

PUBLIC SAMPLE TEARDOWN

Amplitude leads the VP shortlist. The next opportunity is pricing and portability proof.

We sampled how ChatGPT, Claude, Gemini, and Perplexity answer high-intent product analytics prompts. Amplitude already appears as the enterprise-safe choice for retention and experimentation, while PostHog increasingly owns pricing predictability and engineering-led portability.

AI VISIBILITY SNAPSHOT

Amplitude		4/4
Mixpanel		#2
PostHog		Risk

60

LLM answer runs across product analytics buyer prompts

4

AI answer engines: ChatGPT, Claude, Gemini, Perplexity

15/15

Prompt groups where Amplitude appeared

3

Priority actions for content, proof, and reruns

Where Amplitude wins, and where competitors shape the buyer narrative.

LLM answers are buyer-visible perception signals, not verified product facts. The commercial value is seeing which vendor narratives appear before buyers reach a website, demo form, analyst report, or review page.

METRIC	VALUE	WHAT IT MEANS
Prompts tested	60 runs / 15 groups	Coverage across discovery, alternatives, evaluation, pricing, security, risk, lock-in, and final shortlist prompts.
Platforms	4 engines	ChatGPT, Claude, Gemini, and Perplexity were sampled to compare answer patterns across major AI surfaces.
Clusters where Amplitude appears	15/15	STRENGTH Amplitude has broad AI visibility and is not absent from core buyer journeys.
Final shortlist position	#1 in 4/4	STRENGTH AI systems frame Amplitude as the enterprise product analytics default for VP-level shortlists.
Top competitor narrative risk	PostHog	RISK PostHog owns pricing predictability, open-source portability, and engineering-led data ownership.
Proof-density gap	Neutral sources	Many positive Amplitude claims trace to vendor-owned comparison pages instead of analyst, customer, or third-party proof.

WIN

Enterprise shortlist default

Amplitude is repeatedly recommended for retention, cohort analysis, integrated experimentation, and governance.

RISK

Pricing predictability

PostHog is framed as more predictable at 100K-500K MTU scale, while Amplitude is framed as premium.

GAP

PM-friendly speed

Mixpanel still wins some speed and simplicity narratives even when Amplitude wins on depth and governance.

Three moves to protect the #1 position and close bottom-funnel narrative gaps.

The goal is not generic SEO. It is to publish proof that AI systems can retrieve when buyers ask comparison, pricing, migration, and shortlist questions.

P0

Publish an Amplitude vs PostHog TCO guide.

Target pricing and alternative-search prompts with honest MTU/event-volume scenarios. Show when Amplitude's enterprise capabilities justify premium positioning, and where buyers should expect cost tradeoffs.

P0

Make the metric-to-replay workflow highly citeable.

Create a focused page with screenshots, short demos, docs, and customer examples that show how teams move from metric drop to replay evidence to experiment decision without losing governance.

P1

Add third-party proof to defend the final shortlist.

Reinforce Amplitude's #1 AI-visible position with current analyst citations, customer stories, neutral comparison evidence, and implementation proof for retention, experimentation, and executive reporting.

What a full 14-day audit would deliver

This public sample shows the type of finding. A paid AnswerMap audit expands the prompt set, competitor benchmark, source analysis, and action roadmap so teams know exactly what to publish or correct first.

100-200 prompts across discovery, alternatives, evaluation, pricing, security, and shortlist decisions.

3-5 competitors benchmarked by prompt-level share of answer and narrative ownership.

4-5 AI platforms with answer extraction, claim review, source mapping, and rerun guidance.

30-day action roadmap prioritized by buyer intent, narrative risk, and content/proof effort.