MCP SERVER | NO CODE | CLOUD HOSTED

Deep Sleep Estimator MCP

Turn sleep data into biological recovery metrics.

Deep Sleep Estimator MCP uses biological sleep modeling to calculate the restorative value of your rest. By looking at sleep duration, timing, age, and activity, your AI client can determine deep sleep percentages, physical recovery scores, and circadian alignment to help you understand how well you actually recovered overnight.

A+ Quality Score 100/100

sleep | recovery | health-tech | circadian | wellness

MAINTAINED BY VINKIUS CLOUD FOR AI AGENTS · SEPTEMBER 2026
vinkius.com/en/ai-agent-connect/deep-sleep-estimator

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.

Deep Sleep Estimator MCP

4 tools available
Cloud-hosted on Vinkius

You can use this MCP to move past simple sleep duration tracking and actually understand the biological quality of your rest. Instead of just seeing how many hours you slept, your agent can model the distribution of Slow Wave Sleep (SWS) to see if your rest was actually restorative. It takes your age, activity levels, and sleep timing into account to provide a much more accurate picture of physical recovery.

Your AI client can break down whether your deep sleep was concentrated or spread thin, helping you identify if your sleep was fragmented. It also checks how well your sleep schedule aligns with your biological needs, giving you a clear look at your circadian efficiency. Whether you are trying to optimize training recovery or just fix a messy sleep schedule, this MCP provides the specific biological metrics needed to make sense of your sleep data.

Core Capabilities

01 — Sleep Distribution Modeling

Your agent calculates the percentage and concentration of deep sleep.

02 — Recovery Scoring

The MCP determines physical recovery scores based on activity and sleep quality.

03 — Circadian Alignment Checks

Your client evaluates how well your sleep timing matches your biological needs.

04 — Sleep Density Analysis

The tool identifies if sleep was dense or fragmented.

One Click on Vinkius — From Prompt to Execution

Available at **vinkius.com/en/ai-agent-connect/deep-sleep-estimator** — connect your AI agent in three steps.

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

02 Provide your sleep data, including duration, timing, and age.

03 Include your activity levels for more accurate recovery scoring.

04 Your agent runs the specific tool to calculate your metrics.

05 Receive a detailed breakdown of your sleep distribution and recovery.

Connecting this MCP to your AI client gives you immediate access to biological sleep modeling tools.

Built For

This MCP is built for people who want to move beyond basic sleep tracking and into biological modeling.

Athletes

They hand off sleep and activity data to the AI to monitor physical recovery for training.

Biohackers

They use the biological modeling to optimize circadian alignment and sleep density.

Health Researchers

They use the specific metrics like SWS distribution to analyze sleep patterns.

What Changes When You Connect

01 Calculates physical recovery scores using activity and sleep data.

02 Identifies sleep fragmentation through density analysis.

03	Provides circadian alignment scores based on age and sleep timing.
04	Models Slow Wave Sleep distribution for better restorative insights.

Real-World Applications

Training Recovery

An athlete uses the MCP to see if their high activity level today requires more deep sleep for recovery.

Circadian Optimization

A user checks their sleep timing efficiency to find their optimal biological sleep window.

Sleep Quality Audits

You analyze sleep density to see if your sleep is being interrupted or fragmented.

Age-Based Sleep Adjustments

The AI adjusts sleep expectations and timing recommendations based on your specific age.

4 Tools Available

#	TOOL	DESCRIPTION
01	analyze_sleep_density	This tool determines if your deep sleep was concentrated or fragmented. It shows whether your restorative sleep was dense or spread thin throughout the night.
02	calculate_recovery_impact	This tool isolates the specific restorative value of a sleep session. It calculates how much your sleep actually contributed to physical recovery.
03	estimate_sleep_distribution	This tool provides a full breakdown of deep sleep characteristics. It gives you a detailed look at how your deep sleep was distributed for any given night.
04	get_sleep_timing_efficiency	This tool evaluates how well your sleep timing aligns with your biological needs. It checks your circadian alignment based on when you actually went to bed.

See It in Action

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

U I slept for 480 minutes starting at 23:00. I am 30 years old and had high activity today. What was my sleep distribution?

Your sleep distribution included a deep sleep percentage of 18% with a high concentration in the first half of the night. Your physical recovery score was 72/100 due to your high activity level.

U How efficient was my sleep timing if I went to bed at 01:00 and I am 45 years old?

Your circadian alignment score is 45/100. For your age, the optimal window to start sleep for maximum efficiency is 22:30.

U Calculate the recovery impact for 90 minutes of deep sleep for a 25-year-old with moderate activity.

The total restorative contribution is 85 units, with a recovery efficiency of 0.94.

Frequently Asked Questions

01 **What is the difference between sleep duration and sleep distribution?**

Sleep duration is just how long you slept. Sleep distribution, which this MCP models, looks at how your deep sleep (SWS) is spread throughout that time.

02 **How does age affect the sleep calculations?**

The MCP uses your age to determine optimal sleep timing efficiency and to help model how your deep sleep should ideally be distributed.

03 **Can I use this with any AI client?**

Yes, this MCP works with any MCP-compatible client like Claude, Cursor, or Windsurf.

04 **Does this MCP require me to host it?**

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

05 **How is physical recovery calculated?**

The MCP calculates recovery scores by looking at your sleep duration, the amount of deep sleep, and your recent activity levels.

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 `"deep-sleep-estimator": { "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

Deep Sleep Estimator 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

DOCUMENT INFORMATION

Generated	September 2026
MCP Server	Deep Sleep Estimator MCP
Server ID	`01a0b8e3-9bbd-7101-9ce6-800ed59afba9`
Platform	Vinkius Cloud for AI Agents
Endpoint	`https://edge.vinkius.com/{token}/mcp`

LICENSE & USAGE