

CASE STUDY · AI SEARCH VISIBILITY

The 9% Problem: 400 AI Prompts, Measured Against 182 Days of Real Search Demand

Creative Data Engineers · July 2026

Summary

An agency delivered its AI-search strategy for a brand: hundreds of slides of analysis and a 400-prompt workbook to steer content. The prompts were well written, generated from a persona, intent, and topic framework; not one prompt was tied to a measured query. We connected the target property's Google Search Console, pulled 182 days of data (roughly 5 million impressions, 5,000 distinct queries) and matched the agency's question clusters against it. The clusters the set covers carry **roughly 9% of real query impressions**. The other 91%, model prices, availability, technical data, and a technology topic the brand had no page for at all, was absent from the set. The fix was a method change: replace invented questions with measured ones. A per-page fan-out of buyer questions, grounded in the queries each page already receives and scanned against six pages, produced 14 ready-to-adapt content drafts, plus a visibility dashboard and structured-data recommendations down to the named properties.

Client and agency are not named. The numbers are unchanged from the client deliverable and trace to dated artifacts.

The setup

Our client's press portal, built for journalists, had quietly become their best-performing AI-search asset: server-rendered, ranking on question searches, and regularly cited by AI assistants in its market.

To steer content for AI visibility, the brand's agency delivered its AI-search strategy: hundreds of slides of analysis and a workbook of 400 prompts, the questions they believed buyers ask AI assistants about the brand. The prompts were well written and thoughtfully organized.

The question opinion could not settle

Are these the right 400 questions? The agency believed yes. Instinct said parts were missing. Nobody had data. The odd part: this discipline has had the answer for twenty years. SEO always started from what people actually type, and Search Console has always been where that data lives. A prompt set is the new keyword list, and the old rule still applies: check the list against measured demand before it steers a page.

We treated the prompt set as a measurable claim instead of a matter of opinion.

What we did

- **Connected the property's Google Search Console and pulled 182 days of data:** roughly 5 million impressions across 5,000 distinct queries. Real searches by real people that surfaced the client's pages.

- Clustered the demand by topic and matched each of the agency's 400 questions against the clusters.
- Ranked every content decision by measured impressions, instead of by matrix cell.

What the data showed

- **The covered clusters carried roughly 9% of query impressions.** The set hit the two technology clusters squarely: real demand, correctly sensed.
- **The bulk of demand sat in clusters the set never touched:** prices, availability, and technical data. Among them single queries larger than whole covered clusters.
- **One technology topic had real measured demand and no page** anywhere on the client's properties. No framework can surface a gap like that, because no framework knows what is missing from your site.
- **Practically all of the portal's demand was consumer intent,** on a portal built for journalists. Journalist terms such as embargoes and press kits had no measurable Google demand: journalists do not arrive via Google.
- **The agency's instincts held real value.** Their technology-explainer prompts matched real queries. As an editorial blueprint, the set held up. As a priority ranking, it collapsed to roughly 9%.

What replaced the list

Instead of one global prompt list, a per-page engine: for each high-demand page, take the queries that actually land on that page, fan them out into the roughly 12 sub-questions an AI assistant breaks the topic into when it builds an answer, scan the page for which questions it already answers, and draft fill-content for the gaps using only facts already published on that page.

Six pages scanned, 14 ready-to-adapt content drafts delivered, alongside a visibility dashboard and technical structured-data recommendations specified down to the JSON-LD @graph nodes and named properties. Every draft traceable to measured demand, every fact traceable to the client's own published content.

The outcome

The analysis was presented to the client's team the morning it was finished. The measurement framing worked exactly as designed: no debate about whose questions were better. The team member who had commissioned the original question set summarized the result for a new colleague themselves.

And the strongest proof of buy-in: **while the presentation was still running, the client's technical lead connected every European market to Google Search Console live on the call,** so the same demand analysis can now run for every market. The drafted content packs went straight to the editorial team.

The lesson

Generated question sets are hypotheses. Measured demand sets the priorities.

Nobody had to be wrong in the room. The covered share, roughly 9%, was real demand, correctly sensed. From here, every question set, anyone's, is ranked against the measured queries before it steers a single page. That is the repeatable method: instrument prompt sets instead of debating them.

About the numbers

This case study publishes four figures, and only these:

What we publish	Value
Measurement window	182 days (January to July 2026)
Search impressions analyzed	roughly 5 million
Distinct queries analyzed	5,000
Demand covered by the 400-prompt set	~9% of query impressions
Pages scanned / content drafts delivered	6 / 14

The full query-level breakdown stays where it belongs: with the client.

About Creative Data Engineers

Creative Data Engineers is a data and AI consultancy. The AI Search Visibility service measures how visible a brand is in AI-generated answers across ChatGPT, Perplexity, Google AI Overviews, and more, then delivers fixes ready to publish: a Baseline Audit from EUR 4,900, ongoing Monitor + Fix from EUR 1,490 per month.

Read this case study online, with FAQ and sources: creative-data-engineers.com/case-studies/prompt-set-audit

Contact: turan@creative-data-engineers.com, creative-data-engineers.com