MCP SERVER | NO CODE | CLOUD HOSTED

AB Test Sample Size Calculator MCP for AI Agents

Calculating Minimum Detectable Effect and Conversion Rates

The AB Test Sample Size Calculator MCP helps data teams nail down the statistical foundations of any experiment. It calculates exactly how many users you need per variant and projects the precise duration your test must run. Plus, it assesses your peeking risk, so you can confidently declare a winner without risking false positives.

A+ Quality Score 100/100

ab-testing | statistics | data-science | experimentation | conversion-rate

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

AB Test Sample Size Calculator MCP

3 tools available
Cloud-hosted on Vinkius

Running an A/B test isn't just about flipping a switch; it’s about statistics. This MCP provides the foundational tools to ensure your experiments are actually reliable. You tell your agent what your baseline conversion rate is and how big of an effect you want to detect, and the tool figures out the exact sample size required per variant. From there, you can get a solid projection on how many days your test needs to run based on current traffic. The most crucial part is checking for peeking risk; this helps prevent you from making calls too early that are just statistical noise. If you're already using Vinkius as your central catalog, connecting this MCP gives your AI client access to essential CRO math right where you need it.

Core Capabilities

01 — Determine required user counts

Calculates the minimum number of users needed in each test group based on your expected conversion rates and target effect size.

02 — Estimate experiment timeline

Projects the necessary duration for an A/B test, using your site's current average daily traffic.

03 — Gauge risk of early analysis

Evaluates how high your probability of a false positive is if you stop analyzing the data before the planned end date.

One Click on Vinkius — From Prompt to Execution

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

01 Input your known metrics: Provide the baseline conversion rate, desired Minimum Detectable Effect (MDE), and confidence level.

02 The MCP processes these inputs using statistical formulas to calculate the necessary sample size per group and project the required test duration based on your traffic volume.

03 Finally, you receive a risk assessment that tells you if continuing the experiment is critical or if you're safe to analyze the results.

The bottom line is: it moves you past guessing game statistics and gives you actionable timelines for reliable data analysis.

Built For

This MCP is critical for Product Managers, Growth Marketers, and Data Analysts. If your job involves testing changes on a website or app—from button colors to checkout flows—you need this. It stops you from wasting time chasing false positives.

Product Manager

Determines whether enough user traffic exists and how long an experiment needs to run before committing engineering resources.

Growth Marketing Specialist

Calculates the required sample size needed for a specific conversion rate improvement, ensuring their campaign has statistical backing.

Data Analyst

Uses the MCP to validate experimental setups and assess the risk of premature conclusion drawing before running any reports.

What Changes When You Connect

01 Stop basing product decisions on gut feelings. The `calculate_required_sample_size` tool tells you exactly how many users are needed, guaranteeing your results matter.

02 Avoid running tests that never finish. Use `estimate_test_duration` to set realistic timelines and manage stakeholder expectations immediately.

03 Eliminate false positives. By using the `assess_peeking_risk` function, you'll know when it's safe to call an end date on your experiment.

04 Your agent can handle complex statistical inputs—like baseline conversion rates and desired power levels—in a single query, saving manual spreadsheet work.

05 The MCP keeps all your core CRO math centralized. You connect once via Vinkius and get access to the full suite of testing tools.

Real-World Applications

Need to test a new checkout flow

A Product Manager asks their agent, 'If our current conversion rate is 4% and I want to detect at least a 10% lift, how big does this A/B test need to be?' The agent uses the required sample size tool and provides the necessary user count per variant.

Testing seasonal changes with limited traffic

A Growth Marketer asks their agent for a timeline: 'Our site only gets 1,000 visitors daily right now; if I need 50,000 users, how long will the test take?' The agent uses the duration tool and gives a precise number of days.

Deciding whether to wrap up an experiment

A Data Analyst checks their running test results and asks the agent about statistical safety. Using the risk assessment tool, they get immediate feedback: 'High peeking risk; continue for another week.'

Validating multiple simultaneous experiments

The team needs to run three concurrent tests (CTA change, image update, pricing model). The agent runs the sample size tool for all three, ensuring they don't over-allocate resources or under-test critical variables.

Patterns to Avoid

Assuming traffic is enough

✗ AVOID

A team decides to run a test because the product manager says 'we get tons of traffic!' but forgets that the required sample size was never calculated.

✓ INSTEAD

Don't guess. Always start by running `calculate_required_sample_size` using your target MDE and baseline CR. Only after you have the number do you check if enough people are coming in.

Stopping tests based on early trends

✗ AVOID

The conversion rate looks amazing for the first week, so the team declares a winner and ships the change before the test is properly completed.

✓ INSTEAD

Always use `assess_peeking_risk` to validate your timing. It tells you if the current positive trend is statistically significant or just random chance.

Mixing up metrics

✗ AVOID

The analyst uses a tool designed for user counts but plugs in daily revenue numbers, getting garbage results that mislead the entire team.

✓ INSTEAD

Be specific about your inputs. `calculate_required_sample_size` needs rates (e.g., 4%) and percentages (MDE), not absolute dollar amounts.

The Right Fit

Use this MCP if your primary problem is statistical rigor: you need to know *if* an observed change is real, or just random chance. Specifically, use it when setting up a new experiment, checking if current traffic levels support the test size, or deciding whether to call time on a running test. Don't use this if you simply need to visualize data (use standard analytics dashboards). You also shouldn't rely solely on it for *what* change to make; that requires domain knowledge. However, if you have a hypothesis—like 'changing the button color will boost CR by 5%'—this MCP is the necessary math layer between your idea and deployment.

AB Test Sample Size Calculator: Determining Minimum Detectable Effect in CRO

Currently, setting up an A/B test involves a lot of manual guesswork. Teams often estimate the sample size based on 'gut feel' or outdated benchmarks, leading to tests that are either too small (and inconclusive) or unnecessarily huge (wasting resources). It’s tedious work involving cross-referencing statistical calculators and making assumptions about traffic growth.

With this MCP, you simply provide your baseline conversion rate and the minimum lift you want to prove. The tool immediately calculates the required sample size for both variants. You get a definitive number—the exact user count needed—and that’s it.

AB Test Sample Size Calculator: Managing Experimental Duration with Traffic Forecasting

Manually forecasting test duration is painful. You have to calculate total required users, divide by current daily traffic, and then account for potential dips or spikes in visitor volume. This process requires jumping between multiple spreadsheets and making subjective adjustments.

This MCP handles that complexity instantly. After determining the necessary user count, you run the duration estimate tool. It gives you a clear timeline—say, '20 days'—allowing your team to schedule dependencies accurately.

3 AB Test Sample Size Calculator Tools for Conversion Rate Analysis

These three tools let your agent calculate user requirements, project test durations, and assess the risk of early analysis in any A/B experiment.

#	TOOL	DESCRIPTION
01	`calculate_required_sample_size`	Figures out how many users you need in each group for a statistically sound A/B test setup.
02	`estimate_test_duration`	Provides an estimate of how long your experiment must run based on your site's daily visitor traffic.
03	`assess_peeking_risk`	Warns you if analyzing the data too early increases the chance of a false positive result.

See It in Action

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

U I need to test if changing our headline increases conversions by 8%. What sample size should I aim for?

A/B Test Sample Size Calculation

Based on your inputs (4.2% baseline CR, 0.5% MDE, 95% confidence):

- **Required Users per Variant:** ~18,500 users
- **Total Required Users:** ~37,000 users

This is the minimum user count needed to reliably detect that lift.

U If we need 60,000 total users and our site usually gets 2,000 visitors a day, how long should the test run?

Estimated Test Timeline

- **Target Users:** 60,000
- **Daily Traffic:** 2,000
- **Calculated Duration:** 30 days

We recommend scheduling for a minimum of 30 full days to account for weekend variations.

U Our test has been running for 10 days out of the planned 28. Should I stop it already?

Peeking Risk Assessment

- **Current Status:** 10/28 days complete.
- **Risk Level:** Medium-High.
- **Recommendation:** Do not make a final call yet. While the results look promising, continuing until day 25 will significantly lower your false positive risk and give you more statistical certainty.

Frequently Asked Questions

01 **How does the AB Test Sample Size Calculator MCP determine how many users I need per test?**

The tool calculates the minimum number of participants required for each group. It uses your baseline conversion rate and the effect size you want to detect, ensuring that if a real change happens, your test has enough power to prove it.

02 **Can I use this MCP to figure out how long my A/B test must run?**

Yes. You provide the total user count needed and your site's average daily traffic. The calculator then gives you a precise, data-backed estimate of the minimum number of days required.

03 **What is 'peeking risk,' and how does this MCP help me avoid it?**

Peeking risk is the danger of stopping a test early because the numbers look good. This MCP assesses that risk, telling you if you must wait for the full planned duration to prevent making false conclusions.

04 **Do I have to know my baseline conversion rate to use the AB Test Sample Size Calculator?**

Yes, knowing your current performance (the baseline CR) is essential. The tool needs this starting point to accurately calculate how large of a difference you need to detect.

05 **Is this MCP useful for testing different marketing channels?**

Absolutely. Whether the traffic comes from search, social media, or email campaigns, this MCP uses your aggregate daily traffic numbers to provide accurate test duration estimates.

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-sample-size-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

AB Test Sample Size 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 AB Test Sample Size 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	AB Test Sample Size Calculator MCP
Server ID	`019f11d4-f53f-73d9-a117-37a273f97646`
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-sample-size-calculator.