

MCP SERVER

NO CODE

CLOUD HOSTED

Feature Adoption Analytics MCP

Turn raw usage data into clear product growth models.

Feature Adoption Analytics MCP gives your AI client the ability to quantify how users interact with new releases. Instead of staring at spreadsheets, you can ask your agent to calculate adoption velocity, estimate when a feature will hit saturation, or model how changes to onboarding might impact your long-term growth trajectories.

adoption

retention

metrics

product-growth

user-behavior



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.

Feature Adoption Analytics MCP

4 tools available

Cloud-hosted on Vinkius

You shouldn't have to manually calculate adoption curves every time you ship a new update. This MCP turns your AI client into a specialized product analytics engine. You can feed it raw usage numbers and immediately get answers about whether a feature is actually sticking or just generating temporary noise.

It helps you move past basic vanity metrics. You can determine if a specific release is actually driving user loyalty or if it's just a distraction. By modeling different growth scenarios, you can test how changes to discoverability might affect your future numbers before you even commit to a design change. It's built to bridge the gap between messy usage logs and the high-level insights you need to make product decisions.

Core Capabilities

01 — Adoption Forecasting

Your agent uses this to predict when a feature will reach its maximum user base.

02 — Retention Analysis

The AI uses this to measure the direct link between feature usage and user loyalty.

03 — User Tiering

Your client uses this to group users into power, regular, or casual categories.

04 — Growth Modeling

The agent uses this to simulate how changes to product variables impact future adoption.

One Click on Vinkius — From Prompt to Execution

Available at vinkius.com/en/ai-agent-connect/feature-adoption-analytics — connect your AI agent in three steps.

- 01
- Connect the MCP to your preferred client like Claude or Cursor via Vinkius.
- 02
- Provide your raw usage or retention data to your AI client.
- 03
- Ask your agent to run specific tools like `calculate_adoption_metrics`.
- 04
- Review the calculated trajectories, segments, or impact scores.

Connecting this MCP to your AI client gives you instant access to advanced product analytics.

Built For

This MCP is built for product teams who need to move fast without losing sight of user behavior.

Product Managers

You hand off raw usage data to the AI to get immediate adoption and retention insights.

Product Analysts

You use the tool to run complex growth simulations and segmentation models.

Growth Leads

You use the modeling tools to test how onboarding changes might affect long-term growth.

What Changes When You Connect

- 01
- Predicts feature saturation points based on current adoption velocity.
- 02
- Quantifies the specific retention lift provided by new features.
- 03
- Segments users into distinct tiers based on their actual usage frequency.

- 04** Models future growth by simulating changes to product variables.
-

Real-World Applications

Post-Launch Validation

Check if a new feature is actually being adopted as expected after the first few days of release.

User Segmentation

Identify your power users to better understand who is getting the most value from your updates.

Retention Audits

Determine if a specific feature is a key reason why users stay with your product.

Growth Projections

Run simulations to see how improving feature discoverability might change your adoption trajectory.

4 Tools Available

#	TOOL	DESCRIPTION
01	calculate_adoption_metrics	This tool provides a snapshot of current adoption progress and projects the feature's future trajectory.
02	evaluate_retention_impact	Use this to determine if a specific feature is a primary driver of user loyalty.
03	predict_usage_segmentation	This tool categorizes your user base into usage tiers to separate power users from casual ones.
04	simulate_growth_scenarios	This tool lets you model how adjusting product variables affects future adoption rates.

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 current adoption progress for my new feature with 1000 total users, 200 feature users, and 10 days since launch?



The current adoption rate is 20.0%, with an adoption velocity of 2.0% per day and an estimated saturation in 45 days.

U Is the new dashboard feature driving user loyalty? We have 500 feature users with 60% retention and 500 non-feature users with 45% retention.



The feature provides a retention lift of 0.15, which is categorized as a Critical impact on user loyalty.

U How many power users do we have if the usage frequency is: user_1: 50, user_2: 5, user_3: 1, user_4: 12?



Based on the usage frequency, you have 1 power user, 1 regular user, 1 casual user, and 1 one-time user.

Frequently Asked Questions

01 How do I use this MCP with Claude?

You connect to the MCP through the Vinkius platform. Once connected, your Claude client can immediately call the tools to analyze your product data.

02 What kind of data do I need to provide?

You provide usage counts, retention percentages, or user frequency logs. The AI uses this data to run the specific analytics tools.

03 Can I use this to predict future growth?

Yes. You can use the `simulate_growth_scenarios` tool to model how changing product variables will affect your future adoption.

04 Does this replace my existing analytics dashboard?

It acts as an intelligence layer on top of your data. It's designed to turn raw numbers into specific insights like retention lift or user segmentation.

05 Is my data secure?

Vinkius hosts and manages the MCP, ensuring your connection to your AI client is stable and secure.

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>"feature-adoption-analytics": { "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

Feature Adoption Analytics 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 Feature Adoption Analytics. 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	Feature Adoption Analytics MCP
Server ID	01a0a269-710d-71f8-8fe5-2a40866a22de
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/feature-adoption-analytics.