

MCP SERVER

NO CODE

CLOUD HOSTED

# Math Evaluation Engine MCP

Stop LLMs from making math mistakes.

Math Evaluation Engine gives your AI client reliable arithmetic. It fixes the common problem where Large Language Models struggle with complex math, floating-point errors, or operator precedence. You can feed it complicated formulas—like  $(1/3) * 9$ —and get a guaranteed, deterministically precise answer every time, plus force specific decimal rounding for financial reports.

**F** Quality Score 43.65/100

floating-point

math-library

expression-evaluation

deterministic-logic

data-processing



# The infrastructure that powers AI agents in the real world.



Vinkius connects AI to the world's software through secure, enterprise-grade infrastructure — enabling real-world execution at scale, built on the Model Context Protocol (MCP).

# Your AI Connections Run Through Vinkius Cloud

The world's largest  
managed MCP catalog

Vinkius is the cloud infrastructure where AI agents connect 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](https://cloud.vinkius.com) — connect your AI agent in under 60 seconds.

# Math Evaluation Engine MCP

2 tools available

Cloud-hosted on Vinkius

You know the drill: you ask your agent to calculate something, and it gives you an answer that's close, but off by a hair.

LLMs are great at language, terrible at arithmetic. They mess up basic algebra, they struggle with operator order (like PEMDAS), and floating-point math is where they fall apart.

This MCP solves that. It completely separates the calculation from the conversation. Instead of letting your AI client guess the answer, it uses a robust engine to guarantee mathematically perfect results. You send it a complex formula as text, and you get back the exact number you need—no guesses, no rounding errors. If you're working with money or science data that needs absolute precision, this is mandatory. Through Vinkius, you connect this capability to your existing AI client and finally get math you can trust.

---

## Core Capabilities

### 01 — Calculate complex formulas

It safely evaluates any complicated mathematical expression written as text.

### 02 — Force decimal precision

You specify how many decimal places a number must have, and it rounds the result exactly to that count.

# One Click on Vinkius — From Prompt to Execution

Available at [vinkius.com/mcp/math-evaluation-engine](https://vinkius.com/mcp/math-evaluation-engine) — connect your AI agent in three steps.

- 01** Your AI client sends the MCP a string containing the math problem or formula.
- 02** The Math Evaluation Engine runs this input through its dedicated calculation engine, bypassing standard LLM arithmetic limitations.
- 03** You receive back the result: a precise number that adheres strictly to mathematical rules and requested rounding.

The bottom line is your agent stops hallucinating math results; it gets accurate, reliable numbers instead.

---

## Built For

Financial analysts who can't afford a single decimal point error. Scientific researchers running models that depend on pure calculation. Data engineers needing to validate complex formulas before they go live in production code.

### Quantitative Analyst

Calculating bond yields or risk factors, knowing the AI can guarantee precision for multi-step financial formulas.

### Scientific Researcher

Running simulations or statistical models where floating-point inaccuracies could invalidate months of work.

### Data Engineer

Developing validation logic that requires computing complex, nested mathematical expressions reliably before deployment.

---

## What Changes When You Connect

- 01** Accuracy is guaranteed. You don't have to trust the AI's internal calculation; this MCP offloads computation, eliminating guesswork for formulas like  $(1.5 * 3) / 0.2$ .

- 
- 02 Precision control. Use the ability to round values so you can force financial calculations or scientific metrics down to exactly two (or ten!) decimal places.

---

  - 03 Handles complexity. Forget basic addition; this MCP reliably processes operator precedence in complex, multi-step equations that usually crash standard language models.

---

  - 04 Speed and reliability. Because it runs directly within the client's environment, you get instant results without waiting on external API calls or latency.

---

  - 05 Formula validation. You can use `calculate_expression` to test any user-submitted equation instantly, verifying its mathematical soundness before committing to a workflow.
- 

---

## Real-World Applications

### Calculating complex investment returns

A financial analyst needs to calculate the compounded return on an investment using multiple variables and exponents. Instead of asking their agent, they use `calculate_expression` through this MCP to get a mathematically perfect rate, knowing the LLM wouldn't handle the exponentiation correctly.

### Validating code logic math

A data engineer writes Python code that relies on a specific mathematical formula. They use this MCP to test the formula in isolation, using `calculate_expression`, verifying it works exactly as intended before running the whole pipeline.

### Standardizing scientific data output

A researcher collects measurement data in various formats. They send the numbers and a required precision (e.g., 4 decimal places) to `round_value`. The MCP ensures every resulting number is standardized for consistent analysis across all models.

### Handling currency formatting

A content writer drafts a report that includes several calculated figures. Instead of having the agent guess the correct rounding for dollar amounts, they invoke `round_value` to force every number to two decimal places, ensuring compliance.

---

# Patterns to Avoid

---

## Asking the LLM directly

### X AVOID

Prompting your agent: 'What is  $0.1 + 0.2$ ?' and getting a result like '0.3' when it should be exactly '0.3'. The model simply struggles with binary floating-point representation.

### ✓ INSTEAD

Use the `calculate_expression` tool. You send the formula string, and the dedicated engine returns the precise, deterministic answer every single time.

---

## Manual spreadsheet rounding

### X AVOID

Calculating a complex metric in Excel, then manually copying and pasting it into an agent prompt, only to have the formatting or context break the formula.

### ✓ INSTEAD

Use `round_value` via this MCP. The calculation and the precise rounding happen within your agent's workflow, keeping the data clean and contained.

---

## Ignoring operator precedence

### X AVOID

Writing a complex query like ' $10 + 5 * 2 / 3$ '. An unassisted LLM might calculate this in the wrong order, giving you an incorrect final number.

### ✓ INSTEAD

Feed the entire calculation to `calculate_expression`. It respects standard math rules (PEMDAS) and gives you the single correct result.

---

## The Right Fit

Use this MCP if your workflow requires mathematical certainty. If you need to compute, round, or validate any number using formulas where 'close enough' isn't good enough—especially in finance, physics, or data modeling—this is the right tool. You must use it when dealing with floating-point math. Don't use this if your task is purely text generation (e.g., summarizing a document) or simple retrieval of facts. For general arithmetic help, you might just talk to your agent; but for anything requiring determinism, use `calculate_expression` or `round_value`. Never rely on the LLM itself for calculation.

---

## The frustration of math errors in AI outputs

You've asked your agent to calculate a complex metric. It returns an answer, but you immediately feel that gut twitch—that sense that the number is wrong. You spend time checking if it messed up the operator order or if the floating-point arithmetic threw off the final decimal point. Then you have to copy the formula into a separate spreadsheet just to verify the single piece of data.

With this MCP, you skip all of that. You hand the math problem directly to your agent's tools. The engine handles the complexity in isolation, giving you back an answer guaranteed by strict mathematical standards. What you get is absolute confidence in every number.

---

## Math Evaluation Engine delivers precision with `calculate_expression`

You don't have to manually verify formulas across multiple tabs or copy-paste results from a calculator just to get an answer your agent can use. The tool takes the raw formula string and outputs only the pure, correct number.

The result is clean: reliable math that integrates directly into your workflow without any messy workarounds.

---

# Math Evaluation Engine: 2 Tools Available

These two tools allow your AI agent to perform rigorous mathematical calculations, from solving complex expressions to precisely rounding floating-point values.

#	TOOL	DESCRIPTION
01	<code>calculate_expression</code>	This tool evaluates a complicated string of math, like <code>`(2^4 + 10) / 2`</code> , and returns the exact numeric answer.
02	<code>round_value</code>	It takes any number or expression and rounds it precisely to a specific count of decimal places, perfect for currency.

---

## See It in Action

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

**U** Safely calculate ``4.5 * (10 / 3)`` and avoid floating point inaccuracies.



✓ **Math Result:** 15

**U** Evaluate this complex user-submitted equation: ``(2^4 + 10) / 2``.



✓ **Math Result:** The evaluated result is exactly 13 .

**U** Round the floating point number ``14.55556`` down to exactly 2 decimal places.



✓ **Rounded Result:** 14.56

---

## Frequently Asked Questions

### 01 How does Math Evaluation Engine fix floating point errors?

It uses a specialized library designed for mathematical computation, completely bypassing the standard arithmetic limitations of LLMs. This ensures deterministic precision for all calculations.

### 02 Can I use `calculate_expression` for exponents?

Yes. You simply include exponentiation symbols in your string (e.g., ``2^4``) and the engine handles the calculation correctly, unlike a general-purpose language model.

**03 What is the difference between `calculate_expression` and `round_value`?**

`calculate_expression` evaluates an entire formula (e.g., `(10/3) + 5`), while `round_value` specifically takes a number or result and adjusts it to a fixed decimal count.

---

**04 Does Math Evaluation Engine work with currency values?**

Absolutely. Use the MCP's rounding capability (`round_value`) to force all outputs into two decimal places, which is perfect for handling monetary calculations consistently.







---

# Go Live in 60 Seconds

Get your connection token from [cloud.vinkius.com](https://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
 <b>Claude AI</b>	Profile → Customize → Connectors → "+" → Add custom connector → Paste endpoint
 <b>Cursor</b>	Settings → Features → MCP Servers → "+ Add New MCP Server" → Type: SSE → Paste endpoint
 <b>VS Code</b>	Ctrl/Cmd+Shift+P → "MCP: Add Server" → add <code>"math-evaluation-engine": {   "url": "..." }</code>
 <b>Windsurf</b>	MCP Settings → <code>mcp_settings.json</code> → Add endpoint URL
 <b>ChatGPT</b>	Settings → Tools & plugins → Add MCP server → Paste endpoint
 <b>Gemini</b>	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

# Math Evaluation Engine 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](https://vinkius.com) · [support@vinkius.com](mailto: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 Math Evaluation Engine. 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	June 2026
MCP Server	Math Evaluation Engine MCP
Server ID	019e38be-bb00-7224-851c-7301003b3087
Platform	Vinkius Cloud for AI Agents
Endpoint	<a href="https://edge.vinkius.com/{token}/mcp">https://edge.vinkius.com/{token}/mcp</a>

### 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/math-evaluation-engine](https://vinkius.com/mcp/math-evaluation-engine).