

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

# Expense Growth Rate MCP

Track and analyze spending changes through your AI client.

Expense Growth Rate MCP lets your AI agent handle the math behind your spending fluctuations. Instead of manually calculating percentage shifts between budget periods, you can ask your AI to find exact growth rates, generate descriptive reports, or categorize the intensity of cost changes across multiple expense pairs at once.

**A+** Quality Score 100/100

expenses

growth

percentage

spending

analysis



# 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](https://cloud.vinkius.com) — connect your AI agent in under 60 seconds.

# Expense Growth Rate MCP

4 tools available

Cloud-hosted on Vinkius

You can stop manually calculating percentage shifts between budget periods. This MCP gives your AI client the specific math tools it needs to interpret spending trends. If you need to know exactly how much a specific category jumped from last month to this month, your agent can pull that number instantly. You can also ask it to look at a long list of historical changes to find where your spending spiked or dropped most significantly. It doesn't just give you raw numbers, either. You can ask for a written summary of a change or have the AI categorize a price hike as significant or minor based on its magnitude. It turns raw expense data into actionable insights without you having to touch a calculator.

---

## Core Capabilities

### 01 — Percentage Calculation

Your AI client uses this to find the exact delta between two spending periods.

### 02 — Trend Summarization

The agent converts raw percentage shifts into readable text for your reports.

### 03 — Volatility Detection

Your agent identifies the most extreme spending spikes and drops in a dataset.

### 04 — Magnitude Categorization

The AI sorts growth rates into tiers to show how meaningful a change is.

# One Click on Vinkius — From Prompt to Execution

Available at [vinkius.com/en/ai-agent-connect/expense-growth-rate](https://vinkius.com/en/ai-agent-connect/expense-growth-rate) — connect your AI agent in three steps.

- 01 Connect the MCP to your client through the Vinkius dashboard.
- 02 Provide your expense data or specific amounts to your AI agent.
- 03 The agent calls the necessary tool to calculate or categorize the change.
- 04 Receive the calculated rate, summary, or tier directly in your chat interface.

Connecting this MCP to your AI client gives it immediate access to financial math tools.

---

## Built For

This MCP is built for anyone managing budgets or analyzing financial trends through an AI interface.

### Financial Analysts

Hand off historical spending data to the AI to find growth outliers.

### Accountants

Use the agent to quickly summarize period-over-period expense changes.

### Operations Managers

Ask the AI to categorize the intensity of cost increases across different departments.

## What Changes When You Connect

- 01 Eliminates manual percentage math for budget comparisons.
- 02 Converts raw numbers into descriptive text summaries.
- 03 Identifies extreme spending fluctuations within large datasets.
- 04 Categorizes cost changes into meaningful tiers.

---

## Real-World Applications

### Monthly Budget Reviews

Ask your agent to compare this month's utility costs against last month's to see the percentage shift.

### Expense Spike Detection

Run a batch analysis on a year of data to find which months had the highest spending growth.

### Reporting Automation

Have the AI generate a written summary of spending changes to paste directly into a slide deck.

### Cost Intensity Assessment

Use tier categorization to determine if a vendor price increase is minor or significant.

---

## Patterns to Avoid

### Manual math requests

#### ❌ AVOID

You ask your agent to manually calculate the percentage difference between two large sets of numbers in a chat window.

#### ✓ INSTEAD

Let the agent use ``get_growth_rate`` to get the exact percentage change instantly without manual calculation errors.

---

### Vague trend inquiries

#### ❌ AVOID

You ask your agent if spending looks okay this month compared to last month.

#### ✓ INSTEAD

Use `'batch_growth_analysis'` to scan your entire history and pinpoint exactly where the highest and lowest growth points occurred.

---

### Ignoring change magnitude

#### ❌ AVOID

You treat a 2% increase and a 25% increase as the same type of budget shift.

#### ✓ INSTEAD

Run `'compare_expense_tiers'` to assign a qualitative tier to each change so you know which spikes actually matter.

---

## The Right Fit

Connect this MCP if you manage budgets and need to interpret spending trends through an AI interface. Do not connect it if you only need to store raw expense data without analyzing changes. The decision depends on whether you want your agent to perform the math or just hold the numbers.

---

# Expense Growth Rate: 4 Specialized Math Tools

These tools handle the heavy lifting of calculating, categorizing, and summarizing spending shifts between budget periods.

| #  | TOOL                  | DESCRIPTION                                                                                                   |
|----|-----------------------|---------------------------------------------------------------------------------------------------------------|
| 01 | batch_growth_analysis | This tool scans a list of historical expense changes to identify the highest and lowest growth points.        |
| 02 | compare_expense_tiers | This tool assigns a qualitative tier to a growth rate to show how intense a change actually is.               |
| 03 | get_growth_rate       | This tool calculates the exact percentage increase or decrease between a previous and current expense amount. |
| 04 | get_growth_summary    | This tool generates a text-based interpretation of a growth change for use in reports.                        |

---

## See It in Action

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

**U** What was the growth rate if last month's expense was 100 and this month's is 125?



The growth rate is 0.25, which represents a 25% increase.

**U** Give me a summary of the change from 500 to 450.



Expenses decreased by 10%.

**U** Categorize the growth from 1000 to 1500.



The change is categorized as Significant.

---

## Frequently Asked Questions

### 01 What can this MCP do with my expense data?

It calculates percentage growth, summarizes changes in text, categorizes the intensity of shifts, and finds extremes in historical data.

### 02 Which AI clients can use this MCP?

You can use this with any MCP-compatible client like Claude, Cursor, Windsurf, or VS Code.

### 03 Do I need to host this myself?

No, Vinkius hosts the MCP and manages the connection for you.

**04 Can it handle multiple expense changes at once?**

Yes, the batch analysis tool allows your agent to evaluate multiple expense pairs in one go.

---

**05 How does it categorize growth?**

It uses a tiering mechanism to turn a raw percentage into a qualitative description of the change's magnitude.

---

**06 How do I calculate the growth rate?**

You can use the ``get_growth_rate`` tool by providing the base expense and the current expense.

---

**07 Can I analyze multiple expense changes at once?**

Yes, the ``batch_growth_analysis`` tool allows you to evaluate a list of historical expense pairs.

---

**08 What happens if the base expense is zero?**

The tools will return an error to prevent division by zero, as a base expense must be greater than zero.

---

# 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                                                                                  |
|----------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------|
|  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>"expense-growth-rate": {<br/>"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

# Expense Growth Rate 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 Expense Growth Rate. 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 | Expense Growth Rate MCP                           |
| Server ID  | 01a0c785-0498-739f-acf9-511abb846e57              |
| 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/expense-growth-rate](https://vinkius.com/en/ai-agent-connect/expense-growth-rate).