

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

Paddle Stroke Counter MCP

Turn your paddling data into actionable fitness metrics.

Paddle Stroke Counter MCP gives your AI client the ability to process and analyze your paddling performance. It calculates total strokes, determines your rhythm intensity, and generates full session reports. You can also use it to check if your recorded stats are physically possible to ensure your data is accurate.

A+ Quality Score 100/100

paddling

fitness

cadence

sports-metrics

calculation



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.

Paddle Stroke Counter MCP

4 tools available

Cloud-hosted on Vinkius

If you track your time on the water, this MCP helps your AI client make sense of the numbers. Instead of manually calculating your effort, you can hand your session data to your agent and get immediate answers about your performance. It handles the math for total stroke counts and rhythm intensity, so you don't have to. You can ask your agent to pull a full session summary or use it to verify that your recorded metrics aren't impossible. It's a straightforward way to turn raw paddling numbers into a clear picture of your fitness and technique.

Core Capabilities

01 — Stroke Volume Calculation

Your agent uses this to find the exact number of strokes in a session.

02 — Rhythm Analysis

Your AI client uses this to determine how intense your paddling cadence was.

03 — Session Reporting

Your agent uses this to compile all your paddling data into a single summary.

04 — Data Validation

Your AI client uses this to confirm your recorded stats are realistic.

One Click on Vinkius — From Prompt to Execution

Available at vinkius.com/en/ai-agent-connect/paddle-stroke-counter — connect your AI agent in three steps.

- 01
- Connect the MCP to your client like Claude or Cursor via Vinkius.
- 02
- Provide your paddling data, such as duration and strokes per minute, to your agent.
- 03
- The agent calls the specific tool needed to process your data.
- 04
- Your agent delivers the calculated strokes, intensity, or summary directly in the chat.

Connecting this MCP to your AI client gives your agent immediate access to paddling math tools.

Built For

This MCP is built for water sports enthusiasts and fitness trackers who want to analyze their paddling performance using AI.

Paddle Athletes

They hand off session duration and cadence to the AI to track progress.

Fitness Coaches

They use the AI to validate athlete stats and summarize training intensity.

Data Analysts

They use the tool to verify the physical plausibility of sports datasets.

What Changes When You Connect

- 01
- Automates the math for stroke counts and intensity levels.
- 02
- Provides instant validation for recorded paddling metrics.
- 03
- Generates full session reports through simple natural language requests.

04 Works directly within your existing AI client workflow.

Real-World Applications

Post-Paddle Review

Ask your agent for a session summary immediately after you get off the water.

Training Intensity Tracking

Use the cadence tool to see if you are hitting your target effort levels.

Data Integrity Checks

Run your logged stats through the validation tool to catch errors in your tracking device.

Volume Monitoring

Calculate total stroke counts over multiple sessions to monitor your workload.

Patterns to Avoid

Manual Math Errors

❌ AVOID

You try to add up every single stroke from a long session using a calculator and then tell your agent the total.

✓ INSTEAD

Just give your agent the raw session data and use ``get_total_strokes`` to get the exact count without mistakes.

Ignoring Data Glitches

❌ AVOID

You assume your GPS or heart rate monitor recorded everything perfectly even when the numbers look weird.

✓ INSTEAD

Run your stats through ``validate_paddling_stats`` to make sure your recorded metrics are actually physically possible.

Vague Effort Checks

❌ AVOID

You tell your agent you felt tired and ask if you worked hard enough.

✓ INSTEAD

Provide your cadence data and use ``get_cadence_intensity`` to get a specific categorization of your effort level.

The Right Fit

Connect this MCP if you regularly track paddling sessions and want your AI client to handle the performance math. Do not connect it if you only paddle occasionally and don't care about analyzing your specific stroke volume or rhythm. The decision comes down to whether you want to turn raw numbers into actionable training insights.

4 Tools for Paddling Performance and Data Analysis

These tools handle the math and logic for your paddling metrics. Your agent uses them to calculate volume, check intensity, and verify data accuracy.

#	TOOL	DESCRIPTION
01	get_cadence_intensity	This tool determines the intensity level of your paddling rhythm. It analyzes your stroke frequency to categorize your effort.
02	get_session_summary	This tool provides a complete breakdown of your entire paddling session. It aggregates all relevant metrics into one report.
03	get_total_strokes	This tool calculates the total number of strokes you performed. It's useful for quick volume checks during or after a session.
04	validate_paddling_stats	This tool checks if your provided paddling metrics are physically plausible. It acts as a sanity check for your recorded data.

See It in Action

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

U How many total strokes did I perform if I paddled for 30 minutes at 45 strokes per minute?



You performed a total of 1,350 strokes.

U What was the intensity of my 15-minute paddle at 60 strokes per minute?



Your paddling intensity was High.

U Give me a summary for a 20 minute session with 50 strokes per minute.



Your session included 1,000 total strokes with an average cadence of 50 strokes per minute over 20 minutes, resulting in Moderate intensity.

Frequently Asked Questions

01 What AI clients can I use with this MCP?

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

02 How do I calculate my total strokes?

You can ask your AI client to use the `get_total_strokes` tool by providing your session duration and cadence.

03 Can this MCP check if my data is wrong?

Yes, the `validate_paddling_stats` tool checks if your recorded metrics are physically plausible.

04 Do I need to host this myself?

No, Vinkius hosts and manages the MCP for you, so it is ready to use once you connect.

05 What kind of intensity reports can I get?

The `get_cadence_intensity` tool determines your intensity level based on your paddling rhythm.

06 Can I verify if my recorded data is accurate?

Yes, the `validate_paddling_stats` tool checks if your duration, cadence, and total strokes are physically consistent.

07 What information is included in the session summary?

The `get_session_summary` tool provides the total strokes, average cadence, session duration, and the calculated intensity level.

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>"paddle-stroke-counter": { "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

Paddle Stroke Counter 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 Paddle Stroke Counter. 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	Paddle Stroke Counter MCP
Server ID	01a0cd45-c20f-7271-9a87-da7601f1b4da
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/paddle-stroke-counter.