

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

Cohort Analysis MCP for customer data.

Track how user groups behave and stay over time.

Cohort Analysis MCP gives your AI client the ability to dig into customer lifecycle metrics. Instead of manually calculating churn or lifetime value, you can ask your agent to pull retention curves, measure decay rates, and adjust for seasonal shifts. It turns raw user data into actionable lifecycle insights through a few simple prompts.

A+ Quality Score 100/100

cohort

retention

ltv

churn

analytics



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.

Cohort Analysis MCP

4 tools available

Cloud-hosted on Vinkius

You can now give your AI client a direct line into your customer lifecycle metrics. This MCP lets you move past basic spreadsheets and ask your agent to perform complex cohort math. You can use it to see exactly how engagement drops off after a user joins or to figure out the total economic value a specific group brings to your business.

If you need to know if your growth is real or just a seasonal spike, the MCP handles the math. It can calculate decay rates to show how fast you are losing users and apply seasonal adjustment factors to normalize your data. This means you can tell if a cohort has actually reached a steady state or if you are just seeing a temporary trend. It's built to help you understand the velocity of churn and the long term value of your user base without you having to write a single line of analysis code.

Core Capabilities

01 — Retention Visualization

Your agent uses this to map out how engagement fades for specific user groups.

02 — Economic Value Tracking

The AI calculates the total revenue impact of specific user cohorts.

03 — Churn Velocity Measurement

Your agent determines how quickly a cohort is losing members.

04 — Seasonal Normalization

The tool adjusts data to account for predictable yearly shifts in behavior.

05 — Stability Detection

Your agent checks if churn rates have leveled off for a specific group.

One Click on Vinkius — From Prompt to Execution

Available at vinkius.com/en/ai-agent-connect/cohort-analysis — connect your AI agent in three steps.

- 01 Connect your AI client to the Vinkius hosted MCP.
- 02 Provide your data source or context to your agent.
- 03 Ask a specific question about a cohort's LTV, retention, or churn.
- 04 The agent executes the corresponding tool to perform the math.
- 05 Receive the calculated results or data trends directly in your chat.

You connect the MCP to your preferred client and start asking questions about your user data.

Built For

This MCP is for data-driven teams who need to monitor user health without manual calculation. You can hand off the heavy lifting of cohort math to your AI agent.

Growth Marketers

They ask the AI to check if new user cohorts are more or less valuable than previous ones.

Product Managers

They use the agent to monitor how feature updates affect long term retention.

Data Analysts

They offload repetitive LTV and decay rate calculations to the MCP.

What Changes When You Connect

- 01 Eliminates manual math for LTV and decay rate calculations.
- 02 Provides seasonal context to prevent misinterpreting data trends.

-
- 03** Turns raw cohort data into visual retention patterns via your AI client.
-
- 04** Identifies when user churn has reached a steady state.
-

Real-World Applications

Measuring User Stickiness

Ask your agent to pull a retention curve to see if users are sticking around after the first month.

Predicting Churn Trends

Monitor the decay rate to catch rising churn before it becomes a major problem.

Evaluating Marketing Spend

Use the LTV tool to see if the users acquired in a specific campaign are actually profitable.

Accounting for Holidays

Apply seasonal adjustment factors to see if a dip in engagement is just a holiday trend.

4 Tools Available

#	TOOL	DESCRIPTION
01	analyze_decay_rate	This tool measures how fast users are churning within a specific cohort.
02	calculate_cohort_ltv	Use this to find the total economic value a specific group of users generates.
03	get_retention_curve	This tool visualizes how engagement for a cohort drops off over time.
04	get_seasonal_adjustment_factor	This tool provides a multiplier to account for behavioral changes caused by the time of year.

See It in Action

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

U What is the retention curve for cohort C123 for monthly periods?



The retention curve for cohort C123 shows 100% retention in month 0, 45% in month 1, and 20% in month 2.

U Calculate the LTV for cohort C456.



The total lifetime value for cohort C456 is \$5,250, which averages to \$52.50 per user.

U Is cohort C789 stabilized?



Yes, cohort C789 has stabilized because the decay rate stayed consistent over the last three periods.

Frequently Asked Questions

01 What AI clients can I use with this MCP?

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

02 Do I need to host the MCP myself?

No, Vinkius hosts and manages the MCP for you. You just connect and start using the tools.

03 Can this MCP help with seasonal data?

Yes, it includes a tool to get seasonal adjustment factors so you can normalize your cohort data against yearly fluctuations.

04 How does it calculate user churn?

It uses the `analyze_decay_rate` tool to measure the velocity at which users leave a cohort.

05 Is it possible to see if a cohort has stabilized?

Yes, you can ask your agent to check if the decay rate has become consistent, which indicates a steady state.

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>"cohort-analysis": { "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

Cohort Analysis 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 Cohort Analysis. 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	Cohort Analysis MCP
Server ID	01a0a349-9c5a-720c-a859-5ad980d55baf
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/cohort-analysis.