MCP SERVER | NO CODE | CLOUD HOSTED

AppLovin MCP for AI Agents

Analyze App Mediation Revenue and Ad Campaign Performance

AppLovin MCP connects your AI agent directly to your AppLovin and MAX advertising data. You instantly analyze monetization performance by querying revenue, impressions, eCPM, and user cohort trends using natural language prompts. Stop exporting dashboards; start asking questions.

A+ Quality Score 100/100

mobile-advertising | monetization | user-acquisition | ad-performance | revenue-tracking | mediation-platforms

MAINTAINED BY VINKIUS CLOUD FOR AI AGENTS · JUNE 2026
vinkius.com/mcp/applovin

SCAN TO VIEW

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.

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.

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.

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	Ed25519	24h TTL	V8 Isolate
Encryption at rest	PKI vault signatures	Ephemeral session keys	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.

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.

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.

06 — Honeytoken Trap System

Phantom credentials are injected into isolated environments. If a honeytoken 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.

Agent Activity

Which AI clients are connecting, how often, and what they're doing — real-time session tracking.

FinOps

Cost tracking per tool, payload compression savings, budget optimization signals, and consumption trends.

Tool Health

Slowest and most error-prone tools, with actionable root-cause insights and performance baselines.

Firewall & DLP

PII redaction activity, sensitive data protection counters, and security event timeline.

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.

AppLovin MCP

7 tools available
Cloud-hosted on Vinkius

This MCP gives your agent a direct line into your core ad platforms, pulling in crucial metrics from both AppLovin and MAX mediation layers. Instead of logging into separate dashboards or manually compiling spreadsheets, you simply ask your AI client for what you need to know about your app's revenue performance.

It handles complex data pulls like calculating aggregated performance by individual user or tracking long-term value through cohort analysis. You can also monitor the growth and spend on specific ad campaigns using AppDiscovery reports and list every active app in your account. It's a powerful way to get instant, detailed insights without leaving your chat window. All 4,000+ specialized integrations are cataloged at Vinkius, so this MCP makes sure you have real-time access to the data that matters most for mobile growth.

Core Capabilities

01 — Analyze Mediation Performance (MAX)

Retrieves aggregated advertising performance metrics from MAX mediation platforms, including total revenue and eCPM.

02 — Audit User-Level Revenue

Generates detailed reports showing ad revenue broken down either per specific user or per individual impression.

03 — Track Long-Term Retention (Cohort)

Runs MAX cohort reports to monitor how users retain value and what their long-term engagement looks like.

04 — Monitor Campaign Spend & Growth

Pulls performance data specifically for your User Acquisition campaigns tracked via AppDiscovery.

05 — Inventory Management

Lists all apps currently being tracked in your AppLovin account or shows a list of active advertising campaigns.

One Click on Vinkius — From Prompt to Execution

Available at **vinkius.com/mcp/applovin** — connect your AI agent in three steps.

01 Your AI client authenticates with the MCP, confirming connection using required management and report keys.

02 You ask your agent for specific ad performance data, such as 'What was my eCPM last week?' or 'Show me retention metrics for Q3.'

03 The MCP executes the necessary tool calls against AppLovin's APIs and returns structured, clean metrics directly to your AI client.

The bottom line is that instead of writing complex API queries, you just ask questions in natural language, and the system pulls the numbers for you.

Built For

This MCP is essential for AdOps Managers, UA Specialists, and Growth Engineers who are tired of spending hours manually exporting data from multiple advertising dashboards. If your job involves correlating ad spend to actual user revenue or figuring out why a cohort dropped off, you need this.

Ad Operations Manager

Audits monetization trends and eCPM performance across all mediation layers without manually downloading reports.

User Acquisition Specialist

Tracks campaign spend and overall growth metrics in AppDiscovery to determine which channels are driving the best return on ad investment.

Growth Engineer

Analyzes user-level revenue data and cohort reports to pinpoint retention weaknesses and optimize long-term app profitability.

What Changes When You Connect

01 Audit monetization performance immediately. Instead of exporting complex MAX data, you can ask your agent to pull specific eCPM trends or revenue totals using the `get_max_report` tool.

02 Pinpoint retention issues by running the `get_max_cohort_report`. You'll see Day 1 and Day 7 retention metrics instantly, helping you optimize for long-term user value.

03 Manage your ad spend accurately. Use AppDiscovery reporting to monitor campaign performance without digging through multiple dashboards; just ask for a summary of active UA campaigns.

04 Deep dive into profitability by running `get_user_ad_revenue_report`. This tool lets you see revenue broken down per user or per impression, which is impossible with standard reports.

05 Keep track of your assets. Use the `list_apps` and `list_campaigns` tools to get a current inventory check on every app and campaign linked to your account.

Real-World Applications

Determining Month-Over-Month Revenue Changes

An Ad Ops Manager needs to know if the recent ad placement changes are affecting revenue. They ask their agent for a comparative MAX report, which uses `get_max_report` to pull metrics for two different time periods and highlight the percentage change in eCPM.

Identifying User Drop-Off Points

A Growth Engineer notices that users are leaving after day three. They ask their agent to run a cohort report using `get_max_cohort_report` and quickly identify the exact Day 3 retention rate, allowing them to target a specific onboarding fix.

Checking Campaign Budget Health

A UA Specialist needs to confirm if all current ad campaigns are running optimally. They ask their agent to list active campaigns using `list_campaigns` and then request the performance data for a specific campaign via `get_app_discovery_report`.

Assessing Platform Health

A new team member needs an overview of all assets. They ask their agent to list all connected apps using `list_apps`, getting immediate confirmation of the entire app inventory tracked by the platform.

Patterns to Avoid

Relying on Dashboard Exports

✗ AVOID

Exporting a massive CSV file from an ad dashboard, then spending two hours in Excel trying to pivot and correlate eCPM changes against specific user cohorts.

✓ INSTEAD

Instead of exporting data, ask your agent for the analysis directly. Use `get_max_report` combined with natural language prompts to get the exact metric you need—like 'Show me eCPM trends over Q1'—in seconds.

Confusing UA Spend with Revenue

✗ AVOID

Mistaking general campaign spend reports for actual monetization performance, leading to bad investment decisions because the data isn't segmented correctly.

✓ INSTEAD

Always start by using `get_app_discovery_report` for spend tracking, but then cross-reference that against revenue metrics pulled via `get_user_ad_revenue_report` to see true ROI.

Ignoring User Segmentation

✗ AVOID

Calculating overall average revenue without knowing if high-value users are driving the numbers, leading to generic optimization efforts.

✓ INSTEAD

Use `get_user_ad_revenue_report` and specify 'per user' aggregation. This isolates the data for your most profitable segment.

The Right Fit

You should use this MCP if you need real-time, conversational access to highly specific ad metrics like eCPM trends, day-by-day retention rates, or per-user revenue breakdowns. It's perfect for auditing performance quickly and cross-referencing multiple data points without manual effort.

However, don't use it if your primary need is general business intelligence (BI) reporting across unrelated systems (like CRM data combined with ad spend). For those cases, you still need a dedicated BI platform. Also, if you only need to pull one static report that never changes, downloading the raw file might be faster. But if you're analyzing *why* a metric changed or comparing multiple metrics conversationally, this MCP is unmatched.

AppLovin MCP for AI Agents: Solving Ad Mediation Reporting Pain

Today, getting a full picture of your app's revenue means logging into MAX, downloading the report, opening Excel, and manually correlating eCPM spikes with specific user groups. It's tedious, prone to error, and takes hours just to build a simple comparison.

With this MCP, you simply tell your agent what you need—for example, 'Show me my revenue trend vs. last month'—and it executes the `get_max_report` tool instantly. You get an immediate, conversational answer that highlights exactly where the changes are.

AppLovin MCP for AI Agents: Optimizing User Retention Analytics

Normally, tracking how users retain value requires running a complex cohort report and then manually graphing Day 1 versus Day 7 retention rates to find patterns. This is critical data that often gets overlooked in daily reporting.

Now, you ask for the MAX cohort analysis directly through this MCP. The agent runs `get_max_cohort_report` and immediately gives you a clear breakdown of user stickiness, letting your team focus on fixing retention gaps instead of pulling reports.

7 AppLovin Tools for Ad Performance Metrics

Use these tools to calculate specific metrics, retrieve campaign performance reports, and audit your app's monetization assets.

#	TOOL	DESCRIPTION
01	`get_account_check`	Verifies that your AppLovin account connection is set up correctly for reporting.
02	`get_user_ad_revenue_report`	Pulls detailed ad revenue data, allowing you to aggregate metrics either per user or per impression.
03	`get_app_discovery_report`	Retrieves performance reports for User Acquisition campaigns running through AppDiscovery.
04	`get_max_cohort_report`	Generates cohort analysis reports specifically designed to track user retention within MAX mediation.
05	`get_max_report`	Gets aggregated performance data for your entire MAX mediation setup using specific date ranges and parameters.
06	`list_apps`	Provides a list of all applications currently being tracked within your AppLovin account.
07	`list_campaigns`	Fetches a complete list of active User Acquisition campaigns from the management API.

See It in Action

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

U How was the overall app revenue last week compared to the previous week?

MAX Performance Summary

The total revenue for the past 7 days was **$4,512.00**. This represents a slight increase of **3.1%** compared to the prior period's $4,370.00.

- **Average eCPM:** $3.98 (Up from $3.85)
- **Total Impressions:** 2,230,000 (Up from 2,160,000)

This increase suggests successful optimization in mediation placement.

U I need to check the performance for my top users. Can you get a report by individual user?

Here is the revenue breakdown for 5 high-value users identified in your account:

USER ID	TOTAL REVENUE (LAST 30 DAYS)	IMPRESSIONS COUNT
U1249	$85.50	1,200
U7782	$62.10	950
U3011	$45.00	750

These users account for a disproportionate amount of your revenue.

U Give me the cohort report to see how long users stick around from last month.

Cohort Retention Analysis (Last Month)

The data shows strong initial engagement, but retention drops sharply after Day 7.

- **Day 1 Retention:** 42%
- **Day 7 Retention:** 18%
- **Day 30 Retention:** 5.5%

You should investigate improving the onboarding flow to boost that Day 7 rate.

Frequently Asked Questions

01 **How does using the AppLovin MCP help me track ad revenue in my app?**

It pulls real-time, detailed data directly from your advertising platforms. You can ask for aggregated metrics like total eCPM or specific revenue totals without logging into any dashboard.

02 **Can I use the AppLovin MCP to check my user retention rates?**

Yes, this MCP runs dedicated cohort analysis reports. It shows you Day 1, Day 7, and longer-term retention metrics so you can understand how sticky your app is.

03 **Is the AppLovin MCP better than just exporting data from AppLovin?**

It's much faster. Instead of spending time in Excel cleaning up a CSV file, you ask natural language questions and get an immediate answer, making cross-platform comparison trivial.

04 **What kind of campaigns can I track with the AppLovin MCP?**

You can monitor your User Acquisition (UA) campaigns using AppDiscovery reports. This helps you see exactly where your ad spend is going and what performance it's generating.

05 **Does the AppLovin MCP work with other platforms besides MAX?**

Yes, because this MCP uses a secure set of tools that cover both AppLovin and the MAX mediation layer, giving you one unified view of your monetization performance.

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 `"applovin": { "url": "…" }`
Windsurf	MCP Settings → `mcp_settings.json` → 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

AppLovin 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 AppLovin. 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	AppLovin MCP
Server ID	`019d7550-cf49-7060-a5ee-bbf9f82838e1`
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/applovin.