MCP SERVER | NO CODE | CLOUD HOSTED

Deputy MCP for AI Agents

Manage staff schedules, track timesheets, and approve leave requests

Deputy lets your AI client manage the entire workforce lifecycle from within conversation. Your agent can pull employee details, review current shift rosters, track timesheets, and handle leave requests instantly. Stop switching between tabs; keep all scheduling and HR data in one place.

A+ Quality Score 100/100

workforce-management | employee-scheduling | timesheets | leave-management | roster-management | shift-planning

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

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.

Deputy MCP

10 tools available
Cloud-hosted on Vinkius

Managing staff schedules used to mean logging into a complex portal, cross-referencing spreadsheets, and digging through multiple forms just to know who's working when. Now, you can talk to your AI client about your workforce needs instead.

Your agent connects directly to Deputy to handle everything from employee directories to leave approvals using natural language conversation. Need to know which employees are clocked in right now? Ask it. Reviewing last week's payroll hours? Tell it to check the timesheets. It handles that instantly, giving you a clear picture of your operations.

When you connect this MCP via Vinkius, your agent gets access to all these tools—it's like having an Operations Manager on standby 24/7. You don't need to be an expert in the API; you just talk about staffing needs, and it does the heavy lifting for roster management, timesheet tracking, and leave coordination.

Core Capabilities

01 — Get employee details

Retrieve comprehensive personal information for any specific staff member.

02 — Check active shift status

Identify exactly which employees are currently clocked in and working their shifts.

03 — List all staff members

Pull a comprehensive directory of every employee within your organization's records.

04 — Review leave requests

View the full list of time-off and leave applications submitted by staff.

05 — Manage pending approvals

Filter down to only show leave requests that are awaiting your manager approval.

06 — Monitor schedules

Get a list of all current and upcoming shift rosters across locations.

One Click on Vinkius — From Prompt to Execution

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

01 Connect the Deputy MCP to your AI client. You'll authorize it using your Deputy API Token and account Subdomain.

02 Ask your agent a question, like 'Show me who is clocked in right now.'

03 The MCP runs the relevant tool, fetches the data, and returns a natural language answer that you can act on immediately.

The bottom line is, instead of writing code or clicking through menus, you just talk to your AI agent, and it handles complex HR tasks using Deputy's live data.

Built For

This MCP is built for Operations Managers and HR teams who are sick of juggling multiple dashboards. If checking shift coverage or auditing timesheets takes more than five minutes, this tool saves you time.

Operations Manager

Needs to quickly check current staffing levels, review upcoming rosters, and ensure proper site coverage when things go wrong.

HR Specialist

Must audit timesheets for compliance issues or monitor leave balances across the entire employee population without manual data entry.

Team Lead

Requires easy access to specific employee profiles and must manage scheduling adjustments on the fly during shift planning meetings.

What Changes When You Connect

01 Instantly know who is working: Use `list_currently_active_shifts` to get a real-time count of clocked-in employees without manually checking shift boards.

02 Streamline approvals: Instead of sifting through months of data, use `list_pending_leave_approvals` to focus only on requests that need your sign-off.

03 Better planning visibility: Review both current and future assignments by calling `list_active_rosters`, helping you spot coverage gaps weeks before they become problems.

04 Accurate payroll data: Quickly audit time records using `list_completed_timesheets` to ensure the hours billed match the actual shifts worked.

05 Full employee context: Use `get_employee_profile` and `search_employees_by_name` together to pull a complete history for any staff member in one conversation.

06 Organizational oversight: Get an overview of all operational sites using `list_business_locations`, making sure your agent knows where every shift is taking place.

Real-World Applications

Finding coverage gaps for tomorrow's shifts

A manager needs to know if the Downtown location has enough staff next Tuesday. They ask their agent, and it checks `list_active_rosters` and reports back on expected staffing levels across all roles.

Handling a last-minute sick call

A team lead gets an alert that key staff are out. They ask their agent to check `get_employee_profile` for the missing person and then use `list_workforce_employees` to find who can cover the gap.

Auditing timesheets before payroll

The HR specialist needs to verify hours for 50 employees. They ask their agent to pull all records using `list_completed_timesheets`, saving them from manually opening dozens of files.

Reviewing leave backlogs

A manager is overwhelmed by pending time off. Instead of reading 30 requests, they ask the agent to use `list_pending_leave_approvals` and get a concise summary of who needs approval.

Patterns to Avoid

Confusing HR data with payroll records

✗ AVOID

The user assumes that calling `get_employee_profile` will give them the final, approved hours for paychecks, but it only provides personal metadata.

✓ INSTEAD

To get verified work hours, always use `list_completed_timesheets`. If you need basic staff info, start with `search_employees_by_name` or `list_workforce_employees` first.

Overlooking location context

✗ AVOID

Trying to view all shifts using a general command without specifying the site. The agent might return hundreds of irrelevant results.

✓ INSTEAD

Always start by calling `list_business_locations` to confirm your scope, and then use that specific location name when requesting data via `list_active_rosters`.

Mixing up active vs. pending requests

✗ AVOID

The user asks about all leave, but the agent returns old, approved records mixed with current issues.

✓ INSTEAD

If you only care about immediate action, use `list_pending_leave_approvals`. If you need a historical view of everyone's time off, then call `list_leave_requests`.

The Right Fit

Use this MCP if your main pain point is coordinating complex workforce details like schedules, timesheets, and leave approvals across multiple systems. It excels when you need to cross-reference data—for instance, checking a staff member's profile (`get_employee_profile`) before reviewing their pending time off (`list_pending_leave_approvals`). However, don't use this if your primary need is managing payroll finances or deep accounting ledger entries. For those tasks, you need an accounting MCP, not a workforce management tool. If you just need to know who works at certain physical spots, start by calling `list_business_locations` to verify the scope before running any other query.

Deputy: How AI agents manage employee rosters and scheduling

Today, coordinating shifts means logging into Deputy, navigating between roster views, then checking timesheets in a separate tab to see who was actually there. You manually compare the scheduled shift against what was recorded, all while hoping you don't miss a gap or an overlap.

With this MCP connected, your agent handles that entire comparison instantly. Instead of clicking through multiple screens and comparing dates yourself, you simply ask: 'Who is supposed to be working at Location X on Date Y?' The answer comes back with the necessary roster details right away.

Deputy: Using AI agents for timesheet and leave request auditing

The current process forces HR staff to download massive CSVs of time records, then manually filter them by approval status or location. This is slow, error-prone work that takes hours every pay cycle.

Now, you can ask your agent directly: 'Show me all timesheets from the South Branch needing manager review.' It compiles and presents only the actionable data you need, saving massive amounts of time and keeping payroll accurate.

Deputy: 10 tools for workforce management and scheduling

Use these specific functions to get detailed employee records, review shift schedules, audit time sheets, and manage leave requests using your AI agent.

#	TOOL	DESCRIPTION
01	`get_employee_profile`	Pulls detailed information for a single, specified employee.
02	`get_authenticated_user`	Retrieves metadata about the user who is currently logged into Deputy.
03	`list_currently_active_shifts`	Identifies and lists all employees who are clocked in at this moment.
04	`list_workforce_employees`	Generates a complete list of every employee record in your organization.
05	`list_leave_requests`	Shows all submitted time-off and leave requests, regardless of approval status.
06	`list_business_locations`	Retrieves a list of every physical site or company location configured in Deputy.
07	`list_pending_leave_approvals`	Lists only the leave requests that are waiting for your direct managerial approval.
08	`list_active_rosters`	Provides a view of both current and future shift schedules across all locations.
09	`list_completed_timesheets`	Shows timesheets that have been submitted by employees for review.
10	`search_employees_by_name`	Finds and returns records for an employee based on their display name.

See It in Action

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

U How many people are clocked in right now across all locations?

Current Staffing Count:
- **Location A (Downtown):** 14 employees active.
- **Location B (Suburbs):** 6 employees active.
- **Total Clocked In: 20 people.**

You can also check who is clocked in at a specific spot, like 'Downtown', if you need to know the exact names.

U Are there any leave requests pending approval for next month?

Pending Leave Approvals (Next 30 Days):
- Sarah Lee: Personal Leave (Oct 5 - Oct 7)
- Mark Chen: Vacation (Oct 12 - Oct 16)
- Emily Rodriguez: Sick Time (Oct 24)

Please review these requests. They need your sign-off by the end of the week.

U Who is assigned to work at the East Side location next Monday?

East Side Roster - Next Monday:
- **Shift 1 (8:00 AM - 4:00 PM):** John Smith, Jane Doe
- **Shift 2 (11:00 AM - 7:00 PM):** Alex Kim, Maria Garcia

All shifts are currently covered. Let me know if you need the employee IDs for this team.

Frequently Asked Questions

01 **How can I use Deputy with AI agents to check staffing levels?**

You simply ask your agent, 'Who is working at my site?' It instantly queries the system and tells you who is clocked in right now or what the full roster looks like for any future date. This saves you from checking multiple dashboards.

02 **Can Deputy MCP help me manage leave approvals?**

Yes, it filters through all time-off requests and presents only those that require your specific manager approval. You get a clear list of who needs attention without having to sort through the entire employee history.

03 **What if I want to know an employee's full work history?**

You can pull detailed profiles for employees and then review their completed timesheets. This gives you a clear audit trail of everything they have worked, which is crucial for payroll accuracy.

04 **Is Deputy MCP good for multi-location scheduling?**

It's excellent for multi-site management. You can ask about rosters at specific business locations and get full coverage details, allowing you to compare sites easily in one chat session.

05 **Does this help me with payroll compliance checks?**

By accessing completed timesheets and employee profiles, your agent helps audit work hours against policy. You can ensure that the logged time matches the required shift duration before sending it to payroll.

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 `"deputy": { "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

Deputy 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 Deputy. 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	Deputy MCP
Server ID	`019d7584-2ab1-70bc-9248-385b0a98ea41`
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/deputy.