

MCP SERVER

NO CODE

CLOUD HOSTED

FPS Hardware Bottleneck Estimator MCP for AI Agents

Predicting Gaming Framerates and Pinpointing Component Limitations

The FPS Hardware Bottleneck Estimator is a specialized simulation MCP for PC gamers. It predicts your expected game framerates and identifies exactly where performance limits exist—whether it's your CPU or your GPU. You can model how different components interact at specific resolutions like 1080p, 1440p, or 4K. Stop guessing if a new graphics card is worth the money; this MCP gives you concrete data on average FPS and critical 1% lows.

A+ Quality Score 100/100

fps

benchmark

gpu

cpu

bottleneck

pc-gaming



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.

FPS Hardware Bottleneck Estimator MCP

3 tools available

Cloud-hosted on Vinkius

Want to know what frame rates your rig will hit in Cyberpunk 2077 at 4K? This connector simulates that using deep performance analysis. It runs calculations based on how tightly coupled your CPU and GPU throughput are, instantly telling you if your system is bottlenecked by the processor or the graphics card. Instead of buying parts hoping for a boost, you feed in your components and get an expected average FPS range. You can even simulate upgrading one part to see its precise impact before spending a dime. Because Vinkius hosts thousands of specialized connections, this MCP gives you access to detailed gaming performance modeling that just wasn't available before. It's the data point every serious PC builder needs.

Core Capabilities

01 — Estimate expected FPS and bottleneck status

The MCP calculates the predicted average frames per second (FPS) and identifies whether your system is CPU-bound or GPU-bound for a specific game scenario.

02 — Simulate component upgrades

It predicts how replacing one specific part, like the GPU or RAM, will change your overall gaming performance metrics.

03 — Determine optimal hardware upgrade paths

The system analyzes your current setup against various goals to recommend which single component upgrade will yield the greatest performance increase for your titles.

One Click on Vinkius — From Prompt to Execution

Available at vinkius.com/mcp/fps-hardware-bottleneck-estimator — connect your AI agent in three steps.

- 01 Input your computer's full specifications, including CPU model, GPU model, and desired game settings (e.g., Cyberpunk 2077 at 1440p).
- 02 The MCP runs a specialized simulation that calculates the interaction between your primary components under those specific load conditions.
- 03 You receive predicted performance data, showing expected average FPS ranges and pinpointing the exact hardware component causing any bottleneck.

The bottom line is you get objective performance numbers so you know exactly where to spend your money on upgrades.

Built For

This MCP is for serious PC builders, competitive gamers, and tech enthusiasts who refuse to guess at hardware purchases. If wasting hundreds of dollars on a component that won't help feels like a possibility, you need this.

PC Gamer

Needs to know if their current setup can handle the latest AAA title at 4K without sacrificing frame rates.

System Builder

Uses this MCP to stress-test component combinations and identify the most cost-effective upgrade path for a specific budget.

Tech Enthusiast / Streamer

Wants empirical data on performance degradation when adding streaming encoders or running background tasks while gaming.

What Changes When You Connect

- 01 Stop guessing if new gear is worth it. Use `estimate_performance` to get concrete FPS numbers before you buy a single part.

-
- 02 Maximize your budget with `identify_optimal_path`. Instead of upgrading everything, this MCP tells you the one component that will give you the biggest boost for your money.

 - 03 `simulate_upgrade` lets you test hypothetical hardware swaps—like jumping from an RTX 3060 to a 4080—and see the exact percentage increase in FPS.

 - 04 Instantly determine your bottleneck status. It tells you if your CPU or GPU is limiting performance, saving you time and money on parts that won't help.

 - 05 Get detailed data for specific scenarios, whether it's running at 1080p, 1440p, or demanding 4K resolution.
-

Real-World Applications

My current GPU isn't enough for this new game.

A user feeds in their CPU and existing GPU details, then runs `estimate_performance` using the new game's specifications. The agent returns data showing a significant drop in 1% lows, confirming that an upgrade is necessary.

I have \$500 to spend on my PC build.

The user runs `identify_optimal_path` with their full system specs and budget. The MCP determines the CPU needs replacing first, yielding a 30% performance gain for less cost than upgrading the GPU.

I'm thinking of swapping my RAM or my Graphics Card.

Instead of buying parts blindly, the user runs `simulate_upgrade` to compare the impact of two different components. The results show that the GPU upgrade provides a much higher FPS increase than increasing the RAM.

Patterns to Avoid

Assuming more money means better performance

✗ AVOID

The user buys an expensive CPU and new memory without checking if their older, weaker GPU is already the limiting factor.

✓ INSTEAD

Run ``estimate_performance`` first. The MCP will tell you that your existing GPU is bottlenecking the system at 4K, meaning a better GPU is the only necessary upgrade.

Upgrading random parts to feel like an 'upgrade'

✗ AVOID

The user upgrades both their CPU and their RAM because they heard it was good, but the performance increase is negligible.

✓ INSTEAD

Use ``identify_optimal_path`` to force the system to recommend only one primary upgrade. It will point out that a specific GPU model offers the best bang for your buck.

Ignoring resolution impact

✗ AVOID

The user tests performance at 1080p and assumes everything is fine, but then plays a demanding game at 4K and runs into severe slowdowns.

✓ INSTEAD

Always run ``estimate_performance`` for the specific resolution you intend to play at. The MCP adjusts its calculations based on whether you are targeting high-refresh rates or pure graphical fidelity.

The Right Fit

Use this MCP if your primary goal is objective performance data and calculating component ROI (Return on Investment). You need it when you must know *why* your frame rate drops—is it the CPU, the GPU, or something else? If you're simply looking for general advice like 'get a better card,' this won't help. Don't use this if you just want to see marketing hype; use `estimate_performance` to get raw numbers instead. You also don't need it if your system is already balanced and running well across all major titles, but if you suspect one component is dragging the whole thing down, this MCP is essential.

Using FPS Hardware Bottleneck Estimator for Accurate Gaming Performance Benchmarks

Right now, determining why your PC slows down in a specific game feels like guesswork. You buy components based on reviews that test different games or resolutions, and when you play something new, the performance is unpredictable. It's all clicking through benchmarks and comparing vague specs.

With this MCP, you stop guessing. You feed in your exact parts list and select the game and resolution. The system runs a deep simulation to give you predicted FPS numbers—average FPS *and* 1% lows—so you know exactly what performance looks like under real-world stress.

Optimizing PC Builds with FPS Hardware Bottleneck Estimator Upgrade Recommendations

Manual upgrades involve a lot of risk. You might spend money on an expensive new GPU, only to find out the CPU can't keep up and the performance gain is minimal. It's costly trial and error.

This MCP solves that by running `identify_optimal_path`. It forces you to look at your entire system as one unit, pointing directly to the single component upgrade that will give you the most frames per dollar spent.

FPS Hardware Bottleneck Estimator: 3 Tools for PC Gaming Benchmarking

Run simulations to estimate your framerate, simulate component swaps, or find the absolute best upgrade path for any gaming setup.

#	TOOL	DESCRIPTION
01	<code>estimate_performance</code>	Runs a simulation to estimate the expected average FPS and identify if your system is bottlenecked by the CPU or GPU.
02	<code>simulate_upgrade</code>	Predicts how changing one component, like swapping out your graphics card, will affect your game performance metrics.
03	<code>identify_optimal_path</code>	Analyzes all potential hardware upgrades to recommend the single best part replacement for maximum performance gain.

See It in Action

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

U What FPS can I expect playing Starfield on an i7-12700K and RTX 3070 at 1440p?



Starfield Performance Estimate (1440p):

- **Average FPS:** 68 - 75 FPS
- **1% Lows:** 55 FPS
- **Bottleneck Status:** GPU-bound. The RTX 3070 is the limiting factor at this resolution and setting.

Recommendation: Upgrading to an RTX 4070 would provide a solid boost in both average and low FPS.

U If I upgrade my GPU, how much better will Warzone be? (Current setup listed)



Warzone Upgrade Simulation (1440p):

COMPONENT	CURRENT FPS (AVG)	NEW FPS (AVG)
GPU Swap: RTX 3070 → RTX 4080	~90 FPS	~135 FPS

Predicted Gain: A 55% increase in average frames. This confirms the GPU is your primary bottleneck.

U Which component should I upgrade first for Valorant at 1080p?



Optimal Path Analysis (Valorant, 1080p):

Based on your current system specs and target title, the most impactful single upgrade is replacing the Ryzen 5600 CPU.

- **Upgrade:** Ryzen 5600 to Core i5-13600K
- **Projected Gain:** 42% increase in average FPS.
- **Value Score:** Highest (Best balance of cost vs. performance jump).

Frequently Asked Questions

01 How does the FPS Hardware Bottleneck Estimator MCP work?

This MCP simulates how your CPU and GPU talk to each other when running a specific game at certain settings. It doesn't just give you one number; it predicts performance data like average frames per second, telling you exactly if hardware limitations are holding you back.

02 Can I use the FPS Hardware Bottleneck Estimator MCP for different games?

Yes. You specify the game and resolution you want to test. The simulation adapts its calculations to the specific graphical demands of whatever title you input, giving relevant performance data.

03 Is this better than just looking at benchmark scores online?

Absolutely. Online benchmarks are generic; they don't use your exact parts list or account for how different components interact in a single system. This MCP runs a personalized simulation based on *your* specific hardware.

04 How do I know which component to upgrade using the FPS Hardware Bottleneck Estimator?

You use the `identify_optimal_path` tool within this MCP. It analyzes all your parts and tells you, based on performance gain vs. cost, exactly which single component gives you the best bang for your buck.

05 Does this MCP help me decide between upgrading my CPU or my GPU?







Yes, that's its main job. By running simulations and using `estimate_performance`, it will point out whether your system is currently limited by the processor or by the graphics card.

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>"fps-hardware-bottleneck-estimator": { "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

FPS Hardware Bottleneck 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](https://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 FPS Hardware Bottleneck Estimator. 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	July 2026
MCP Server	FPS Hardware Bottleneck Estimator MCP
Server ID	019f2d2c-4708-7399-b886-5088c26eae09
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/mcp/fps-hardware-bottleneck-estimator.