MCP SERVER | NO CODE | CLOUD HOSTED

Kameleoon MCP

Manage your A/B testing and personalization from chat.

Kameleoon MCP lets your AI agent manage A/B testing and personalization workflows without logging into a dashboard. It gives you direct control over running experiments, checking audience segments, and requesting performance reports for any property or site within your account. Use it to automate the monitoring of live tests and analyze complex variations right from your development environment.

A+ Quality Score 98.33/100

ab-testing | personalization | feature-flagging | experimentation | audience-segmentation | conversion-optimization

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

SCAN TO VIEW

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.

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.

Kameleoon MCP

10 tools available
Cloud-hosted on Vinkius

Your AI client now controls your Kameleoon experimentation platform. This MCP lets agents handle everything from setup to analysis, managing A/B tests, audience segments, and personalization across all your digital properties. You can ask it to list every active experiment or pull the latest performance reports for a specific campaign. If you're used to bouncing between different dashboards just to check status updates, this changes that. Connect via Vinkius and let your agent do the heavy lifting. It reads site metadata, lists variations associated with an ID, and even requests asynchronous results reports so you can analyze experiment performance on the fly. You get full control of your testing lifecycle without ever leaving your chat interface.

Core Capabilities

01 — Initiate new experiments

Creates a brand-new A/B test with a specified name and site ID.

02 — Check current experiment status

Lists all active and past experiments to monitor campaign statuses.

03 — Inspect variations and segments

Retrieves details on specific A/B test variations or defined audience segments used for targeting traffic.

04 — Identify required sites

Queries your entire account to list all registered websites and properties, ensuring accurate environment targeting.

05 — Analyze performance reports

Triggers the retrieval of detailed, asynchronous results reports for any given experiment ID.

One Click on Vinkius — From Prompt to Execution

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

01 First, subscribe to this MCP on Vinkius and enter your unique Kameleoon Client ID and Secret credentials.

02 Next, prompt your AI client with a request, like 'List all active experiments for my checkout flow.'

03 Your agent sends the query through the MCP, receives structured data about experiment status or variations, and reports the actionable findings back to you.

The bottom line is, it lets your AI client talk directly to Kameleoon's backend tools using natural language instructions.

Built For

This MCP is for the Product Manager tired of manually checking 15 different dashboards just to know if a test passed. It's for the Growth Engineer who needs to query targeting rules directly from their IDE, and the Data Analyst who wants automated performance report retrieval without writing complex API calls.

Product Manager

Checks the status of live experiments and reviews variant details instantly by asking your agent instead of logging into the Kameleoon dashboard.

Growth Engineer

Quickly queries complex targeting rules or custom data dimensions right from their IDE, making environment setup faster.

Data Analyst

Automates the retrieval of experiment results and metadata. Instead of manual export/import, they request reports using a simple prompt.

What Changes When You Connect

01 Monitor live tests instantly. Instead of navigating to the Kameleoon dashboard, you can use `list_experiments` to see all active campaigns and their current status in a single prompt.

02 Targeting accuracy improves. You can quickly list defined audience segments or check specific targeting rules using `list_segments` and `list_targeting_rules`, ensuring your tests hit the right users every time.

03 Analyze performance data faster. When an experiment wraps up, you don't waste time setting up reports. Just prompt for results and use `get_experiment_results` to initiate fetching the latest metrics.

04 Know your environment at a glance. Need to confirm if a site is ready for testing? Use `list_sites` or `get_site` to verify all properties are correctly registered before building an experiment.

05 Manage complex variations simply. Instead of opening the test details page, you can use `list_variations` to see every single variant tied to an experiment ID, confirming setup instantly.

Real-World Applications

Debugging a broken personalization flow.

The Growth Engineer notices a segment isn't working correctly. They prompt the agent: 'List all segments and check the targeting rules for the checkout page.' The agent uses `list_segments` and `list_targeting_rules`, pinpointing that the rule is too restrictive, saving hours of debugging.

Getting a performance snapshot pre-meeting.

The Product Manager has a meeting about last quarter's testing. They ask: 'What are the results for the homepage test from May?' The agent uses `get_experiment_results` and pulls up the necessary data report, instantly ready to present.

Verifying site readiness.

The team is launching a new feature across multiple properties. A developer asks: 'List all sites in our account.' The agent uses `list_sites`, confirming that both the main e-commerce platform and the microservice landing page are registered and ready for testing.

Comparing test setups.

A PM wants to know what variations exist for a specific experiment ID. They prompt: 'What are the variations for campaign 456?' The agent uses `list_variations`, providing a clean list of all possible test groups instantly.

Patterns to Avoid

Manual dashboard hopping

✗ AVOID

Opening Kameleoon, navigating to Experiments > Select Experiment ID > Click Variations tab > Manually checking status. This takes 5-10 clicks and five minutes of context switching.

✓ INSTEAD

Instead, ask your agent directly: 'What are the variations for experiment ID X?' The agent uses `list_variations` in one step, giving you a clean list immediately.

Forgetting which sites exist

✗ AVOID

Attempting to create an experiment and forgetting to select the correct site ID from a dropdown menu, leading to deployment errors.

✓ INSTEAD

Before starting any test, prompt: 'List all registered sites.' Use `list_sites` to get every valid environment name and ensure your new test is scoped correctly.

Overcomplicating segmentation

✗ AVOID

Trying to remember the exact rule structure for a complex audience segment, leading to incorrect setup or manual consultation of documentation.

✓ INSTEAD

Just ask: 'List all defined audience segments.' The agent uses `list_segments` and shows you every available group name without needing to navigate deep into the platform.

The Right Fit

Use this MCP if your workflow revolves around monitoring, comparing, or reporting on multiple active A/B tests. It's perfect when you need an agent to act as a universal dashboard aggregator—a place where checking experiment status (`list_experiments`), verifying site readiness (`get_site`), and pulling performance reports (`get_experiment_results`) are all simple prompts.

Don't use this if your primary job is pure content creation or single-platform data entry. If you only need to read a basic piece of text from an external source, a generic API connector will suffice. You also don't need it if you're just managing users; then a CRM MCP would be better suited.

The pain of checking status across multiple testing platforms

Right now, checking on your tests means logging into the Kameleoon dashboard. You navigate to the 'Experiments' tab, then you have to click through active campaigns one by one. If you need results for five different experiments, that’s at least 25 clicks and a quarter hour of context switching just to gather the status update.

With this MCP, your agent handles it all. You simply ask, 'List all active experiments.' The system uses `list_experiments` to pull every current campaign's metadata into one clean response. You get the entire status report without opening a single browser tab.

Get immediate insights with Kameleoon MCP

The manual process of checking variations requires you to first find the experiment, then click into it, and finally review the list of variations. This is slow, error-prone work that burns valuable time before a launch.

Now, you just tell your agent which experiment ID you're interested in. The system uses `list_variations`, providing an immediate, structured breakdown of every variation associated with that test. It’s instantaneous insight.

Kameleoon MCP: 10 Tools for Experiment Management

These tools allow you to programmatically interact with Kameleoon, enabling your agent to create tests, check status, and pull reports directly via natural language prompts.

#	TOOL	DESCRIPTION
01	`create_experiment`	Creates a new A/B test when given a name and the site ID.
02	`list_custom_data`	Retrieves a list of available custom data dimensions for use in testing.
03	`list_experiments`	Lists all experiments within Kameleoon, allowing you to monitor various campaign statuses.
04	`get_experiment`	Fetches detailed information about a single, specific experiment.
05	`get_site`	Retrieves the full details for one particular site property.
06	`get_experiment_results`	Requests an asynchronous report containing the latest performance metrics for a finished experiment.
07	`list_targeting_rules`	Lists all predefined rules used to segment and target specific user groups.
08	`list_segments`	Provides an overview of all defined audience segments in the account.
09	`list_sites`	Gathers a complete list of every site registered across your entire account.
10	`list_variations`	Lists all specific A/B variations associated with a given experiment ID.

See It in Action

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

U Show me all active experiments in my Kameleoon account.

I've fetched your experiments. Currently, you have 4 active experiments, including 'New Checkout Flow (AB)' and 'Home Page Personalization'. Which one would you like to inspect?

U What are the variations for experiment ID '12345'?

For experiment 12345, I found 3 variations: 'Reference' (Original), 'Variant A (Red Button)', and 'Variant B (Blue Button)'. All variants are currently receiving traffic.

U List all sites registered in my Kameleoon profile.

I've retrieved 2 sites from your profile: 'Main E-commerce (Code: SEC-123)' and 'Marketing Blog (Code: BLOG-456)'.

Frequently Asked Questions

01 **How do I find out what experiments are running using the Kameleoon MCP?**

You simply prompt your agent to list all active tests. The system uses `list_experiments` to retrieve a comprehensive overview of every experiment currently set up in your account.

02 **Can I get results for an old test using the Kameleoon MCP?**

Yes, you can request performance metrics on past tests. Prompting for reports triggers `get_experiment_results`, which initiates fetching the necessary asynchronous data for analysis.

03 **What if I want to check a site ID before creating an experiment?**

Use the agent to fetch details about specific sites. The `get_site` tool retrieves all metadata for one property, confirming its status and readiness before you commit to running a test.

04 **How do I see which user groups are available for testing?**

You can list all defined audience segments using the `list_segments` tool. This gives you an immediate overview of every group ready for precise traffic allocation in your A/B tests.

05 **What is the difference between listing variations and getting experiment details with Kameleoon MCP?**

Getting experiment details (`get_experiment`) gives general metadata about the whole test. Listing variations (`list_variations`) specifically provides a breakdown of every individual testing group or variant attached to that single ID.

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 `"kameleoon": { "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

Kameleoon 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 Kameleoon. 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	Kameleoon MCP
Server ID	`019d75bf-170a-7067-a1fd-1704272880af`
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/kameleoon.