MCP SERVER | NO CODE | CLOUD HOSTED

A/B Test Significance Calculator MCP for AI Agents

Calculating Statistical Proof of Conversion Rate Changes

The A/B Test Significance Calculator MCP instantly computes critical metrics for any A/B test result. It determines if observed changes are due to true user behavior or just random chance, giving you p-values, confidence intervals, and clear uplift measurements. You'll get actionable recommendations on whether to stop a test or keep collecting data.

A+ Quality Score 100/100

ab-testing | statistics | p-value | conversion-rate | experimentation

MAINTAINED BY VINKIUS CLOUD FOR AI AGENTS · JUNE 2026
vinkius.com/mcp/ab-test-significance-calculator

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.

A/B Test Significance Calculator MCP

3 tools available
Cloud-hosted on Vinkius

This MCP provides the specialized engine needed to evaluate A/B test results accurately. It helps your AI client compute critical experimental metrics like p-values, 90%, 95%, and 99% confidence intervals, relative uplift, and statistical power immediately. You can use this tool to measure the true magnitude of change between control and variant groups, which is key for product optimization. Furthermore, it lets you determine the probability that observed differences occurred purely by chance. Finally, instead of just giving numbers, it generates an experiment verdict, providing actionable business advice on whether your test should end or if you need more data. By connecting to Vinkius, you get access to this specialized statistical engine alongside thousands of other operational tools for your AI client.

Core Capabilities

01 — Measure Statistical Significance

Calculate core probability metrics and confidence ranges for an experiment using `calculate_significance_metrics`.

02 — Quantify Group Performance Change

Compute the relative and absolute uplift, showing the exact magnitude of change between your control and variant groups via `calculate_uplift_analysis`.

03 — Receive Actionable Test Verdicts

Get a clear business recommendation on whether to end the test or continue collecting data using `generate_experiment_verdict`.

One Click on Vinkius — From Prompt to Execution

Available at **vinkius.com/mcp/ab-test-significance-calculator** — connect your AI agent in three steps.

01 First, feed your AI client the raw results from your A/B test (e.g., visitor counts and conversions for both groups).

02 Next, prompt your agent to run the analysis through the MCP, asking it to calculate core metrics like p-values and confidence intervals.

03 Finally, review the comprehensive output, which includes a clear business recommendation on what action to take next.

The bottom line is that you get statistical certainty and immediate business direction for your product experiments.

Built For

This MCP is built for Product Managers and Data Analysts who run conversion rate optimization (CRO) efforts. Are you tired of guessing whether a small uplift in conversions was real or just random noise? This tool gives you the statistical proof you need to make high-stakes decisions.

Product Manager

Uses this MCP to determine if a new feature launch is genuinely better than the old one, preventing wasted development cycles on inconclusive tests.

Data Analyst

Runs complex statistical checks to validate marketing campaign performance metrics and identify true drivers of conversion rate changes.

CRO Specialist

Leverages the confidence intervals to optimize landing pages, knowing exactly how much lift is statistically guaranteed before pushing a change live.

What Changes When You Connect

01 Stop guessing if your changes worked. The `calculate_significance_metrics` tool gives you the p-value and confidence ranges, so you know when a result is statistically solid.

02 Don't rely on gut feelings about performance. Use `calculate_uplift_analysis` to measure the precise magnitude of change, knowing if the lift was 1% or 20%.

03 Get immediate direction instead of just data points. The `generate_experiment_verdict` tool tells you exactly whether to stop testing or if you need more users.

04 Saves weeks of wasted effort. By validating results quickly, your team avoids committing resources to features that show no real statistical improvement.

05 Builds trust in product decisions. Every launch decision is backed by rigorous statistics and clear probability metrics.

Real-World Applications

Testing a New Checkout Flow

The CRO Specialist runs an A/B test on the checkout page. Their agent uses this MCP to calculate significance metrics, confirming that the new flow's higher conversion rate is not just luck, but a statistically proven improvement.

Comparing Two Pricing Models

The Product Manager tests two pricing tiers and needs to know which model generates the most reliable revenue lift. They use `calculate_uplift_analysis` to quantify the absolute difference in average customer value.

Determining Test Conclusion

After collecting a large data set, the Data Analyst feeds the results into the MCP and asks for a verdict. The system responds with clear guidance: 'Continue testing' or 'Stop now'.

Patterns to Avoid

Comparing raw numbers directly

✗ AVOID

An analyst sees Group A had 50 conversions and Group B had 70. They conclude, 'B is obviously better!' without checking the underlying visitor counts or variability.

✓ INSTEAD

Use `calculate_uplift_analysis` first to get a true relative increase, then use `calculate_significance_metrics` to check if that difference is statistically significant at your chosen confidence level.

Ignoring test duration

✗ AVOID

A team ends an A/B test because the initial results look good, but they didn't wait long enough for full user cycle data.

✓ INSTEAD

Use `generate_experiment_verdict` to get advice on whether the current sample size and time frame are sufficient. It guides you on when your findings are reliable.

Using simple percentages

✗ AVOID

Saying, 'Our new button increased clicks by 15%,' without knowing if that increase was random chance.

✓ INSTEAD

The MCP calculates the confidence intervals and p-value to prove that the 15% lift is reliable and repeatable.

The Right Fit

Use this A/B Test Significance Calculator when your business decisions depend on statistical proof. You need more than just averages; you need certainty regarding conversion rate changes. If you are comparing two groups (A vs B) and the outcome determines a major product pivot or feature launch, use this MCP. Don't use it if you only need descriptive statistics, like calculating means for non-comparative data sets. For simple reporting, basic spreadsheet functions will suffice; but when you hit the ambiguity of 'Is this real?', you absolutely require the confidence intervals and p-values provided here.

A/B Test Significance Calculator: Solving Conversion Rate Proof Problems

Today, product teams spend hours wrestling with data dashboards. They manually pull visitor counts and conversion totals into spreadsheets, spending time calculating basic percentages and comparing raw numbers. The process is tedious and always leaves the team asking the same question: 'Are these differences real?'

With this MCP, your agent handles the entire statistical heavy lifting. You provide the inputs, and the system outputs a fully calculated significance report, including confidence intervals and uplift measurements. You get actionable proof, not just raw numbers.

A/B Test Significance Calculator: Determining Optimal Test Conclusions

Previously, deciding when to end a test felt like guesswork. Teams often quit too early, or worse, waste months collecting data that was already conclusive. This uncertainty slowed down the entire product release cycle.

Now, you simply ask for an experiment verdict. The MCP delivers a clear, objective business recommendation—whether to stop and deploy, or if you need to gather more traffic. You move from indecision to execution.

A/B Test Significance Calculator: 3 Tools for CRO Analysis

These tools let your AI client calculate key statistics, measure performance lift, and provide a definitive verdict on any A/B test result.

#	TOOL	DESCRIPTION
01	`generate_experiment_verdict`	Provides an immediate business recommendation on whether the test should end or if more data needs to be collected.
02	`calculate_significance_metrics`	Calculates core probability metrics and confidence ranges, giving you a deep dive into the experiment's statistical foundation.
03	`calculate_uplift_analysis`	Computes the exact relative and absolute change in performance between your control and variant groups.

See It in Action

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

U I ran a test: Control had 500 visitors/15 conversions; Variant had 620 visitors/24 conversions. Is the difference significant?

Significance Analysis

- **P-Value:** 0.035
- **Confidence Interval (95%):** [1.12%, 1.89%]
- **Verdict:** Statistically Significant.

The observed increase in conversion rate is likely real, suggesting you should proceed with the variant.

U What's the uplift if Control got 400 visits/10 conversions and Variant got 550 visits/23 conversions?

Uplift Analysis Report

- **Absolute Uplift:** +0.86 percentage points
- **Relative Uplift:** ~96%
- **Interpretation:** This shows a strong, measurable improvement in conversion rate compared to the control group.

U Based on these results (p-value 0.15, power 0.7), should I end my test with 95% confidence?

Experiment Verdict

❌ **Recommendation: CONTINUE COLLECTING DATA.**

The p-value of 0.15 exceeds the standard 0.05 threshold. Furthermore, the statistical power of 0.7 is too low to make a definitive call. Keep running the test until you reach better metrics.

Frequently Asked Questions

01 **Does this A/B Test Significance Calculator help me decide if my changes are real?**

Yes, it gives you statistical proof. Instead of guessing, the MCP calculates p-values and confidence intervals, showing you whether an observed change in conversion rate is just random noise or a genuine improvement.

02 **How do I use the A/B Test Significance Calculator to measure lift?**

You use the uplift analysis feature. It computes both relative and absolute changes, telling you exactly how much better your variant group performed compared to the control group in clear percentage terms.

03 **Can this MCP tell me when I should stop my A/B test?**

Absolutely. By using the verdict tool, it reviews all your data—the p-value, power, and significance thresholds—and tells you if the result is conclusive enough to launch or if you need more time.

04 **What kind of metrics does this MCP calculate for my experiments?**

The MCP calculates core probability metrics like p-values, along with confidence ranges (90%, 95%, and 99%). This gives you a complete picture of the statistical certainty behind your data.

05 **Is this A/B Test Significance Calculator only for conversion rates?**

No. While it excels at conversion rates, its underlying math applies to any measurable outcome where you compare two distinct groups against a baseline metric.

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 `"ab-test-significance-calculator": { "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

A/B Test Significance Calculator 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 A/B Test Significance Calculator. 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	A/B Test Significance Calculator MCP
Server ID	`019f11d5-4cc5-7049-bcb9-ed8a56d1e6db`
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/ab-test-significance-calculator.