

MCP SERVER

NO CODE

CLOUD HOSTED

# VWO MCP

Control A/B tests and features via chat.

VWO MCP connects your AI client directly to Visual Website Optimizer. It lets you manage A/B tests, tweak feature flags, and analyze conversion goals through natural conversation. Get real-time insights on user behavior without leaving your chat window.

**A+** Quality Score 100/100

ab-testing

conversion-rate-optimization

feature-flags

personalization

experimentation

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

# VWO MCP

10 tools available  
Cloud-hosted on Vinkius

Managing web experiments used to mean jumping between dashboards, running reports, and cross-referencing spreadsheets just to see if a change actually worked. Now you can talk to your agent about it. This MCP lets you list active A/B tests and check their statistical results instantly. Need to know the status of a new payment gateway feature? Just ask; your AI client handles the flag management across staging or production environments. You can also pull lists of conversion goals, track which audience segments are performing best, and understand why users aren't completing key actions. With Vinkius, you connect once from any compatible agent, giving your AI a full view of your experimentation stack.

---

## Core Capabilities

### 01 — List and monitor all running tests

Retrieve comprehensive lists of every active A/B test and personalization campaign configured in VWO.

### 03 — Analyze experiment performance metrics

Pull detailed statistical results from any campaign, including identifying which variation is statistically significant or performing best.

### 05 — Review user segmentation groups

Get a list of pre-defined audience segments, helping you see exactly how different user groups are being targeted by your tests.

### 02 — Check feature flag status across environments

View the current configuration and toggle the enabled state for specific features (flags) across multiple system environments.

### 04 — Audit conversion tracking setup

Browse all defined conversion goals and key performance indicators (KPIs) to understand what success looks like for your site.

# One Click on Vinkius — From Prompt to Execution

Available at [vinkius.com/mcp/vwo](https://vinkius.com/mcp/vwo) — connect your AI agent in three steps.

- 01 Subscribe to this MCP and provide your VWO API Token and Account ID.
- 02 Connect your preferred AI client (like Cursor or Claude) through the Vinkius platform.
- 03 Ask your agent a question in natural language; it executes the necessary calls and returns structured data.

The bottom line is, you talk to your agent like talking to a coworker, and it handles all the complex API interactions for you.

---

## Built For

This is built for growth teams. If you're constantly toggling flags or running reports on different dashboards just to prove an idea worked, this MCP saves hours of context switching.

### Product Manager

Uses the MCP to manage feature rollouts and toggle flags across environments during a launch, ensuring stability before going live.

### Growth Manager

Monitors experiment results and statistical significance for active A/B tests without manual data exports or spreadsheet juggling.

### Data Analyst

Quickly surfaces conversion goal metrics, audience segment performance, and campaign details via simple chat commands.

---

## What Changes When You Connect

- 01 You stop manually exporting data. Instead of pulling reports for `get_campaign_results`, you simply ask your agent, and it returns the statistical significance right away.

- 
- 02** Feature rollouts get safer. Use `list_feature_flags` and `toggle_feature_flag` to manage flag status across different environments without logging into three different dashboards.
- 
- 03** Auditing is instant. Need to know what goals are tracked? Run `list_conversion_goals` in a chat prompt, getting a full inventory of KPIs instantly.
- 
- 04** Context switching ends. Listing all campaigns or checking audience segments via `list_optimization_campaigns` means you never have to open the VWO UI just for an overview.
- 
- 05** Better targeting control. By listing available user groups using `list_audience_segments`, your agent helps you confirm who exactly is seeing a test variation.
- 

---

## Real-World Applications

### A/B Test Review

The Growth Manager needs to know if the 'Checkout Flow Optimization' campaign worked. They ask their agent, and it uses `get_campaign_results` to report: Variation B is the winner with 98% statistical significance, showing a 14.5% lift.

### Pre-Launch Audit

The CRO Specialist needs to confirm which conversion goals are active before a major launch. Running `list_conversion_goals` provides an instant inventory of every KPI being tracked, ensuring nothing is missed.

### Emergency Feature Kill

The Product Manager realizes 'new\_payment\_gateway' is causing errors in production. They prompt their agent to `toggle_feature_flag` for that ID, immediately disabling the problematic feature across all users.

### Segment Verification

A Data Analyst wants to ensure the test only hits high-value users. They run `list_audience_segments` and confirm that the 'VIP Customer' segment ID exists, validating their targeting scope.

---

# Patterns to Avoid

---

## Trying to manually check flag status

### X AVOID

Opening VWO, going to Settings > Feature Flags, clicking through multiple environments (Staging, Production) just to confirm if the 'beta\_ui' is on.

### ✓ INSTEAD

Use the MCP's `list_feature_flags` tool first, then ask your agent to check its status or use `toggle_feature_flag` directly in a single conversation.

---

## Misunderstanding campaign scope

### X AVOID

Assuming that running a general report will show all test results when the test was run months ago and archived.

### ✓ INSTEAD

Always start by using `list_optimization_campaigns` to confirm the exact name or ID of the active campaigns you need data for before pulling specific results with `get_campaign_results`.

---

## Ignoring environment differences

### X AVOID

Toggling a feature flag in Staging, assuming that automatically updates the Production setting without confirmation.

### ✓ INSTEAD

Always check all configured environments first using `list_vwo_environments`, and specify the target environment when you run `toggle_feature_flag`.

---

## The Right Fit

Use this MCP if your workflow involves frequently comparing live data against configuration settings. You need to know *what* tests are running, *if* they're statistically significant, and *where* a feature is currently enabled across environments. Don't use it if you simply need general analytics—for that, look at raw web traffic tools. If your primary task is writing ad copy or generating hypotheses, stick to pure LLM chat models. But if the job requires auditing an existing experiment setup, managing flags, or retrieving concrete metrics from VWO, this MCP gives your agent all the necessary hands-on access through its specific tools like `get_campaign_results` and `toggle_feature_flag`.

---

---

## The Pain of Context Switching

Right now, checking on an experiment's status is a multi-step chore. You have to log into the VWO dashboard, navigate to the specific campaign, find the results tab, and then manually compare that data against the feature flag settings in a different section of the site. This means multiple clicks, opening separate tabs, and cross-referencing dates.

With this MCP, your agent handles all those steps for you. You just tell it what you need—like checking campaign results or listing active flags—and the information is returned immediately in plain text. It's like having a dedicated expert sitting next to you who knows exactly where every piece of data lives.

---

## VWO MCP Gives You Control Over Your Experimentation

Instead of manually running reports or clicking through pages to confirm a feature's status, you can now use the `list_optimization_campaigns` tool to get an instant overview. Then, if needed, you execute `get_campaign_results` without ever leaving your chat window.

This isn't just about retrieving data; it's about command execution. You control rollouts and audits by using tools like `toggle_feature_flag`, making the whole process as simple as typing a short instruction.

---

# VWO MCP: 10 Tools for Growth Optimization

These tools let you programmatically interact with every core function of VWO—from listing campaigns to toggling features—all through your AI client.

#	TOOL	DESCRIPTION
01	<code>get_vwo_account_info</code>	Retrieves basic information about the connected VWO account.
02	<code>get_campaign_details</code>	Fetches specific operational details for a single VWO campaign.
03	<code>get_feature_flag_details</code>	Retrieves the current configuration settings for one VWO feature flag.
04	<code>get_campaign_results</code>	Pulls statistical performance results for a specific VWO campaign, identifying winners and metrics.
05	<code>list_optimization_campaigns</code>	Lists all A/B test and personalization campaigns currently running in VWO.
06	<code>list_vwo_environments</code>	Provides a list of every configured VWO environment, like Staging or Production.
07	<code>list_feature_flags</code>	Lists all available feature flags managed within the VWO platform.
08	<code>list_conversion_goals</code>	Displays every conversion goal and measurable KPI that is currently being tracked by VWO.
09	<code>list_audience_segments</code>	Retrieves a list of all user segments defined for targeted advertising or testing.
10	<code>toggle_feature_flag</code>	Changes the status of a feature flag, either enabling it (true) or disabling it (false).

---

## See It in Action

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

### **U** List all my active A/B testing campaigns.



I found 4 active campaigns: 'Homepage Redesign' (ID: camp-101), 'Checkout Flow Optimization' (ID: camp-102), 'Pricing Page Test' (ID: camp-103), and 'Newsletter Signup Personalization' (ID: camp-104). Which one would you like to review?

### **U** What are the results for the 'Checkout Flow Optimization' campaign?



Results for camp-102: Variation B is the winner with 98% statistical significance. It shows a 14.5% improvement in conversion rate compared to the control. Variation A showed no significant change. Would you like me to pull more detailed metrics?

### **U** Disable the feature flag for 'new\_payment\_gateway'.



I've disabled the feature flag 'new\_payment\_gateway' (ID: feat-555) as requested. The status has been updated successfully across all environments. Is there anything else you need me to adjust?

---

## Frequently Asked Questions

### **01** How do I check if an A/B test is finished with VWO MCP?

You can use ``list_optimization_campaigns`` to see all current tests. For results, run ``get_campaign_results``, which provides the statistical significance needed to determine if a campaign should be marked as complete.

---

---

**02 Can I disable a feature flag using VWO MCP?**

Yes, you can use the `toggle_feature_flag` tool. You just need to provide the specific feature ID and tell your agent if you want it enabled (true) or disabled (false).

---

**03 Which tools do I use for conversion tracking in VWO MCP?**

To see what's tracked, run `list_conversion_goals`. If you need to know who is seeing the test, check out `list_audience_segments`.

---

**04 Does VWO MCP support multiple environments like staging and production?**

Absolutely. You can use `list_vwo_environments` to see all configured settings, giving you visibility across your entire deployment lifecycle.

---

**05 What if I want details on a specific campaign? Do I use VWO MCP?**

Yes. Use the `get_campaign_details` tool for general info, and then run `get_campaign_results` to pull the actual performance metrics.







---

# 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
 <b>Claude AI</b>	Profile → Customize → Connectors → "+" → Add custom connector → Paste endpoint
 <b>Cursor</b>	Settings → Features → MCP Servers → "+ Add New MCP Server" → Type: SSE → Paste endpoint
 <b>VS Code</b>	Ctrl/Cmd+Shift+P → "MCP: Add Server" → add <code>"vwo": { "url": "..." }</code>
 <b>Windsurf</b>	MCP Settings → <code>mcp_settings.json</code> → Add endpoint URL
 <b>ChatGPT</b>	Settings → Tools & plugins → Add MCP server → Paste endpoint
 <b>Gemini</b>	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

# VWO 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 VWO. 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	VWO MCP
Server ID	019d761d-edd4-7384-8815-a11bcb8e84d8
Platform	Vinkius Cloud for AI Agents
Endpoint	<a href="https://edge.vinkius.com/{token}/mcp">https://edge.vinkius.com/{token}/mcp</a>

### 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/vwo](https://vinkius.com/mcp/vwo).