource planning—the whole mess of figuring out who has time and who's swamped—using natural conversation. You connect your Float account through Vinkius, giving your agent access to all your team data. Instead of manually cross-referencing project boards with calendars, you ask your AI client directly: 'Is Jane available next week for the Website Redesign?' It checks availability and reports back instantly. You can manage everything from listing departments and clients to creating specific task allocations or retrieving actual logged hours worked. This means you never have to guess if a developer has bandwidth before committing them to a new project, keeping your whole operation running smoothly.

Core Capabilities

01 — Check team availability

Fetch detailed profiles and real-time status for every team member.

02 — View current projects and assignments

See a list of active client projects, including which department owns them and who is assigned to work on them.

03 — Schedule specific work tasks

Assign exact hours and dates for team members across multiple projects with one command.

04 — Manage time away records

Check schedules of planned vacations, sick days, or public holidays to plan capacity accurately.

05 — Analyze logged working hours

Compare the actual hours recorded against the scheduled project time to track progress and efficiency.

06 — Map organizational structure

List all clients, departments, or user accounts within your agency's framework.

One Click on Vinkius — From Prompt to Execution

Available at **vinkius.com/mcp/float** — connect your AI agent in three steps.

01 Subscribe to this MCP through the Vinkius Marketplace and provide your Float Personal Access Token.

02 Connect your preferred AI client (like Cursor or Claude) and authorize access using that token.

03 Ask your agent a natural language question, such as 'What is Jane's capacity next week?' and it executes the necessary calls.

The bottom line is you give your AI client the keys to Float, letting it read and write resource data on your behalf.

Built For

This MCP is built for operational leads who are tired of toggling between calendar apps, project management software, and time tracking sheets. If you spend more than 15 minutes a day just checking basic availability or scheduling follow-ups, this tool saves your sanity.

Project Manager

Quickly verify if a specific developer is available for a new task and assign the necessary hours to keep milestones on track.

Resource Planner

Get an immediate, comprehensive overview of team capacity across departments and account users, factoring in scheduled time off.

Agency Director

Automate the retrieval of complex reports, like total logged time or a list of all active clients, for quarterly operational reviews.

What Changes When You Connect

01 Saves you from manual data compilation. You can ask your agent to 'List all active projects and who is assigned,' instantly getting a full overview without opening multiple tabs or running complex reports.

02 Guarantees accurate scheduling. Use the `create_allocation` tool to schedule hours for specific tasks, updating team availability in real-time so you never overcommit someone.

03 Keeps you informed on capacity limits. By using `list_time_offs`, your agent automatically filters out time away, ensuring any allocation plan respects scheduled vacations and holidays.

04 Provides deep historical context. Running a query with `get_logged_time` lets you compare what was planned versus what was actually worked, giving clear data on project efficiency.

05 Knows your company structure. You can list departments or clients using dedicated tools so the agent always has full context about your agency's setup.

Real-World Applications

A developer needs to be assigned immediately, but you aren't sure who is free.

Instead of checking every team member's calendar manually, ask your agent to list all people and filter by current capacity. This lets you instantly identify the right fit for the task allocation.

The client asks for a report on last month's total billable hours.

Instead of exporting and summing time sheets, ask your agent to get logged time across all relevant projects. It compiles the actual hours worked instantly, ready for reporting.

You need to update a project timeline because a key team member is taking leave.

First, use list time offs to confirm Sarah's vacation dates. Then, ask your agent to adjust the project details and reallocate tasks around her absence.

You're onboarding a new account and need to see all possible departments.

Simply tell your agent to list departments or clients. It pulls the full organizational structure, giving you immediate context for resource planning.

Patterns to Avoid

Treating scheduling as a calendar function

✗ AVOID

Trying to schedule time off by just changing an event on a shared calendar. This leaves no record of the request and messes up capacity planning.

✓ INSTEAD

Use `list_time_offs` first, then ask your agent to manage time away records so that Float correctly updates everyone's official availability status.

Over-relying on project names for scope

✗ AVOID

Assuming a 'Marketing' project means all tasks are marketing. Without specific instructions, the agent might pull irrelevant task labels.

✓ INSTEAD

Always confirm available task labels first by calling `list_project_task_names`. This keeps your allocations accurate to Design or Development.

Manually checking individual availability

✗ AVOID

Having to ask five different people if they're free before scheduling a meeting. This is slow, inefficient, and requires too much back-and-forth.

✓ INSTEAD

Ask your agent to check team availability for the entire group at once. It runs checks across all listed people and reports on overall capacity.

The Right Fit

Use this MCP if your primary bottleneck is resource visibility. This tool excels when you need to know *who* has time, *when* they have time, and *what* project that time should be spent on. It's perfect for

Project Managers and Resource Planners who live in the overlap between calendars and task lists. Don't use this if your core problem is billing reconciliation or deep financial accounting; those require dedicated ERP connectors. If you just need a list of all client names, simple database tools might suffice, but Float adds the crucial layer: capacity planning. Always verify allocations using `list_allocations` to make sure nothing conflicts.

The time sink is checking if someone has bandwidth.

Today, figuring out who can work on a new task involves juggling multiple systems. You open the project management platform to see tasks, then switch to the team calendar to check availability, and finally jump into the department directory just to confirm job roles. It's a constant cycle of tabs, clicks, and cross-referencing that eats up hours.

With this MCP, you talk to your agent like talking to a colleague. You simply ask: 'Do we have anyone free next Tuesday who knows Python?' The agent handles the complex checks across all those systems, giving you one direct answer instead of six different dashboards.

Float MCP gives you total resource visibility.

The days of sending out 'Are you free?' emails that get ignored are over. You no longer need to manually cross-check if a developer's time off conflicts with an urgent project deadline, because the agent checks and reports it instantly via `list_time_offs`.

Now, your team’s capacity is visible in natural language commands. It’s immediate, comprehensive, and accurate.

Float: 12 Tools for Resource Planning

Use these tools in your agent to list projects, check staff availability, assign hours, or retrieve logged working time from Float.

#	TOOL	DESCRIPTION
01	`create_allocation`	Schedules specific hours for a task to a team member.
02	`get_logged_time`	Retrieves the actual number of hours worked by an employee.
03	`get_person`	Fetches detailed profile information for a specific team member.
04	`get_project`	Retrieves full details about a single client project.
05	`list_user_accounts`	Provides a list of all user accounts within the organization.
06	`list_allocations`	Lists existing scheduled tasks and allocations across the team.
07	`list_clients`	Pulls a comprehensive list of all active clients.
08	`list_departments`	Retrieves a categorized list of departments within the company structure.
09	`list_people`	Lists every team member, including their general profile details.
10	`list_project_task_names`	Retrieves a list of common task labels used in projects (e.g., Design or Development).
11	`list_projects`	Generates a complete list of all ongoing client projects.
12	`list_time_offs`	Checks and lists scheduled time off, including vacations or personal days.

See It in Action

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

U List all active projects in Float and the team members assigned to them.

Fetching projects... I found 5 active projects, including 'Website Redesign' (3 members) and 'Mobile App Launch' (5 members). Would you like to see the specific allocations for any of these?

U Schedule John Doe for 4 hours a day on the 'Q3 Marketing' project from Monday to Friday.

Allocation created! John Doe has been scheduled for 4 hours per day on the 'Q3 Marketing' project for the requested dates. His availability has been updated in Float.

U Who is scheduled for time off this month?

Checking time off... This month, Sarah is on vacation from the 10th to the 15th, and Mike has a personal day on the 22nd. Would you like me to check capacity for the rest of the team?

Frequently Asked Questions

01 **How do I use Float to check if people are available?**

You ask the agent to check team availability for a specific date range or group of users. This uses data from `list_people` and checks against scheduled time off.

02 **Can Float automatically assign tasks to John Doe?**

Yes, you can schedule specific work tasks by using the `create_allocation` tool. You just need to tell it the hours, dates, and project name.

03 **What information does Float keep about projects?**

It tracks detailed records via `list_projects`, including which departments own the projects and what types of tasks (like Design or Development) are involved.

04 **Does this MCP help with time tracking?**

Absolutely. You can use `get_logged_time` to pull actual hours worked against scheduled project times, which is key for billing and efficiency reporting.

05 **What if I need to see all the clients in my account?**

Use the `list_clients` tool. It gives you a full list of every client, helping you maintain context when planning resource 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 `"float": { "url": "…" }`
Windsurf	MCP Settings → `mcp_settings.json` → 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

Float 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 Float. 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	June 2026
MCP Server	Float MCP
Server ID	`019d759b-6cb2-739e-bfa6-f690bf602445`
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/mcp/float.