MCP SERVER | NO CODE | CLOUD HOSTED

Runn MCP

Get real-time resource and project data in chat.

Runn gives your AI client total visibility into resource planning, capacity forecasting, and project pipelines. Instead of clicking through dashboards to find out who's assigned where or how many hours were billed, you simply ask your agent in natural chat language. Get real-time data on people, projects, assignments, and actual timesheets instantly.

A+ Quality Score 100/100

resource-planning | capacity-forecasting | timesheet-tracking | project-pipelines | team-allocation | resource-management

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

SCAN TO VIEW

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.

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.

Runn MCP

12 tools available
Cloud-hosted on Vinkius

Think about managing a large client engagement. You need to know if the design team has bandwidth next month, who's working on Project X right now, and whether the hours logged actually match what was planned. This MCP connects your agent directly to Runn’s core data set. It lets you pull essential project metadata, current resource allocations, people details, and timesheet logs straight into your conversation.

You can ask for a list of all projects or check specific milestones for an initiative. Need to compare budgeted work against actual billed hours? Your agent pulls the `list_actuals` report instantly. Because this integration is hosted on Vinkius, you connect once from any compatible AI client and gain access to this deep project data without needing dedicated software logins.

This means your conversations become operational dashboards. Instead of juggling spreadsheets or navigating complex menus, you just ask: 'Who needs to sign off on the Beta project?' The answer comes back with names, roles, and dates. It's about getting immediate, actionable clarity on resource capacity.

Core Capabilities

01 — Identify all projects and clients

Get comprehensive lists of every project or client managed in the organization.

02 — Track team assignments and roles

Find out exactly who is assigned to which project phase, listing people, teams, and defined job roles.

03 — Review resource capacity and availability

Check if specific people are free or overloaded by querying their details across multiple projects.

04 — Analyze project timelines and goals

Retrieve key project milestones, track assigned team members, or see the full history of a single operational scope.

05 — Compare planned work versus billed hours

Pull detailed timesheet data to compare recorded actuals against anticipated labor costs.

One Click on Vinkius — From Prompt to Execution

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

01 Install the core Runn integration into your secure workspace.

02 Log into your primary Runn account online, navigate to Settings, and generate a permanent API access string.

03 Securely provide this specific Runn API Token in the integration's configuration section.

The bottom line is that providing an API token lets your AI agent read live project data from Runn and use it in conversation.

Built For

Project Managers, Agency Operations Leads, and Directors who are sick of clicking through multiple dashboards to answer simple questions. You're the person whose job requires knowing instantly if a resource is overbooked or when the next critical milestone hits.

Project Manager

Determining scheduling conflicts, checking resource constraints, and getting status updates on multiple concurrent projects without leaving your chat window.

Agency Operations Lead

Comparing billed actuals against forecasted hours to manage profitability across client accounts instantly.

Executive Director

Quickly requesting high-level summaries on overall initiatives, like which department is leading major efforts or what the next three critical milestones are.

What Changes When You Connect

01 Stop guessing about capacity. Instead of clicking through multiple dashboards to see who's available, ask your agent directly: 'Who is assigned to the internal design project?'

02 Manage budgets accurately by calling `list_actuals`. You get an instant comparison between planned work and the billed hours logged in timesheets.

03 Get a clear view of team structure using `list_people` and `list_roles`. You can immediately know who has which title without navigating HR systems.

04 Track project progress effortlessly. Use `list_milestones` to ask for all upcoming deadlines or use `get_project` to get a deep dive into one specific scope.

05 Avoid scheduling conflicts by checking team availability and roles simultaneously, combining data from `list_assignments` with your natural query.

Real-World Applications

Checking for resource bottlenecks before kickoff

The PM needs to know if the Beta project can start next week. They ask their agent, 'What capacity do we have?' The agent uses `list_assignments` and compares it against the available people from `list_people`, instantly flagging that three key developers are already allocated elsewhere.

Billing discrepancy investigation

The Agency Operations Lead notices a client account looks under-billed. They ask their agent to run a comparison, and the tool uses `list_actuals` against project metadata to show exactly which hours were logged but not invoiced.

Executive status report generation

The Executive needs a summary of all major initiatives. They prompt their agent to list all active projects using `list_projects`, and the tool gathers key details, including the next critical milestone for each one.

Understanding team structure for new hires

A new employee needs context on departmental roles. They ask their agent to list all defined roles using `list_roles` and then check which teams are currently utilizing those roles, providing an immediate organizational map.

Patterns to Avoid

Comparing multiple spreadsheets

✗ AVOID

Exporting timesheets to Excel, manually cross-referencing them with the project scope in a separate Google Sheet, and then trying to calculate capacity for next month.

✓ INSTEAD

Just ask your agent to use `list_actuals` or `get_project`. It aggregates that data live. Don't copy-paste anything; just talk to it.

Searching through internal wiki pages

✗ AVOID

Trying to find out who was assigned to a project six months ago by searching vague keywords in the company knowledge base.

✓ INSTEAD

Use `list_assignments` or `get_project`. The system tracks this relationship directly, giving you the authoritative answer instead of dusty wiki pages.

Relying on memory for dates

✗ AVOID

Remembering if a project was delayed because of a national holiday and manually checking an old calendar.

✓ INSTEAD

Use `list_holidays`. The MCP gives you the official, up-to-date list of non-working days to factor into any scheduling query.

The Right Fit

You should use this MCP if your core job involves answering complex questions about people, time, and resources (e.g., 'Who is available for X next month?' or 'Were we billed correctly on Y?'). This tool excels at connecting operational data points—people, projects, roles, actuals—into a single conversational answer.

Don't use it if you just need to search unstructured documents, like finding an old contract clause. For that, you need a document retrieval MCP. Also, don't use this if your primary goal is generating new content; focus on data retrieval and analysis first. If you only need a simple contact list, `list_people` handles it, but if you need to see *who is working* on what, the full suite of tools is necessary.

The Headache of Resource Visibility

Right now, figuring out resource capacity feels like a scavenger hunt. You open the project dashboard to see milestones, then you switch tabs to check who's assigned (the assignments list), and finally, you pull up the timesheet system just to compare actual hours. You spend twenty minutes clicking between three different tools just to answer one question: 'Can we launch this next month?'

With this MCP, that entire process collapses into a single query. Instead of opening multiple tabs or exporting data, you ask your agent a simple question and get the complete picture—milestones, resource roles, and current allocations—delivered right back to you.

Get Project Data with Runn MCP

The manual steps that vanish are opening the project scope, cross-referencing people lists against assignment tables, and then manually auditing billed hours. No more jumping between client portals and internal HR systems.

Now, you ask your agent to compare planned work versus actuals using `list_actuals`, and it gives you a definitive answer immediately. It's the difference between spending an hour building a report and getting the truth in five seconds.

Runn: 12 Tools for Project Operations

These twelve tools let your agent pull every piece of project metadata you need—from people rosters to billed hours—into your conversation.

#	TOOL	DESCRIPTION
01	`get_person`	Retrieves specific details for any person in the system.
02	`get_project`	Pulls detailed information about a single, specified project.
03	`list_actuals`	Generates a list of logged hours from timesheet data for billing review.
04	`list_assignments`	Lists all people and resources currently assigned across every project.
05	`list_clients`	Retrieves a list of all clients associated with the organization.
06	`list_holidays`	Provides a record of public holidays and non-working days for scheduling accuracy.
07	`list_milestones`	Lists all major completion points planned for a specific project.
08	`list_people`	Gets a complete roster of every person and resource in the Runn system.
09	`list_phases`	Shows all defined project phases for a specific operational scope.
10	`list_projects`	Retrieves a list of every active and completed project managed in the system.
11	`list_roles`	Lists all official job roles or positions defined within the organization's structure.
12	`list_teams`	Retrieves a list of every team set up within your workspace.

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 mapped.

I securely reliably naturally organically verified the Runn platform logs appropriately seamlessly dynamically pulling precisely 5 active tracked project namespaces successfully accurately flawlessly organically explicitly gracefully dynamically successfully smartly meticulously correctly comprehensively smartly efficiently naturally securely mapped seamlessly faithfully smoothly efficiently gracefully successfully smoothly.

U Which team is assigned to the Alpha project next week?

The 'Frontend Engineering' and 'Design' teams are assigned to Project Alpha next week. Lead designer Jane Doe is allocated for 20 hours, and developer Mark Smith for 35 hours.

U What are the upcoming milestones for the Beta project?

Project Beta has 2 upcoming milestones: 'Design Sign-off' on April 15th and 'MVP Release' on April 30th. Both are currently tracking on schedule.

Frequently Asked Questions

01 **How do I find out who is assigned to Project Alpha using Runn MCP?**

You can use the `list_assignments` tool. This pulls all current resources and people linked to that project, giving you a complete picture of the team structure.

02 **Can Runn MCP help me with billing questions?**

Yes, it uses `list_actuals`. You can ask your agent to pull detailed timesheet data to compare what was worked against what was budgeted or billed for a client.

03 **What is the easiest way to see all open projects?**

Just use `list_projects`. This tool gives you an immediate overview of every active project name and ID in the system, helping you narrow down your focus areas.

04 **How do I check who is available next month with Runn MCP?**

Check capacity using `list_people` and then cross-reference that list against current assignments. You can ask the agent to analyze both data sets for you.

05 **Do I need to worry about holidays when scheduling? Is Runn MCP accurate?**

No. The `list_holidays` tool provides an authoritative list of public and company non-working days, ensuring any schedule your agent builds accounts for proper time off.

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 `"runn": { "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

Runn 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 Runn. 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	Runn MCP
Server ID	`019d7600-e3fe-7133-9fba-2b8081d5242e`
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/runn.