

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

Appointment Capacity Calculator MCP

Get exact scheduling limits and buffer math instantly.

Appointment Capacity Calculator handles the math for your service business. It calculates how many clients you can actually see within specific time windows while accounting for mandatory buffer times and service durations. Your AI client uses these tools to prevent overbooking and optimize your daily or weekly schedule.

A+ Quality Score 100/100

scheduling

capacity-planning

time-management

service-providers

optimization



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.

Appointment Capacity Calculator MCP

4 tools available

Cloud-hosted on Vinkius

Managing a service-based business means constantly juggling time. You need to know exactly how many clients you can fit into a morning block without cutting into your lunch or skipping necessary prep time. This MCP takes that math off your plate. Instead of manually calculating gaps, you tell your AI client your available hours, the length of your appointments, and how much buffer you need between them. It handles the logic of overlapping slots and constraints. You can use it to test different service models or check if a last-minute request actually fits into your existing window. It's about knowing your limits before you commit to a booking.

Core Capabilities

01 — Buffer Management

Your agent calculates the necessary downtime between sessions to prevent back-to-back burnout.

02 — Constraint Modeling

The AI uses specific start and end times to find realistic scheduling limits.

03 — Plan Comparison

Your agent compares different service lengths to see which yields more revenue or volume.

04 — Real-time Availability Checks

The AI checks if a new client request fits into your remaining time window.

One Click on Vinkius — From Prompt to Execution

Available at vinkius.com/en/ai-agent-connect/appointment-capacity-calculator — connect your AI agent in three steps.

- 01 Connect the MCP to Claude, Cursor, or Windsurf via Vinkius.
- 02 Provide your appointment duration and required buffer time to your AI client.
- 03 Ask the agent to calculate capacity or check a specific time window.
- 04 Receive the exact number of available slots or confirmation of availability.

Connect the MCP to your AI client and start asking about your schedule.

Built For

Service providers who need to manage tight schedules without manual math.

Service Providers

They hand off time-window math to the AI to plan their workdays.

Clinic Managers

They use the tool to determine how many patients can be seen in a shift.

Freelancers

They check if a new project or client fits into their existing weekly availability.

What Changes When You Connect

- 01 Eliminates manual math for appointment buffers.
- 02 Prevents overbooking by checking real-time slot availability.
- 03 Compares service models to find the most efficient schedule.

Real-World Applications

Daily Planning

Calculate how many clients you can see between 9 AM and 5 PM with 15-minute breaks.

Service Optimization

Compare a 30-minute service plan against a 45-minute plan to see the capacity difference.

Last-minute Bookings

Check if a specific client request fits into a remaining time gap.

Staffing Requirements

Determine how many staff members are needed based on total appointment capacity.

Patterns to Avoid

Ignoring buffer requirements

❌ AVOID

You tell your agent you have 8 hours and 30-minute appointments, assuming you can fit 16 people in.

✓ INSTEAD

Account for downtime between sessions. Use ``calculate_capacity_with_constraints`` to include specific start and end times along with necessary breaks.

Guessing on new bookings

❌ AVOID

A client calls for a 45-minute slot and you say yes without checking your actual remaining time.

✓ INSTEAD

Verify the gap first. Use ``check_slot_availability`` to see if that specific request actually fits your current window.

Sticking to one service model

❌ AVOID

You keep running 60-minute sessions even though they aren't maximizing your daily volume.

✓ INSTEAD

Test different configurations. Use ``compare_service_plans`` to see which service length yields a higher appointment count.

The Right Fit

Service providers and clinic managers should connect this MCP if they manage tight, back-to-back schedules. Freelancers who need to protect their personal time should use it to avoid overcommitting. If you don't have specific time constraints or buffer needs, you don't need this.

4 Tools for Precise Appointment Capacity Management

These tools handle the math of time blocks, buffer requirements, and service model comparisons.

#	TOOL	DESCRIPTION
01	calculate_max_capacity	This tool finds the total number of appointments that fit into a specific time block.
02	check_slot_availability	Use this to see if your current window has enough room to add one more appointment.
03	compare_service_plans	This tool compares two different service configurations to see which one maximizes your appointment count.
04	calculate_capacity_with_constraints	This tool calculates how many appointments fit when you provide specific start and end times.

See It in Action

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

U How many 30-minute appointments can I fit in a 4-hour window with a 5-minute buffer between each?



You can fit 7 appointments in a 4-hour window with those constraints.

U I have a window from 09:00 to 12:00. How many 45-minute appointments can I schedule with a 10-minute buffer?



You can schedule 4 appointments between 09:00 and 12:00.

U Is there enough time left for a 20-minute appointment if I have already used 100 minutes of a 150-minute window and need a 5-minute buffer?



Yes, there is enough time remaining to accommodate the appointment.

Frequently Asked Questions

01 How does this MCP handle buffer times?

The tools include buffer time in the calculation to ensure you have the necessary gap between appointments.

02 Can I use this with Claude?

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

03 What is the difference between the capacity tools?

One tool finds the maximum possible slots, while another handles specific start and end time constraints.

04 Does this help with comparing different service lengths?

Yes, the `compare_service_plans` tool lets you see which configuration allows for more appointments.

05 Do I need to host this myself?

No, Vinkius hosts and manages the MCP so it is ready to use immediately after connection.

06 How does the tool handle buffer times?

The tool treats the buffer as a required trailing component for every appointment slot to ensure operational readiness for the next possible slot.

07 Can I compare two different service models?

Yes, you can use `'compare_service_plans'` to determine which configuration of duration and buffer allows for more appointments within the same window.

08 What happens if the appointment is longer than the available time?

The tool will return an error if the requested appointment duration exceeds the total available window.

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>"appointment-capacity-calculator": { "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

Appointment Capacity Calculator 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 Appointment Capacity Calculator. 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	Appointment Capacity Calculator MCP
Server ID	01a0d090-3496-705a-b824-b3a3ef252edb
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/appointment-capacity-calculator.