

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

Beach Towel Calculator MCP

Handle group towel logistics without the math errors.

Beach Towel Calculator MCP handles the math for group towel requirements and service tiers. Your AI client uses this to calculate exact towel counts, check if your inventory can handle a group, and categorize groups into Small, Medium, or Large service tiers based on their needs.

A+ Quality Score 100/100

towel-calculator

group-planning

logistics-tool

beach-prep

capacity-management



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 — Honeytokens Trap System

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

Beach Towel Calculator MCP

4 tools available

Cloud-hosted on Vinkius

Managing towel inventory for large groups shouldn't be a guessing game. This MCP gives your AI agent the specific logic needed to handle beach towel logistics. Instead of you manually counting heads and multiplying by towel counts, your agent does the heavy lifting. You can quickly figure out how many towels to order for a specific crowd or check if you have enough surplus left over after a group is served. It also helps you organize service levels by automatically sorting groups into Small, Medium, or Large tiers. Whether you are managing a resort, a beach club, or a large event, this tool ensures your inventory stays aligned with your guest counts.

Core Capabilities

01 — Inventory Math

Your agent calculates total towel needs based on participant counts.

02 — Surplus Tracking

The AI identifies how many extra towels remain after fulfilling a request.

03 — Capacity Checks

Your agent verifies that towel requests don't exceed safe limits.

04 — Tier Categorization

The tool assigns service tiers like Small or Large to different groups.

One Click on Vinkius — From Prompt to Execution

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

- 01 Connect your preferred MCP-compatible client to Vinkius.
- 02 Provide your group size or towel count to your AI agent.
- 03 The agent calls the specific tool needed for the calculation.
- 04 Your agent delivers the final towel count or tier classification.

Connect your client to Vinkius and start sending logistics data to your agent.

Built For

This MCP is built for anyone managing beachside equipment or group hospitality.

Beach Club Managers

They hand off group size and towel requirements to the AI to plan inventory.

Event Coordinators

They use the agent to ensure enough supplies are available for large guest lists.

Resort Operations Staff

They use the tool to categorize service tiers and manage daily towel distribution.

What Changes When You Connect

- 01 Eliminates manual multiplication errors when planning for groups.
- 02 Automates the categorization of service tiers.
- 03 Prevents over-requesting through capacity validation.
- 04 Provides instant visibility into surplus inventory.

Real-World Applications

Daily Inventory Planning

Calculate exactly how many towels to bring to the beach based on the day's guest list.

Service Tier Management

Automatically group incoming requests into Small, Medium, or Large tiers for staffing.

Surplus Auditing

Check how many towels are left over after a large group departs.

Safety Compliance

Verify that a large group request doesn't exceed the total available capacity.

Patterns to Avoid

Manual Math Guessing

❌ AVOID

You try to multiply guest counts by towel ratios in your head or on a notepad before asking the agent for help.

✓ INSTEAD

Let the agent handle the math directly by using `get_total_towels` to get an exact count based on your group size.

Ignoring Capacity Limits

❌ AVOID

You approve a massive group request without checking if your current stock can actually support them.

✓ INSTEAD

Always run `validate_group_capacity` first to ensure the request stays within your safe operational limits.

Unorganized Service Tiers

❌ AVOID

You treat every group the same, regardless of whether they are a small family or a massive event crowd.

✓ INSTEAD

Use `estimate_group_tiers` to automatically sort groups into Small, Medium, or Large categories for better staffing.

The Right Fit

Connect this MCP if you manage beachside equipment, resort hospitality, or large event logistics. Do not connect it if you are managing indoor inventory or non-group based supplies. The decision depends entirely on whether your inventory needs are driven by guest counts and service tiers.

Beach Towel Calculator 4 Tools

These tools handle the math and logic for towel inventory and group service tiers.

#	TOOL	DESCRIPTION
01	estimate_group_tiers	This tool sorts groups into Small, Medium, or Large service tiers based on their towel needs.
02	get_total_towels	This tool calculates the exact number of towels required for a specific group size.
03	validate_group_capacity	This tool checks if a requested towel count stays within your safe operational limits.
04	calculate_excess_towels	This tool finds out how many extra towels you have left in your inventory.

See It in Action

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

U How many towels do I need for 10 people if everyone gets 2 towels?



You will need 20 towels in total.

U I have 50 towels. If I have 15 people and each needs 3 towels, how many extra do I have?



You have 5 extra towels.

U What tier is a group of 5 people needing 2 towels each?



This group falls into the Small tier.

Frequently Asked Questions

01 What can this MCP do?

It calculates towel requirements, checks group capacity, finds surplus towels, and assigns service tiers.

02 Which AI clients can use this?

You can use this with any MCP-compatible client like Claude, Cursor, Windsurf, or VS Code.

03 How does it handle group sizes?

The agent uses the `get_total_towels` tool to find the exact quantity needed for any number of participants.

04 Can it help with service levels?

Yes, it uses the `estimate_group_tiers` tool to sort groups into Small, Medium, or Large categories.

05 Does it check for inventory limits?

Yes, the `validate_group_capacity` tool checks if towel requests are within safe operational limits.

06 How do I calculate the total towels needed?

You can use the ``get_total_towels`` tool by providing the number of participants and the towels assigned per person.

07 Can I check if my towel request is too large?

Yes, use the ``validate_group_capacity`` tool to check if your requested towel count is within safe operational limits.

08 How are service tiers determined?

The ``estimate_group_tiers`` tool categorizes groups into Small, Medium, or Large tiers based on the total towel requirement.

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>"beach-towel-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

Beach Towel 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 Beach Towel 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	Beach Towel Calculator MCP
Server ID	01a0cdf7-f0a3-7080-8887-6f04ff15423b
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/en/ai-agent-connect/beach-towel-calculator.