MCP SERVER | NO CODE | CLOUD HOSTED

Interoception Training Protocol MCP

Build structured plans to improve your physiological awareness.

Interoception Training Protocol MCP gives your AI agent the framework to design and manage personalized physiological regulation plans. It handles the math for training intensity, tracks your readiness to advance to harder levels, and suggests quick exercises to help you apply body awareness to your daily routine.

A+ Quality Score 100/100

interoception | body-awareness | wellness | physiology | mindfulness

MAINTAINED BY VINKIUS CLOUD FOR AI AGENTS · SEPTEMBER 2026
vinkius.com/en/ai-agent-connect/interoception-training-protocol

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.

Interoception Training Protocol MCP

4 tools available
Cloud-hosted on Vinkius

You can use this MCP to turn your AI client into a specialized coach for physiological regulation. Instead of generic mindfulness advice, you get structured protocols tailored to your specific awareness levels and goals. The system manages the technical side of your progress, calculating how much cognitive or physical effort a session requires so you don't overextend yourself. It also monitors your consistency and accuracy to decide when you're actually ready to move from foundational exercises to more advanced training tiers. You can ask your agent for quick, low-effort exercises to use during work or busy periods, making it easier to integrate these skills into your actual life rather than just treating them as a separate task. It's a way to move from vague wellness goals to a data-driven training regimen for your nervous system.

Core Capabilities

01 — Custom Plan Generation

Your agent uses this to build specific training schedules based on your unique physiological goals.

02 — Intensity Management

The AI calculates the effort required for sessions to keep your training within safe limits.

03 — Progression Tracking

Your agent evaluates your performance data to determine when to increase training difficulty.

04 — Real-world Application

The system suggests quick exercises that fit into your existing daily schedule.

One Click on Vinkius — From Prompt to Execution

Available at **vinkius.com/en/ai-agent-connect/interoception-training-protocol** — connect your AI agent in three steps.

01 Connect the MCP to your preferred client like Claude or Cursor.

02 Ask your agent to generate a protocol based on your goals.

03 Use the agent to calculate the load for upcoming sessions.

04 Request daily exercises to apply your skills in real time.

05 Check your readiness for new tiers using your performance data.

You connect the MCP to your AI client and start requesting specific training actions.

Built For

This MCP is for individuals looking to move beyond general mindfulness into structured physiological training. It's built for those who want an agent to manage the complexity of training loads and progression.

Wellness Practitioners

They hand off the creation of structured, tiered training protocols to the AI.

Biofeedback Enthusiasts

They use the agent to track readiness and manage training intensity based on performance.

High-performance Professionals

They use the tool to get quick, low-effort exercises to maintain awareness during work.

What Changes When You Connect

01 Calculates training load to prevent overexertion.

02	Automates the decision to advance to higher training tiers.
03	Provides actionable exercises for use during daily tasks.
04	Creates structured plans based on specific awareness levels.

Real-World Applications

Stress Regulation Training

Generate a specific plan for someone starting with low awareness to help manage stress through regular practice.

Workday Integration

Get a 30-second exercise to perform while typing to maintain physiological awareness during a busy shift.

Tier Advancement

Use your consistency and accuracy data to see if you can move from foundational to advanced training.

Load Management

Check the estimated effort of a new training session to ensure it doesn't exceed your current capacity.

Patterns to Avoid

Ignoring training load limits

AVOID

You decide to jump into a high-intensity session because you feel motivated, ignoring how much mental energy you actually have left.

✓ INSTEAD

Check the estimated intensity first by using `calculate_training_load` to ensure you don't burn out.

Guessing your readiness to advance

AVOID

You try to move to a harder training tier just because you've been practicing for a week, even if your accuracy is still low.

✓ INSTEAD

Let the agent decide when you've actually hit the benchmarks by using `evaluate_progression_readiness`.

Treating training as a separate chore

AVOID

You try to schedule hour-long sessions during your busiest work hours, which makes you quit after two days.

✓ INSTEAD

Use `suggest_daily_integration_exercise` to find quick, low-effort movements that fit into your actual workday.

The Right Fit

Connect this MCP if you want a data-driven way to train your nervous system rather than following generic wellness advice. Do not connect it if you are looking for casual mindfulness apps. The deciding factor is whether you want an agent to manage your training intensity and progression for you.

Interoception Training Protocol Tools

These tools handle the heavy lifting of calculating training loads, managing progression, and generating personalized physiological exercises.

#	TOOL	DESCRIPTION
01	`calculate_training_load`	This tool estimates the cognitive and physiological effort needed for a specific session. It helps you avoid burnout by quantifying the intensity of your planned practice.
02	`evaluate_progression_readiness`	This tool checks if you have met the necessary benchmarks to move to a more advanced training tier. It looks at your current level and performance metrics to decide when to level up.
03	`generate_protocol`	This tool creates a personalized interoception training plan. It builds a structured schedule based on your specific goals and current awareness levels.
04	`suggest_daily_integration_exercise`	This tool provides specific, low-effort exercises for real-world use. These are designed to help you apply your training during normal daily activities.

See It in Action

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

U Generate a training protocol for someone with low awareness who wants to focus on stress regulation, practicing 3 times a week.

Your personalized protocol is 'Foundational Somatic Scanning'. It includes 3 sessions per week focusing on broad body scans to build baseline awareness.

U I am currently doing Passive Awareness. Suggest an exercise I can do during work.

Try the 'Micro-Respiration Check': take 30 seconds to notice the sensation of air entering your nostrils while you type.

U Am I ready for Tier 2? My current level is Tier 1, my consistency is 0.9, and my accuracy is 0.85.

Yes, you are ready to move to Tier 2 training.

Frequently Asked Questions

01 **What can this MCP do for my training?**

It allows your AI agent to generate personalized training protocols, calculate session intensity, and track your readiness to move to advanced levels.

02 **How do I use this with Claude or Cursor?**

You connect the MCP through the Vinkius platform. Once connected, your AI client can immediately use the tools provided by this MCP.

03 **Can I get exercises for when I am busy?**

Yes, you can ask your agent to suggest daily integration exercises that are designed to be low-effort and easy to do during work or other activities.

04 **How does the progression work?**

The MCP uses tools to evaluate your consistency and accuracy to determine if you are ready to move to a more advanced training tier.

05 **Does this help prevent burnout?**

Yes, the calculate_training_load tool estimates the cognitive and physiological effort required for a session to help you manage your practice intensity.

06 **How do I start a training plan?**

You can use the `generate_protocol` tool by providing your current awareness level, your primary goal, and how many times per week you wish to practice.

07 **Can I move to more advanced exercises?**

Yes, you can use `evaluate_progression_readiness` to check if your consistency and accuracy scores meet the requirements for a higher training tier.

08 **How do I prevent over-training?**

Use the `calculate_training_load` tool to estimate the cognitive and physiological effort of your planned session to ensure it stays within safe limits.

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 `"interoception-training-protocol": { "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

Interoception Training Protocol 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 Interoception Training Protocol. 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	Interoception Training Protocol MCP
Server ID	`01a0b6c0-69ee-70f8-afde-be358a1cfbf2`
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/interoception-training-protocol.