

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

Recipe Batch Scaler MCP

Scale recipes and manage production batches with precision.

Recipe Batch Scaler MCP handles the math for professional kitchen scaling. It calculates ingredient amounts for specific serving sizes, accounts for prep waste, checks if your batch fits in your equipment, and determines your final portion yield. Your AI client does the heavy lifting so you can focus on the cooking.

A+ Quality Score 100/100

recipe

scaling

kitchen

batching

cooking

inventory



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.

Recipe Batch Scaler MCP

4 tools available

Cloud-hosted on Vinkius

Scaling a recipe from four people to forty isn't just about multiplication. You have to account for the flour that sticks to the bowl, the sauce that evaporates, and whether or not your largest stockpot can actually hold the finished product. This MCP handles those variables for you.

Instead of manual calculations, you tell your AI client how many servings you need and what your waste margin looks like. It gives you the exact ingredient weights required. It also helps you avoid kitchen disasters by verifying that your scaled batch won't overflow your containers. If you're managing inventory or planning a large event, you can use it to figure out exactly how many portions you'll get from a specific amount of raw product. It's a practical way to turn a single recipe into a production plan without the math errors.

Core Capabilities

01 — Ingredient Scaling

Your agent calculates ingredient totals for any number of servings.

02 — Waste Management

The MCP adjusts ingredient requirements based on your specific waste margins.

03 — Equipment Verification

Your AI client checks if a batch size exceeds your container limits.

04 — Yield Calculation

The tool determines the total number of portions available from a batch.

One Click on Vinkius — From Prompt to Execution

Available at vinkius.com/en/ai-agent-connect/recipe-batch-scaler — connect your AI agent in three steps.

- 01
- Connect your preferred MCP-compatible client to Vinkius.
- 02
- Ask your AI client to scale a recipe or check a batch size.
- 03
- The MCP runs the specific calculation tool requested.
- 04
- Your AI client delivers the precise measurements or capacity checks.

Connect your AI client to Vinkius and start scaling immediately.

Built For

This MCP is built for professionals who need to move from small-scale recipes to large-batch production.

Executive Chefs

They use it to scale daily prep lists and manage large event production.

Catering Managers

They use it to ensure they have enough portions and the right equipment for large orders.

Kitchen Managers

They use it to calculate ingredient needs and manage inventory waste.

What Changes When You Connect

- 01
- Eliminates manual math errors during large batch scaling.
- 02
- Prevents equipment overflow by checking container capacity.
- 03
- Accounts for prep waste to ensure you don't run short on ingredients.

04 Provides clear portion counts for better service planning.

Real-World Applications

Event Catering

Scale a standard recipe up to feed hundreds of guests while accounting for waste.

Kitchen Equipment Checks

Verify if a large batch of soup or sauce will fit in your specific stockpots before you start cooking.

Production Planning

Determine if your current stock of a prepared item will cover the expected number of portions.

Inventory Management

Calculate exactly how much raw product to order based on desired portion yields and waste margins.

Patterns to Avoid

Ignoring equipment limits

❌ AVOID

You scale a massive batch of chili for a wedding and realize halfway through that it won't fit in your largest stockpot.

✓ INSTEAD

Always run your scaled recipe through `validate_batch_capacity` first to ensure your equipment can handle the volume.

Underestimating prep waste

❌ AVOID

You calculate ingredients based on perfect yields, but end up short because of flour sticking to bowls or vegetable trimmings.

✓ INSTEAD

Use `get_scaling_efficiency` to see how your waste margin impacts your total ingredient requirements.

Guessing portion counts

❌ AVOID

You assume a batch will feed fifty people, but the actual portion sizes leave you with only forty servings.

✓ INSTEAD

Run `calculate_portion_yield` to get an accurate count of available servings based on your specific portion sizes.

The Right Fit

Professional kitchen managers and catering leads should connect this MCP to handle large-scale production. Home cooks or small cafes won't find much value here. The decision comes down to whether you are moving from single recipes to high-volume batch cooking.

Recipe Batch Scaler: 4 Precision Tools

These tools handle the math for ingredient weights, container limits, and portion yields.

#	TOOL	DESCRIPTION
01	calculate_portion_yield	Determines how many full portions you can get from a batch based on your specific portion size.
02	get_scaled_ingredients	Calculates the exact weight or volume of each ingredient needed for a specific number of servings.
03	get_scaling_efficiency	Analyzes how much your waste allowance increases the total amount of ingredients you need to buy or prep.
04	validate_batch_capacity	Checks if your scaled recipe will actually fit inside a specific container or piece of equipment.

See It in Action

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

U I have a recipe for 4 servings that uses 200g of flour. How much flour do I need for 10 servings with a 5% waste allowance?



You will need 525g of flour.

U Will a 5000ml batch fit in a 4.5L container?



No, the batch will not fit. It exceeds the container capacity by 500ml.

U If I have 1200g of sauce and each portion is 150g, how many portions can I serve?



You can serve 8 portions.

Frequently Asked Questions

01 How does this MCP handle ingredient waste?

The `get_scaled_ingredients` tool allows you to include a waste allowance, which increases the calculated ingredient amounts to ensure you have enough after prep.

02 Can I check if my pot is big enough for a recipe?

Yes, use the `validate_batch_capacity` tool to check if a specific batch volume will fit in your container.

03 What AI clients can I use with this MCP?

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

04 Does this MCP calculate how many people I can feed?

Yes, the `calculate_portion_yield` tool tells you exactly how many portions you can get from a batch based on your portion size.

05 How do I get started?

Since Vinkius hosts the MCP, you just connect your client once and you have instant access to all the scaling tools.

06 How does the waste allowance work?

The waste allowance adds a percentage to the scaled ingredients to compensate for loss during preparation, such as peeling or residue left in bowls.

07 Can I check if my batch will fit in my pot?

Yes, you can use the `validate_batch_capacity` tool to compare the total volume of your scaled recipe against the volume of your container.

08 How many servings will I get from a batch?

You can use `calculate_portion_yield` to find out exactly how many full portions can be extracted from your total batch quantity.

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>"recipe-batch-scaler": { "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

Recipe Batch Scaler 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 Recipe Batch Scaler. 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	October 2026
MCP Server	Recipe Batch Scaler MCP
Server ID	01a0f0fc-76ac-7090-9c1b-fed312c13ddd
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/recipe-batch-scaler.