MCP SERVER | NO CODE | CLOUD HOSTED

Multivariate Test Analyzer MCP for AI Agents

Finding Optimal Element Combinations in Conversion Rate Optimization

The Multivariate Test Analyzer performs complex 2k factorial analysis, letting you move past basic A/B testing. It identifies the optimal combination of multiple elements—like headlines, colors, and CTAs—and measures how they interact to maximize conversion rates.

A+ Quality Score 100/100

multivariate | factorial-design | ab-testing | conversion-optimization | statistical-analysis

MAINTAINED BY VINKIUS CLOUD FOR AI AGENTS · JULY 2026
vinkius.com/mcp/multivariate-test-analyzer

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.

Multivariate Test Analyzer MCP

3 tools available
Cloud-hosted on Vinkius

When your conversion rate optimization efforts get complicated, simple A/B tests fall short. This MCP handles 2k Factorial Design of Experiments (DOE), letting you analyze how several different factors impact performance all at once. Instead of just knowing which variant is 'better,' you learn *why* it's better and if certain elements only perform well when paired with others. You can isolate the direct effect of a headline, detect dependencies between button colors and text sizes, and pinpoint the single best configuration for your product pages. Connecting this through Vinkius gives your AI client the power to process these complex statistical models so you don't have to manually crunch data in spreadsheets. It’s about figuring out the true optimal setup.

Core Capabilities

01 — Calculate Main Effects

Determines the direct impact of a single factor on conversions, regardless of other factors.

02 — Analyze Interaction Effects

Detects statistically significant dependencies when two or more elements are combined (e.g., does 'Green' only work with 'Large Text?').

03 — Identify Winning Combination

Processes all test data to output the single most statistically optimal configuration for your experiment.

One Click on Vinkius — From Prompt to Execution

Available at **vinkius.com/mcp/multivariate-test-analyzer** — connect your AI agent in three steps.

01 Your AI client feeds the MCP three pieces of information: the list of factors being tested, the corresponding visits count, and the conversion results (conversions).

02 The MCP runs a 2k factorial analysis to calculate main effects and interaction effects across all element pairs.

03 It generates actionable insights, identifying which specific combination yields the highest expected conversion rate.

The bottom line is that it takes raw test data—visits and conversions for multiple elements—and outputs a statistically proven optimal configuration.

Built For

This MCP is built for CRO Specialists, Product Managers, and Data Scientists. If you're tired of running endless simple A/B tests that only tell half the story, this tool gets you to true optimization.

CRO Specialist

You use it to analyze complex test data, moving beyond basic comparison testing to find optimal element combinations across an entire user journey.

Product Manager

You rely on this MCP when launching a new feature. You feed it initial A/B results and ask your agent to predict the best setup before rolling out company-wide.

Data Scientist

You use it as a specialized statistical engine, feeding in factors and metrics (visits/conversions) to run formal 2k factorial designs of experiments.

What Changes When You Connect

01 You move past simple A/B testing. Instead of comparing two variants, you analyze dozens of possible combinations to find the true best-performing setup.

02 The MCP calculates main effects for every factor, letting you isolate whether a headline's performance is inherent or dependent on another element like button color.

03 Using `analyze_interaction_effects`, your agent finds dependencies—for example, knowing that 'Large Text Size' only works well when paired with the 'Green' button color.

04 The `identify_winning_combination` tool eliminates guesswork. It processes all data and delivers one statistically optimal configuration for immediate implementation.

05 You save time by automating complex DOE analysis. Instead of manual statistical modeling, your agent runs the full 2k factorial design in seconds.

Real-World Applications

A/B Testing a New Landing Page Layout

The marketing team ran tests on headlines, hero images, and CTAs but doesn't know the ideal mix. They ask their agent to run the Multivariate Test Analyzer. The MCP analyzes the data, detecting that 'Image B' only improves conversions when paired with a 'Benefit-driven headline,' giving them the exact optimal combination.

Optimizing Checkout Flow Elements

The product team wants to improve checkout conversion. They use `calculate_main_effects` to see if changing shipping options alone is enough, then run interaction effects to see if a specific trust badge only boosts sales when paired with expedited shipping.

Comparing Multiple Funnel Steps

A company needs to decide which combination of pricing presentation and form length works best. By feeding the data into this MCP, they can use `identify_winning_combination` to get a single, mathematically proven path for maximum sign-ups.

Patterns to Avoid

Treating all tests as simple A/B comparisons

✗ AVOID

Only testing Headline A vs. B, or Button Red vs. Blue, and declaring the winner based on a single comparison. This misses critical dependencies between elements.

✓ INSTEAD

Use this MCP's tools to run full factorial designs. Instead of comparing headlines alone, use `analyze_interaction_effects` to see how 'Headline A' performs specifically when paired with the 'Blue Button.'

Ignoring statistical significance

✗ AVOID

Assuming that because Variant B performed 3% better than Variant A in a small test, it is definitively superior. This ignores potential confounding variables and interaction effects.

✓ INSTEAD

Run the data through `calculate_main_effects` to ensure the impact of each factor meets statistical significance thresholds before declaring a winner.

Over-analyzing non-critical elements

✗ AVOID

Spending time determining if changing the footer copyright year has any measurable effect. This wastes resources on factors with zero meaningful variance.

✓ INSTEAD

Focus your inputs and use `identify_winning_combination` to constrain the analysis only to high-impact, core conversion levers.

The Right Fit

Use this MCP if your current testing approach is limited to simple A/B comparisons or if you suspect that element performance isn't linear. If you need to know how a headline interacts with button color, or how two separate elements combine to create a massive lift, this is the tool. Don't use it if you just need to compare two

variants side-by-side; for that, a basic A/B testing solution works fine. However, if your goal is true optimization—finding the single best combination of factors across multiple variables—you absolutely need the specialized factorial analysis provided by this MCP.

Multivariate Test Analyzer: Solving Complex Conversion Dependencies

Right now, running a large-scale experiment involves endless manual data aggregation. You collect visits and conversions for headlines, buttons, images, and CTAs across multiple test groups. Then you spend hours in statistical software trying to figure out if the headline improvement is real, or if it's actually just because you paired it with a high-performing button.

With this MCP, your agent handles that complexity automatically. You feed it the raw data, and it uses tools like `analyze_interaction_effects` to immediately map dependencies. The output isn't just 'good'; it tells you exactly *why* it's good—for instance, proving a specific headline combination generates a 15% lift.

Multivariate Test Analyzer: Pinpointing Optimal CRO Configurations

The biggest time sink is the iterative process of refining tests. You run a test, find one winner, declare it 'better,' and then start another test on that new element, never getting to the truly optimal mix across all variables simultaneously.

This MCP changes that by centralizing the analysis. It uses `identify_winning_combination` to synthesize every finding into a single, statistically proven recommendation. You stop guessing and start implementing the maximum-conversion setup immediately.

Multivariate Test Analyzer: 3 Tools for Factorial Design Analysis

These tools allow AI agents to calculate the direct impact of individual factors, analyze how pairs of elements interact, or pinpoint the single best-performing combination from a massive dataset.

#	TOOL	DESCRIPTION
01	`analyze_interaction_effects`	Detects how different pairs of tested elements influence each other's performance.
02	`calculate_main_effects`	Determines the core impact of one factor on conversions, based on provided visits and conversion data.
03	`identify_winning_combination`	Pinpoints the single best-performing setup when analyzing all factors and their interactions simultaneously.

See It in Action

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

U What's the optimal mix of headline, color, and CTA for my service page?

Optimal Combination Found:

- **Headline:** 'Boost Your Conversions by 30%' (Main Effect: +12%)
- **Button Color:** Green (Interaction with Headline: +8%)
- **CTA Text:** 'Start Free Trial' (Main Effect: +5%)

Expected Conversion Rate Increase: 19.5%

U Are there any dependencies between our pricing structure and the hero image variant?

Interaction Analysis Results:

Yes. A significant dependency was detected.

- **Pricing Structure 'Tiered':** Performs best when paired with **Hero Image 3 (Team)**, resulting in a combined lift of 18%.
- The combination of 'Simple' pricing and any image shows no statistically significant interaction effect.

U Show me the direct impact of changing our primary value proposition.

Factor: Value Proposition
Levels: 'Speed', 'Reliability', 'Cost-Effective'

LEVEL	MAIN EFFECT (AVG. LIFT)
Speed	+15%
Reliability	+2%
Cost-Effective	-3%

The direct impact of changing the value proposition is highest when focusing on 'Speed.'

Frequently Asked Questions

01 **How does the Multivariate Test Analyzer MCP help with complex CRO?**

It lets you run 2k factorial designs, which is much deeper than basic A/B testing. Instead of just seeing if one element wins, it tells you *how* elements interact to create a high-performing combination.

02 **What kind of data does the Multivariate Test Analyzer MCP need?**

It requires structured test results: the list of factors (e.g., headlines), corresponding visits, and the number of conversions for each tested variant.

03 **Can I use this MCP to find out which button color works best?**

Yes, you can run main effect calculations on the button color factor alone. You can also run interaction effects to see if that color only performs well when combined with a specific headline.

04 **Is this better than just running multiple A/B tests separately?**

Yes, it is far superior. Separate A/B tests miss out on dependencies; this MCP analyzes all elements simultaneously to find the single optimal configuration that no individual test could predict.

05 **What if I have too many variables for the Analyzer?**

The MCP handles complex factorial designs. You simply provide the list of factors and their levels, and it structures the analysis to identify key interactions efficiently.

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 `"multivariate-test-analyzer": { "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

Multivariate Test Analyzer 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

DOCUMENT INFORMATION

Generated	July 2026
MCP Server	Multivariate Test Analyzer MCP
Server ID	`019f11d7-05d6-726e-a083-f6783216db54`
Platform	Vinkius Cloud for AI Agents
Endpoint	`https://edge.vinkius.com/{token}/mcp`

LICENSE & USAGE