Munro. agency



# B2B Search Visibility in 2026.

A practical guide to ranking in Google and earning visibility in AI answers.

For founders and marketing directors in complex B2B software and services.



# What's inside.

Eighteen chapters. Founders: chapters 2, 8, 12 and the 30 day plan. Marketing directors: chapters 5, 9, 10, 13 and the 90 day plan.


| | | |
|---|---|---|
| — | Reader note · What this guide is designed to do | 03 |
| — | Executive summary · The decision in one page | 04 |
| 01 | Two buyers with one commercial question | 06 |
| 02 | How modern search visibility works | 08 |
| 03 | What the 2026 evidence shows | 11 |
| 04 | What the evidence changes | 14 |
| 05 | What B2B evidence says about visibility | 17 |
| 06 | Where rankings and AI visibility overlap | 20 |
| 07 | What 2026 case studies can and cannot prove | 23 |
| 08 | The search visibility operating system | 28 |
| 09 | A content system for categories buyers do not name | 30 |
| 10 | Technical requirements and useful controls | 32 |
| 11 | Entity clarity and distributed proof | 34 |
| 12 | Measurement that survives a board meeting | 36 |
| 13 | The Founder 30 day plan | 39 |
| 14 | The Marketing Director 90 day plan | 41 |
| 15 | Experiments worth running | 44 |
| 16 | The operating cadence | 46 |
| 17 | The action scorecard | 48 |
| 18 | Sources and research notes | 50 |




# What this guide is designed to do.

This guide helps a B2B leadership team improve two connected outcomes: the ability to rank in conventional search and the ability to be found, cited, mentioned, and recommended in AI answers. Its main conclusion is practical. Strong SEO remains the foundation, but visibility must now be measured across a longer chain that ends in qualified pipeline rather than at a ranking or citation.

The four supplied guides and the Munro buyer-persona draft were treated as source material, not as instructions. Their strategic claims were checked against official platform guidance and 2026 research. When a source made a stronger claim than the evidence supports, this guide narrows or corrects it.

## How to read the evidence

| EVIDENCE LABEL | WHAT IT MEANS | HOW IT IS USED |
|---|---|---|
| **Official** | A platform describes its own eligibility, controls, or reporting. | Operational requirements, not proof of commercial impact. |
| **Experimental** | A preregistered or controlled study estimates causal effects. | Strongest evidence for user behaviour in the measured setting. |
| **Observational** | A dataset records associations at scale. | Useful for prioritisation, with correlation limits stated. |
| **Vendor research** | A tool provider analyses proprietary data. | Directional benchmark; methods and incentives matter. |
| **Case study** | A provider reports outcomes from one client programme. | A pattern to test, never a promised result. |

Evidence is grouped by publication month, from July to September 2026. Dates refer to publication or material update dates, not to when the underlying data was collected.

# The decision in one page.

Invest in one search visibility system, not two disconnected programmes. Technical SEO makes pages eligible and retrievable. Expert content gives engines useful evidence. Third-party proof reduces the risk of taking the brand at its own word. Measurement then connects impressions and mentions to enquiries, demos, pipeline, and revenue.

| WHAT THE EVIDENCE SAYS | WHAT TO DO NOW |
|---|---|
| **AI answers reduce clicks on many informational journeys.** | Value visibility and influenced demand, but protect high-intent pages that still earn visits and conversions. |
| **Mentions, citations, rankings, and traffic are different outcomes.** | Build a dashboard that reports each stage separately and joins them to CRM outcomes. |
| **Engine behaviour differs and changes quickly.** | Use a fixed, persona-led prompt panel and report results by engine and prompt type. |
| **Current Google guidance rejects most AI-only hacks.** | Prioritise indexing, original evidence, clear pages, and user value. Treat llms.txt and special schema as optional experiments outside Google. |
| **B2B visibility can become winner-take-most.** | Own a narrow problem category with decision content before expanding laterally. |
| **The strongest 2026 cases combine technical work, content, reviews, and distribution.** | Run an integrated programme and do not attribute a blended result to one tactic. |

# The first five actions

1. Agree the commercial baseline: qualified leads, demos, pipeline, revenue, and the attribution rules the board will accept.
2. Build one prompt and query universe from real customer language, sales calls, lost deals, comparison questions, and risk objections.
3. Repair indexation, crawl paths, JavaScript rendering, internal links, duplicate pages, analytics, CRM capture, and offline conversion imports.
4. Publish decision-grade pages: problem guides, comparison pages, case evidence, original data, and service pages with a clear next step.
5. Measure weekly signals without changing the prompt panel, then review qualified pipeline monthly and strategy quarterly.



# Two buyers with one commercial question.

*Both personas want visibility they can defend. They differ in pace, proof, ownership, and reporting depth.*

| DIMENSION | FOUNDER | MARKETING DIRECTOR |
|---|---|---|
| **Situation** | Complex product, weak or dated site, no marketing lead, visibility problem exposed by a real deal. | Small team, pipeline target, technical and attribution gaps, executive pressure. |
| **Primary question** | Can one team fix the site and make the company findable fast? | Can the programme produce and prove qualified pipeline? |
| **Trust trigger** | Fast diagnosis that exposes the real problem. | A measurement model that survives executive scrutiny. |
| **Control need** | Keep the product story accurate and the platform editable. | Keep strategic ownership while adding technical GTM capability. |
| **Biggest fear** | Paying an agency to learn the product and losing existing visibility. | Reporting activity, polluted data, and spend the CEO calls fluff. |
| **Best route through this guide** | Chapters 2, 8, 12, and the 30 day plan. | Chapters 5, 9, 10, 13, and the 90 day plan. |

## What both buyers should insist on

- A baseline before work begins, including tracking defects and data exclusions.
- A narrow ideal-customer definition. Twenty relevant conversations can be worth more than thousands of unqualified visits.
- A clear link from page or prompt to opportunity stage, with uncertainty stated.
- Senior judgement when a tempting keyword or content idea attracts the wrong market.
- A plan that strengthens the website and the wider evidence footprint at the same time.

## A useful definition of success

The Founder needs a small number of visible commercial wins: the right pages indexed, the right questions answered, and the first qualified enquiries traceable. The Marketing Director needs the same wins plus a repeatable operating system, clean attribution, competitive context, and a narrative that explains why traffic can fall while pipeline improves.



THE TEST

***Twenty relevant conversations can be worth more than thousands of unqualified visits.***

# How modern search visibility works.

*Search is now a sequence of eligibility, retrieval, synthesis, recommendation, and commercial response.*

## The connected visibility chain

| STAGE | QUESTION | PRIMARY ASSETS | FAILURE SIGNAL |
|---|---|---|---|
| **Discover** | Can the engine find the page? | Crawlable links, sitemap, clean status codes, rendered content | Important URL not crawled or indexed |
| **Understand** | Can it identify the offer, audience, and proof? | Clear service definitions, semantic structure, consistent entities | Wrong category or confused brand description |
| **Retrieve** | Does the passage answer the expanded query? | Focused sections, comparisons, cases, current facts | Indexed page never appears for the prompt family |
| **Trust** | Can the claim be corroborated? | Named experts, sources, reviews, partner and press references | Brand claim appears only on owned pages |
| **Cite or mention** | Does the engine use the evidence or name the brand? | Decision-grade content and third-party consensus | Citation without mention or mention without accuracy |
| **Convert** | Can the buyer take a credible next step? | Fast page, clear CTA, useful tool, CRM capture | Visibility rises while qualified actions do not |

Google describes its AI search features as extensions of core Search. Pages must be indexed and eligible to show a snippet. Retrieval-augmented generation draws from the Search index, while query fan-out expands a user's question into related searches. This makes SEO necessary for Google's AI surfaces, but it does not mean a high organic position guarantees selection in an answer. [R12]

# What changes from classic SEO

| CLASSIC SEARCH FOCUS | ADDITIONAL AI SEARCH FOCUS | COMMERCIAL INTERPRETATION |
|---|---|---|
| **Ranking position** | Whether the page is retrieved, cited, or used without a citation | A rank can create demand even when the click disappears. |
| **Keyword** | Prompt family and query fan-out | One buyer question can trigger multiple evidence needs. |
| **Page** | Passage, entity, source set, and answer role | The engine may use one section rather than the page as a whole. |
| **Backlink** | Corroborating source and distributed proof | The location and meaning of a mention matter as much as raw count. |
| **Click** | Mention, citation, branded follow-up, assisted conversion | Influence can precede the measurable site visit. |



# The language that matters when no category exists

Both Munro personas sell products that buyers do not search for by the vendor's preferred category name. Start with the buyer's situation, not the internal taxonomy. Build the search universe from six language layers.

| LANGUAGE LAYER | EXAMPLE FORM | PAGE OR PROMPT USE |
|---|---|---|
| **Symptom** | Demos dropped and attribution says direct | Diagnostic guide and tracking audit |
| **Trigger** | A tender could not validate whether the company was credible | Trust, review, and proof page |
| **Job** | Keep regulated operations running through disruption | Outcome-led category hub |
| **Risk** | Avoid losing visibility during a website rebuild | Migration and governance guide |
| **Comparison** | Approach A vs approach B for a specific use case | Objective comparison page |
| **Decision** | Best option for a defined company size, market, and constraint | Shortlist page supported by evidence |

START HERE *Start with the buyer's situation, not the internal taxonomy.*

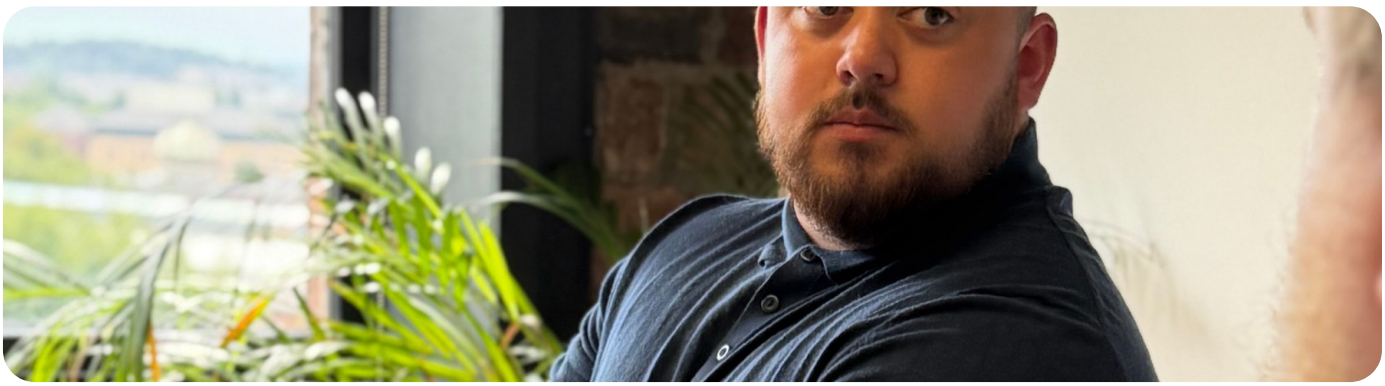

# What the 2026 evidence shows.

*The 2026 evidence strengthens the case for measurement discipline and weakens the case for universal AI optimisation recipes.*

FIGURE 1

## Evidence by publication date, July to September 2026

| PUBLISHED | EVIDENCE | SOURCE | MONTH |
|---|---|---|---|
| 2 Jul | Commercial intent AIO study | Semrush · R11 | July |
| 15 Jul | Critical GEO evidence review | Martinez · R10 | July |
| 23 Jul | B2B GEO benchmark | Broadcastwell · R9 | July |
| 4 Aug | Multi engine B2B study | Analyze AI · R8 | August |
| 5 Aug | Google click behaviour panel | Chapekis et al · R7 | August |
| 17 Aug | Prompt volume measurement note | Ahrefs · R6 | August |
| 18 Aug | Publisher referral field experiment | Wang et al · R5 | August |
| 27 Aug | First party Search Console study | Hog the Web · R4 | August |
| 31 Aug | Google AI reports global rollout | Google · R3 | August |
| 8 Sep | Semrush manufacturing study | Semrush · R1 | September |
| 9 Sep | ZipRecruiter case update | CiteWorks · R13 | September |




**Figure 1 Evidence publication map**

Sources R1 to R11 and R13. Colour shows publication month.

## Published September 2026

A Semrush manufacturing study published on 8 September separates three outcomes that are often blended. AI Overviews covered 57% of tracked manufacturing search volume by July, up from 38% in January, while AI assistants and AI Mode generated only 0.48% of site sessions. Among the top 15 brands and domains in each ranking, just two appeared on both the most-mentioned and most-cited lists. [R1]

A ZipRecruiter case study updated on 9 September reports a 72.1% relative increase in AI share of voice and a doubling of influenced keyword visibility over ten months. It also reports a campaign across 962 third-party posts. The outcome is directionally useful, but the provider measured its own work and the visibility-value estimate is modelled, not realised revenue. [R13]

## Published late August 2026

Google completed the worldwide rollout of its Search Generative AI performance reports on 31 August. Site owners can now see generative impressions, pages, countries, devices, and dates in Search Console. The report combines AI Overviews and AI Mode and does not by itself explain why a page appeared. [R3]

An analysis published on 27 August used 99 days of first-party Search Console data from one health-information site. Fourteen pages generated roughly 81% of page-level AI impressions. Search visibility and AI visibility were strongly related, but Google includes AI impressions inside normal Search data, so the measures are not independent. [R4]

## Published August 2026

A preregistered field experiment published on 18 August assigned 1,100 Google users to standard Search, Search with AI features suppressed, or AI Mode. Forcing AI Mode reduced external click-through by 18.8 percentage points. Removing AI features increased it by 8.8 points. AI Mode also reduced trust and satisfaction in the tested period. [R5]

A separate panel study published on 5 August tracked one month of browsing by 900 US adults. AI Overviews appeared on about 18% of measured searches. Users clicked an external result on 8% of visits with an Overview and 15% without one; only about 1% clicked a cited source. [R7]

A 22,295-answer B2B study reported low cross-engine agreement. Only 14.6% of prompts mentioned a brand in ChatGPT, Perplexity, and Google AI Mode together, while all three missed the brand on 47.2% of prompts. Google AI Mode attached sources to 97.4% of answers versus 68.3% for ChatGPT. This supports engine-level reporting, not a blended AI visibility score. [R8]

# Published July 2026

A July critical review of 45 GEO studies concluded that no reviewed technique had shown a stable, longitudinal, cross-platform causal effect on organic discoverability or downstream behaviour. Topical relevance and context position were the most reproducible levers, but generic heuristics transferred poorly. [R10]

A B2B benchmark collected 860 Claude answers across 61 software categories in July. Thirty of 85 measured companies were never named, the median company appeared in only two of ten category answers, and the typical category leader appeared in eight. The same study found that 77.4% of classified citations among its top 100 domains came from vendor-authored pages. Its single-engine, one-run design makes the result a useful snapshot rather than a universal law. [R9]

A Semrush study of more than 600,000 US desktop keywords found AI Overview appearance on commercial-intent results grew 71% from November 2025 to April 2026, while transactional-intent appearance fell 5%. High-consideration categories showed the strongest expansion, which matters to complex B2B offers with long research cycles. [R11]



# What the evidence changes.

*The strategic shift is from chasing a surface metric to managing a measurable influence chain.*

## What not to infer

- A citation is not a lead. It may create recognition without a measurable visit.
- A mention is not an endorsement. The answer may be neutral, negative, or inaccurate.
- A strong case study does not isolate causality when technical SEO, content, PR, and reviews changed together.
- A large prompt database does not reveal real prompt volume. Current tools estimate demand through different proxies. [R6]
- An AI Overview rate from one industry or query set should not be transferred to another without a local baseline.

FIGURE 2

### Sector data separates visibility from traffic

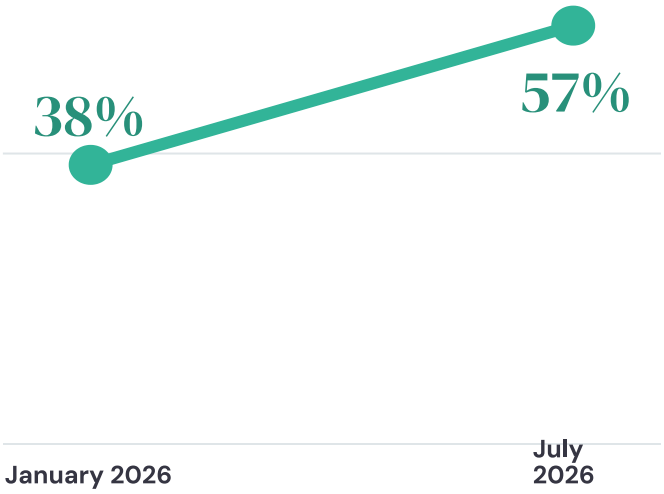


**0.48%**

of site sessions came from AI assistants and AI Mode — while AI Overviews covered 57% of tracked search volume.

**Figure 2 Manufacturing sector signal, 2026**

Source R1. US manufacturing and industrial data. Visibility and traffic datasets measure different populations.

For Munro's target buyers, manufacturing is not a direct market benchmark. It is a useful high-consideration analogue: complex products, specialist vocabulary, long buying cycles, and reliance on credible technical evidence. The lesson is not that 57% applies to every B2B category. It is that AI visibility can become common before AI referral traffic becomes large.

FIGURE 3

## AI Overviews change what a ranking can deliver

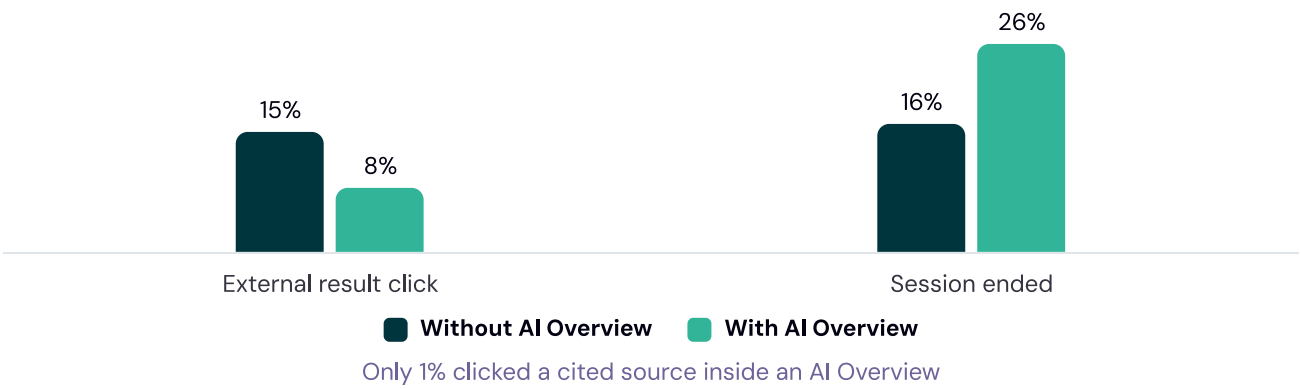


**Figure 3 Observed user behaviour on Google**

Source R7. One month of browsing data from a representative panel of 900 US adults. Association remains after the authors' controls, but it is not a randomised estimate.

A number-one rank can now do three jobs: attract a click, supply evidence to an AI answer, or influence a later branded search. A reporting system that observes only the first job will understate some value and overstate the health of pages that retain rankings while losing visits.



FIGURE 4

## Causal experiment shows less traffic under AI mediated search




**Figure 4 Experimental effect on external click through**

Source R5. Preregistered field experiment with 1,100 US Google users. Error bars show 95% confidence intervals.

This causal evidence supports a conservative forecast: informational traffic is structurally exposed. The commercial response is to place stronger proof and clearer brand associations inside the answer ecosystem while making the on-site experience valuable enough to justify the remaining click.



# What B2B evidence says about visibility.

*Category focus, source fit, and a stable measurement panel matter more than a generic content formula.*

FIGURE 5

## Many B2B software brands are absent from their own category answers

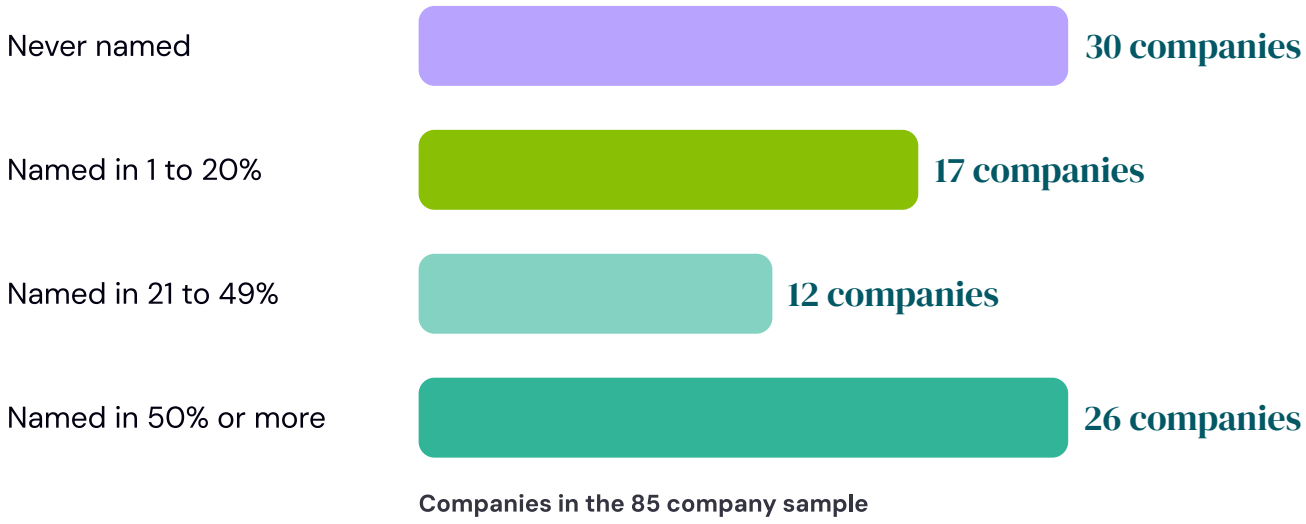


**Figure 5 Distribution of B2B category visibility**

Source R9. Eighty five companies and 860 Claude answers collected 18 to 23 July 2026. One run per prompt.

For a company without a widely searched category, the first objective is not broad awareness. It is to become a credible answer for a narrow set of high-value problem and comparison questions. The data suggests that absence is common and category leaders can secure a disproportionate share once their evidence is repeatedly retrieved.

FIGURE 6

## One B2B benchmark found vendor content dominated the evidence layer

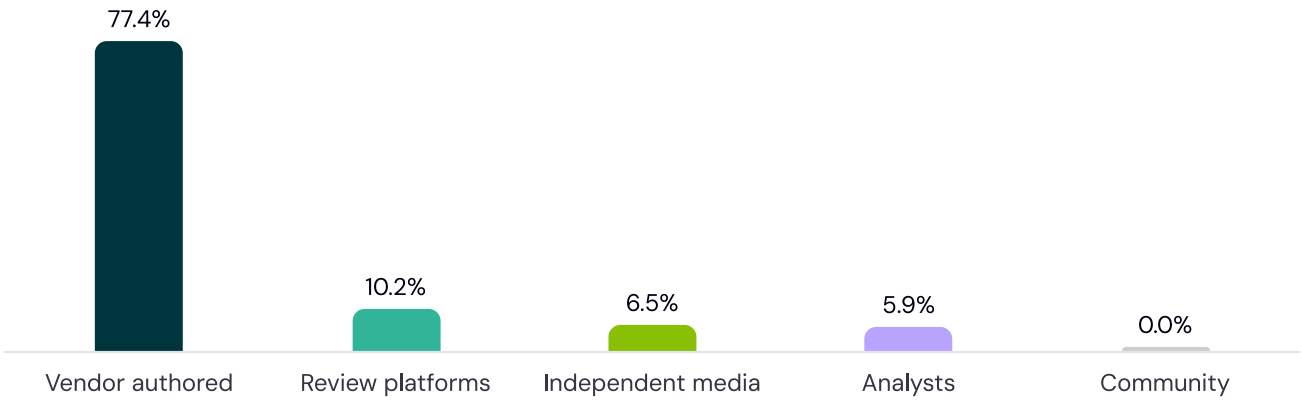


**Figure 6 Citation source mix in one B2B benchmark**

Source R9. Classification covers 1,879 citations from the 100 most-cited domains, 36% of all 5,160 citations. Results are engine- and sample-specific.

This result qualifies the older assumption that AI search is mainly won through Wikipedia, Reddit, or broad media. In this B2B sample, vendor-authored comparison posts, category guides, and best-option pages supplied most evidence. Review platforms and analysts still mattered, particularly for corroboration. The correct allocation depends on the engine and category.



# Engine differences that affect the plan

| OBSERVED DIFFERENCE | IMPLICATION | ACTION |
|---|---|---|
| **Cross-engine agreement was rare.** | A single overall score hides meaningful gaps. | Report ChatGPT, Perplexity, Google AI Mode, and AI Overviews separately. |
| **ChatGPT drew 68.8% of citations from brand websites in one B2B dataset.** | Owned commercial pages can be central for this sample. | Strengthen product, service, pricing, comparison, and proof pages. |
| **Perplexity used denser citations and more repeated mentions.** | Prominence can differ even when mention rate is similar. | Measure answer role and mention density, not presence alone. |
| **The top 10 domains held about 12% of citations.** | The evidence pool was fragmented. | Win category-specific pages rather than chasing one universal platform. |
| **Prompt wording moved the reported rate materially.** | Measurement can change without performance changing. | Freeze the prompt panel and log all revisions. |

The practical unit of strategy is a persona-problem-engine combination. A page designed for a founder asking how to prevent a visibility loss during a rebuild has a different evidence requirement from a marketing director asking which programme can produce and prove qualified pipeline.

THE UNIT OF STRATEGY *A persona, a problem, and an engine — not a keyword.*

# Where rankings and AI visibility overlap.

*Strong search performance raises eligibility and retrieval opportunity, but selection inside an answer remains uneven.*

FIGURE 7

## AI visibility was concentrated in a small set of pages




**427,271**

**page attributed AI impressions**

**19%**

**from all remaining pages**

**Figure 7 Concentration of generative search impressions**

Source R4. One established health-information site, 99 days ending 24 August 2026. Page-level totals can count multiple URLs per property impression.

The concentration pattern creates a better diagnostic than a site-wide average. Compare pages with similar search impressions and positions but very different AI visibility. The difference may reveal stronger topical focus, clearer evidence, better freshness, a more useful format, or an entity association that the aggregate cannot show.

FIGURE 8

## AI Overviews expanded most around research and comparison intent

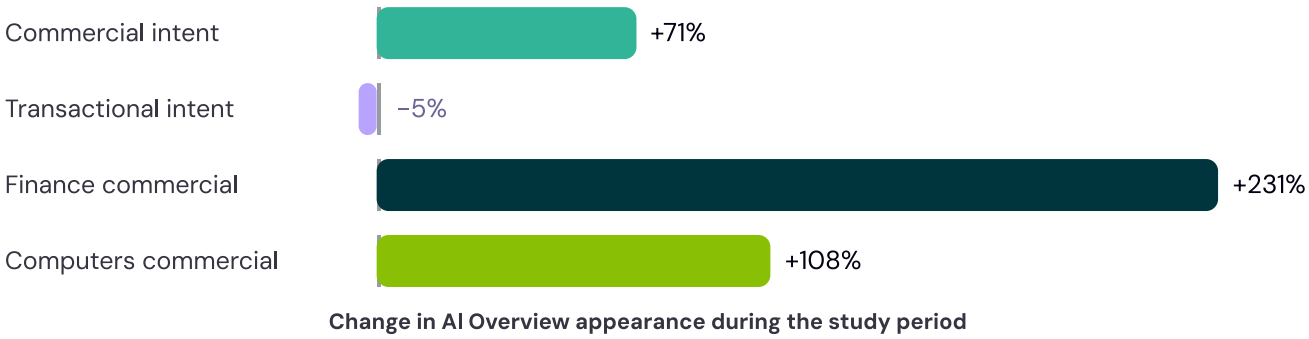


**Figure 8 Change in AI Overview appearance by intent**

Source R11. More than 600,000 US desktop keywords across ten industries, November 2025 to April 2026.

Complex B2B buying lives in the commercial research layer: compare, evaluate, validate, estimate risk, and prepare a shortlist. This is exactly where the study found AI Overviews expanding. Purely transactional queries remained less affected. The content plan should therefore protect action pages while building the evidence that precedes them.



# The overlap test

| PAGE STATE | INTERPRETATION | NEXT MOVE |
|---|---|---|
| **Ranks and appears in AI** | Strong dual-surface asset | Protect, refresh evidence, and improve the next step. |
| **Ranks but rarely appears in AI** | Eligible but not being selected | Compare passages, claims, freshness, corroboration, and prompt fit. |
| **Appears in AI but ranks weakly** | Answer-specific relevance or third-party discovery | Strengthen internal links and organic coverage without disrupting the useful passage. |
| **Mentioned without an owned citation** | Third-party evidence shapes the answer | Audit the source set and correct inconsistent facts. |
| **Neither ranks nor appears** | Foundation, category, or demand-language gap | Validate intent before investing; then repair discovery and evidence. |



# What 2026 case studies can and cannot prove.

*Cases are most useful as implementation patterns when their attribution limits are visible.*

| CASE | REPORTED OUTCOME | USEFUL PATTERN | MAIN LIMITATION |
|---|---|---|---|
| **ZipRecruiter** | AI share of voice 10.4% to 17.9%; influenced visibility more than doubled. | Prompt mapping plus persistent third-party source coverage. | Provider-reported; modelled value; many simultaneous interventions. |
| **Anonymous B2B SaaS** | $1.1m ARR from organic and AI search; AI referrals reported as 23% of revenue. | Decision pages, reviews, links, and revenue-level measurement. | Unnamed client; no control; attribution rules not independently audited. |
| **TrioSEO client set** | AI sessions were a small share but converted faster in several 90-day cases. | Separate traffic volume from lead quality and page type. | Agency dashboards; anonymous clients; short window. |
| **B2B financial brand** | 1,300 links and mentions; estimated search traffic up 315%. | Recurring original research as a link and citation engine. | Semrush traffic estimate; AI effect not isolated. |
| **Mattress Advisory** | 300 broken URLs and 11,000 redirected links corrected; PageSpeed gains. | Technical repair before expansion. | No reported AI visibility or revenue outcome. |

FIGURE 9

## ZipRecruiter case study reported gains across search and AI measures

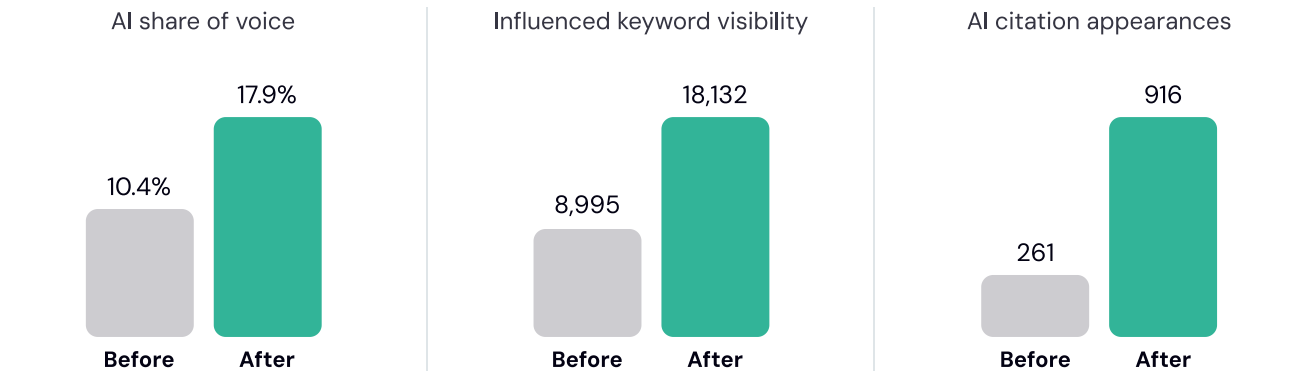


**Figure 9 ZipRecruiter case study measures**

Source R13. Provider-reported programme from October 2025 to July 2026, updated 9 September 2026. Visibility value is excluded from this chart.

The useful lesson is the architecture: map high-intent prompt themes, establish a baseline, create a persistent set of relevant sources, and track search and AI outcomes separately. The case does not prove that 962 third-party engagements caused the reported gain, nor that the same tactic is safe or effective in another category.



FIGURE 10

## One B2B SaaS case reports influence rising down the funnel

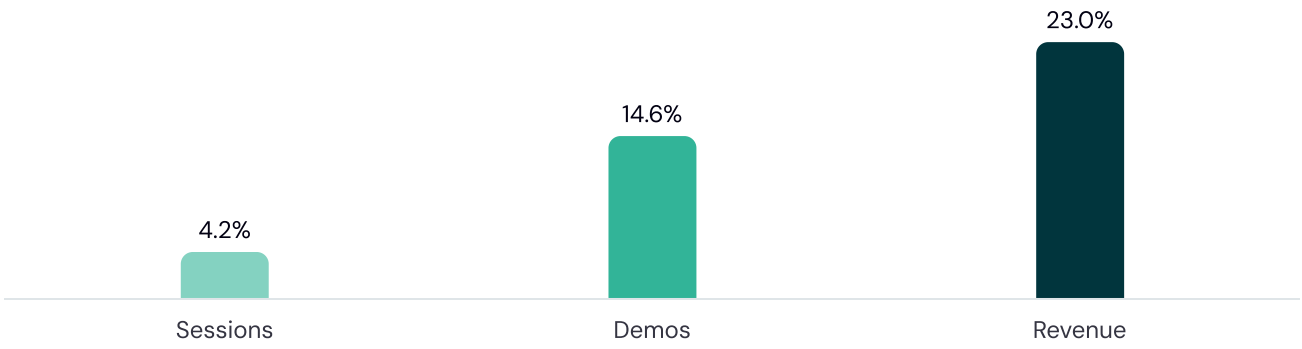


**Reported share attributed to AI referrals**

**Figure 10 Reported AI share by funnel stage**

Source R15. Anonymous B2B SaaS, November 2025 to June 2026. Provider-reported and not independently audited.

This case aligns with the Munro buyer personas because it values revenue over traffic. The client reportedly built about 25 decision-stage pages, refreshed 154 pages, added 22 G2 reviews, and gained about 365 linking sites. AI referrals represented 4.2% of organic sessions, 14.6% of demos, and 23% of revenue. Those figures are commercially interesting, but the blended programme prevents single-tactic attribution.

READ IT AS *A pattern to test in your own category — not a conversion multiplier to quote.*



FIGURE 11

## Three agency cases report higher conversion from low volume AI traffic




**Figure 11 Reported conversion rate in three 90 day cases**

Source R14. Agency-reported GA4 data for anonymous clients. AI session counts are small, so rates can be volatile.

These cases suggest a testable hypothesis: AI-referred visitors may arrive later in the decision process and convert at a higher rate. The sample does not justify a general benchmark. Track your own count of AI sessions, qualified actions, opportunities, and revenue before quoting a conversion multiplier.



FIGURE 12

## Original research case reported compounding authority and search growth

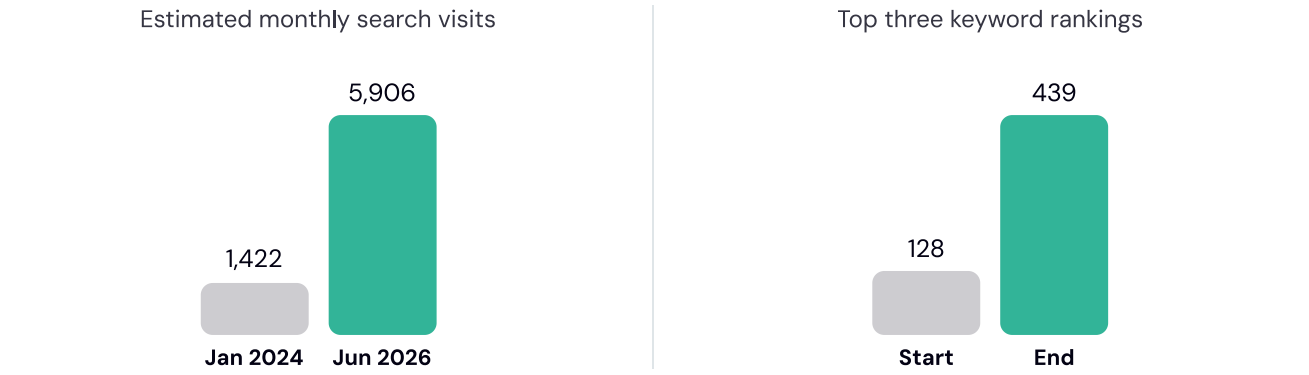


**Figure 12 Original research case measures**

Source R16. Provider-reported case using Semrush traffic estimates. The programme produced recurring research and cannot isolate AI search effects.

The transferable pattern is strong. Original research creates an asset that can earn editorial links, support expert claims, rank for the topic, and supply evidence to AI answers. For a niche B2B company, one narrow benchmark based on proprietary operational data is often more defensible than a large survey with weak relevance.



# The search visibility operating system.

***Build in the order that reduces measurement risk and compounds authority.***

## Workstream one · Commercial baseline

Define a qualified conversion with Sales before content is commissioned. Exclude known pollution, such as consumer or patient enquiries entering a B2B form. Capture source, landing page, first known touch, self-reported discovery, opportunity stage, expected value, and closed revenue.

## Workstream two · Technical eligibility

Verify indexation, canonicalisation, status codes, internal links, sitemap quality, JavaScript rendering, speed, accessibility, and content parity between visible pages and structured data. Do not launch a rebuild without an inventory, redirect map, baseline, and post-launch monitoring plan.

## Workstream three · Decision content

Create pages that help the buyer decide: problem diagnosis, method comparison, category explanation, use-case fit, cost and implementation questions, case evidence, and an outcome-led service page. A page should answer a coherent buyer task, not a collection of keywords.

## Workstream four · Distributed proof

Align the description of the company, category, audience, outcomes, and evidence across the website, review platforms, partner pages, event profiles, media, and expert biographies. Seek real corroboration; do not manufacture mentions or create a Wikipedia page as a marketing asset.

## Workstream five · Measurement and learning

Use Search Console and Bing Webmaster Tools for first-party search and citation visibility where available. Add a version-controlled prompt panel for other engines. Join those signals to analytics and CRM results, then use controlled page or distribution tests to learn what changes locally.

# Priority decision rules

| PRIORITY | CONDITION | DECISION RULE |
|---|---|---|
| **Fix now** | Tracking broken, key pages blocked, indexation unstable, wrong ICP traffic dominates | Do not scale content until the baseline is trustworthy. |
| **Build next** | Foundation sound but decision questions lack a strong owned answer | Publish the smallest complete category and proof set. |
| **Distribute** | Owned pages are strong but the public evidence is thin or inconsistent | Earn reviews, partner citations, expert coverage, and real references. |
| **Experiment** | Baseline stable and enough volume exists to detect a useful change | Change one material variable and predefine the success signal. |
| **Defer** | Tactic has no clear mechanism, owner, or measurement route | Log it in the evidence backlog, not the delivery plan. |



# A content system for categories buyers do not name.

*Organise content around buyer decisions and the evidence required to make them safely.*

## The category coverage map

| DECISION TASK | FOUNDER QUERY | MARKETING DIRECTOR QUERY | BEST ASSET |
|---|---|---|---|
| **Diagnose** | Why did our demos fall even though traffic rose | Why is attribution reading direct | Diagnostic guide plus analytics checklist |
| **Define** | What category are we really in | Which problem language has measurable demand | Category hub plus language map |
| **Compare** | Agency vs freelancer vs internal hire | SEO platform vs specialist partner | Transparent comparison with fit criteria |
| **De-risk** | Will a rebuild lose our rankings | How do we keep attribution clean during migration | Migration method and risk controls |
| **Prove** | Show me a similar complex product | Show qualified pipeline, not activity | Case evidence with baseline, method, outcome, limits |
| **Act** | What can be fixed this month | What can the team deliver this quarter | Prioritised plan with owners and milestones |

# The minimum viable authority set

1. One category hub that defines the problem, audience, approaches, evidence, and next steps.
2. Three comparison pages built around choices buyers actually make.
3. Three proof pages that document a baseline, intervention, measurable outcome, and limitations.
4. Five question pages covering objections that regularly delay or kill a sale.
5. One original research or benchmark asset with a transparent method and reusable charts.
6. One service page per distinct commercial job, with fit, process, proof, and a clear action.

# Decision page blueprint

| SECTION | PURPOSE | QUALITY TEST |
|---|---|---|
| **Direct answer** | Resolve the buyer's question in plain language. | Can Sales use the opening in a call without rewriting it? |
| **Fit and exclusions** | State who should and should not choose the option. | Does it protect the ICP rather than maximise traffic? |
| **Decision criteria** | Make the comparison dimensions explicit. | Are criteria linked to buyer risk and outcomes? |
| **Evidence** | Show data, cases, named expertise, and external sources. | Can each material claim be checked? |
| **Tradeoffs** | Explain conditions, costs, and limitations. | Would a sceptical buyer call it fair? |
| **Next step** | Offer the smallest credible action. | Is the CTA appropriate to the reader's stage? |

Google's current guidance says not to create pages for every fan-out variation or rewrite content purely for an AI system. Organise pages for people, use headings where they improve navigation, and create non-commodity material based on first-hand expertise. [R12]

# Technical requirements and useful controls.

*Technical clarity is an eligibility condition, not a substitute for relevance and proof.*

| AREA | REQUIRED CHECK | EVIDENCE TO RETAIN |
|---|---|---|
| **Indexation** | Priority URLs indexed and eligible to show snippets | URL inspection exports and coverage baseline |
| **Crawl** | Important pages reachable through stable internal links | Crawl map, depth, orphan report, status codes |
| **Canonical** | One intended version for each decision asset | Canonical audit and duplicate cluster decisions |
| **Rendering** | Visible text, links, and proof available after rendering | Rendered HTML comparison and JavaScript error log |
| **Speed** | Templates meet practical performance targets on real devices | Field and lab metrics by template |
| **Structured data** | Markup matches visible content and supported rich-result use | Validation results and deployment owner |
| **Migration** | Complete inventory, mapping, redirects, analytics, and rollback controls | Prelaunch baseline and daily postlaunch checks |
| **Analytics** | AI referrals, search, forms, CRM stages, and offline conversions connected | Data dictionary and test transactions |

# Current myth check

| CLAIM | CURRENT EVIDENCE | DECISION |
|---|---|---|
| **Every site needs llms.txt to appear in Google AI search** | Google says it ignores llms.txt for visibility and ranking. | Do not make it a Google priority. Maintain it only for another documented consumer. |
| **Content must be split into tiny chunks** | Google says there is no chunking requirement; the best length depends on the audience and task. | Use coherent sections and useful headings, not artificial fragments. |
| **Special schema causes AI citations** | Google requires no special schema. A matched Ahrefs study found no material citation lift after pages added JSON-LD. | Use supported structured data for clarity and rich-result eligibility, not as a citation promise. |
| **More mentions always improve AI visibility** | Google warns against inauthentic mentions; recent B2B studies show source effects vary by engine. | Earn relevant proof and track source quality. |
| **Wikipedia is a standard brand tactic** | Wikipedia requires independent notability and neutral editorial governance. | Do not create or control a page as a marketing deliverable. |

The Ahrefs schema analysis tracked 1,885 pages that added JSON-LD and compared them with 4,000 controls. It found a -4.6% relative effect in AI Overviews and effects statistically indistinguishable from zero in AI Mode and ChatGPT. The study predates the rest of the evidence in this guide but directly tests a claim made in older guidance. [R18]

THE RULE *Technical clarity gets you into the room. Relevance and proof get you named.*

# Entity clarity and distributed proof.

*The objective is a consistent, verifiable public description of the company and what it can credibly deliver.*

| EVIDENCE SURFACE | WHAT SHOULD BE CONSISTENT | WHAT SHOULD REMAIN INDEPENDENT |
|---|---|---|
| **Website** | Name, category, audience, capabilities, locations, people, proof | Honest tradeoffs and qualification criteria |
| **Review platforms** | Correct product identity and category | Customer language, ratings, and criticism |
| **Partner pages** | Relationship, integration, geography, and joint outcome | Partner's own description and editorial control |
| **Expert profiles** | Role, subject expertise, publications, and talks | Platform-specific voice and genuine contribution |
| **Media and events** | Accurate facts and attributed expertise | Editorial judgement and independent context |
| **Directories** | Category, location, services, and contact details | Verified customer feedback and comparison context |

## The proof inventory

- Customer outcome data with the baseline, time period, denominator, and attribution method.
- Named specialists with credentials that are relevant to the claim, not decorative biography.
- Product or service documentation that answers implementation and risk questions.
- Independent reviews and partner references that can be checked.
- Original datasets or benchmarks with a method, sample, limitations, and update date.
- Consistent company facts across regional sites and profiles, with an owner for corrections.

# Regional and global governance

For a distributed organisation such as Scotwork, the network can amplify authority or fragment it. Define a global core for the name, category, differentiators, evidence standards, and measurement. Allow territories to localise examples, regulations, language, and calls to action. Audit drift quarterly and route changes through a named owner rather than forcing every region into identical copy.

| GLOBAL CORE | LOCAL ADAPTATION | GOVERNANCE CONTROL |
|---|---|---|
| **Company and service definitions** | Market vocabulary and language | Approved fact sheet with change log |
| **Primary proof standards** | Local cases and credentials | Case template and evidence review |
| **Technical requirements** | Regional CMS and hosting realities | Automated crawl and schema validation |
| **Prompt taxonomy** | Local buyer questions and competitors | Core panel plus controlled local extension |
| **Reporting definitions** | Territory targets and pipeline values | Shared data dictionary and monthly QA |

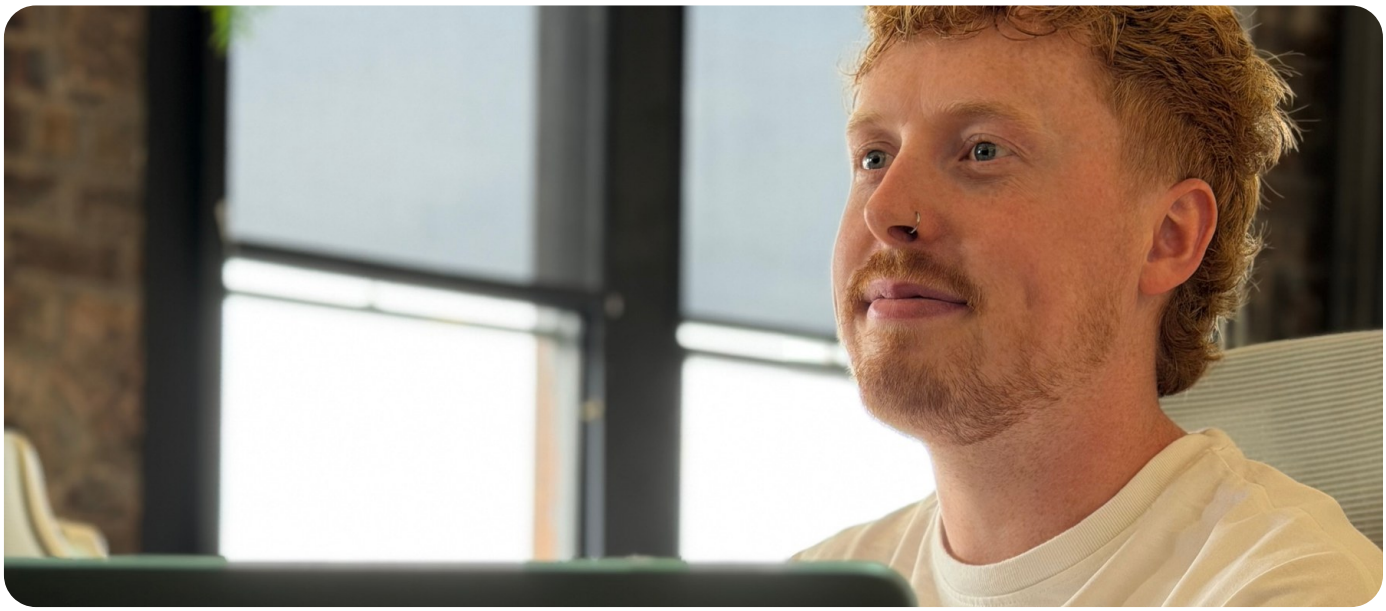

# Measurement that survives a board meeting.

*Report the commercial chain, preserve uncertainty, and prevent the prompt panel from moving under the result.*

## The metric ladder

| LAYER | METRIC | WHAT IT PROVES | WHAT IT DOES NOT PROVE |
|---|---|---|---|
| **Eligibility** | Indexed priority pages, crawl success, template health | Content can enter retrieval | That it will rank or be cited |
| **Search** | Impressions, rank distribution, nonbrand clicks, qualified conversions | Organic demand capture | AI influence before the click |
| **AI visibility** | Mention rate, citation rate, cited pages, answer role, accuracy | Observed presence on a fixed panel | Real audience volume or revenue |
| **Behaviour** | AI referrals, engaged visits, branded search, direct traffic trend | Some downstream response | Complete attribution |
| **Pipeline** | Qualified leads, demos, opportunities, value, revenue | Commercial outcome | Single-touch causality |

THE NARRATIVE *Explain the chain. Do not celebrate a score.*

# Prompt panel design

| PROMPT FAMILY | EXAMPLE | WHY IT MATTERS | WEIGHT |
|---|---|---|---|
| **Problem diagnosis** | Why are demos falling despite more traffic | Captures unbranded pain before a category is chosen | 20% |
| **Category education** | What type of software solves this continuity problem | Tests category association | 15% |
| **Best option** | Best platform for a regulated midmarket team | Tests shortlist inclusion | 20% |
| **Comparison** | Provider A vs Provider B for global deployment | Tests competitive framing | 20% |
| **Implementation and risk** | How long does rollout take and what can fail | Tests proof and operational depth | 15% |
| **Brand validation** | Is this company credible and who uses it | Tests entity accuracy and trust evidence | 10% |

Freeze the panel for at least one reporting cycle. Store the exact prompt, engine, account state, location, date, model or surface, and whether live web retrieval was used. Run repeated observations rather than a single prompt. Report the median or rate with the numerator and denominator. Treat personalisation and stochastic output as measurement conditions, not noise to hide.



# Board dashboard template

| QUESTION | MEASURE | MONTH RESULT | INTERPRETATION AND ACTION |
|---|---|---|---|
| **Are we technically eligible** | Priority pages indexed and healthy | _____ | _____ |
| **Are we visible for ICP demand** | Nonbrand commercial share of voice | _____ | _____ |
| **Are AI engines using our evidence** | Mention and citation rate by engine | _____ | _____ |
| **Is the brand described accurately** | Accurate answer rate and issue count | _____ | _____ |
| **Are visitors qualified** | Lead rate and ICP acceptance by source | _____ | _____ |
| **Is influence becoming pipeline** | Opportunities and value with search or AI evidence | _____ | _____ |

Attributed Influence Value, introduced in the supplied Munro material, can be useful as an internal concept. It is not a standard industry metric. If retained, define the formula, weighting, data source, and limitations. Do not present a weighted mention score as revenue.

TEST IT *Hand the page to the CRO without commentary. If they can read it back to the CEO correctly, it is finished.*

# The Founder 30 day plan.

***Create a trusted baseline, repair the commercial path, and publish the first complete evidence set.***

| WEEK | OBJECTIVE | ACTIONS | EXIT EVIDENCE |
|---|---|---|---|
| 1 | Find the real break | Interview Sales and founder; test analytics and forms; export ranking and index baseline; run 20 core prompts; identify non-ICP pollution. | Signed baseline, issue log, prompt panel, qualified-conversion definition |
| 2 | Protect discovery | Fix blocks, broken links, redirects, canonicals, rendering, tracking, and CRM fields; document migration risk if rebuilding. | Priority URLs healthy; test leads reach CRM with correct source |
| 3 | Make the product legible | Rewrite the core service page; publish one category hub, one comparison, and one case proof page; align company facts. | Pages indexed; Sales approves accuracy and fit language |
| 4 | Add proof and measure | Request appropriate reviews; update partner and expert profiles; repeat prompt panel; report early movement and next backlog. | First monthly scorecard with caveats and 60-day priorities |

# Founder decision rules

- Do not approve a rebuild without a URL inventory, redirect ownership, analytics test plan, and rollback path.
- Do not publish agency language that changes the product meaning. The founder or product owner signs off the category and proof statements.
- Do not use traffic as the main success measure. Report qualified enquiries, demos, pipeline, and the pages or prompts that assisted them.
- Prefer a small complete evidence set over thirty thin pages created to look active.

# The first monthly report

| RESULT | BASELINE | CURRENT | WHAT CHANGED | NEXT DECISION |
|---|---|---|---|---|
| **Qualified enquiries** | _____ | _____ | _____ | _____ |
| **Demos booked** | _____ | _____ | _____ | _____ |
| **Commercial pages indexed** | _____ | _____ | _____ | _____ |
| **Top ten nonbrand terms** | _____ | _____ | _____ | _____ |
| **AI mention rate on 20 prompts** | _____ | _____ | _____ | _____ |
| **Tracking defects open** | _____ | _____ | _____ | _____ |

# The Marketing Director 90 day plan.

*Build a system the team can operate and the executive group can evaluate.*

## The executive narrative

A defensible board update explains the chain rather than celebrating a score. For example: commercial search visibility improved on the target problem cluster; Google generative impressions concentrated on three decision pages; ChatGPT mentions remain weak; AI referrals are small but have produced two accepted opportunities; and the next quarter will extend proof and third-party coverage while protecting conversion pages.

FIGURE 13

### The 90 day operating sequence

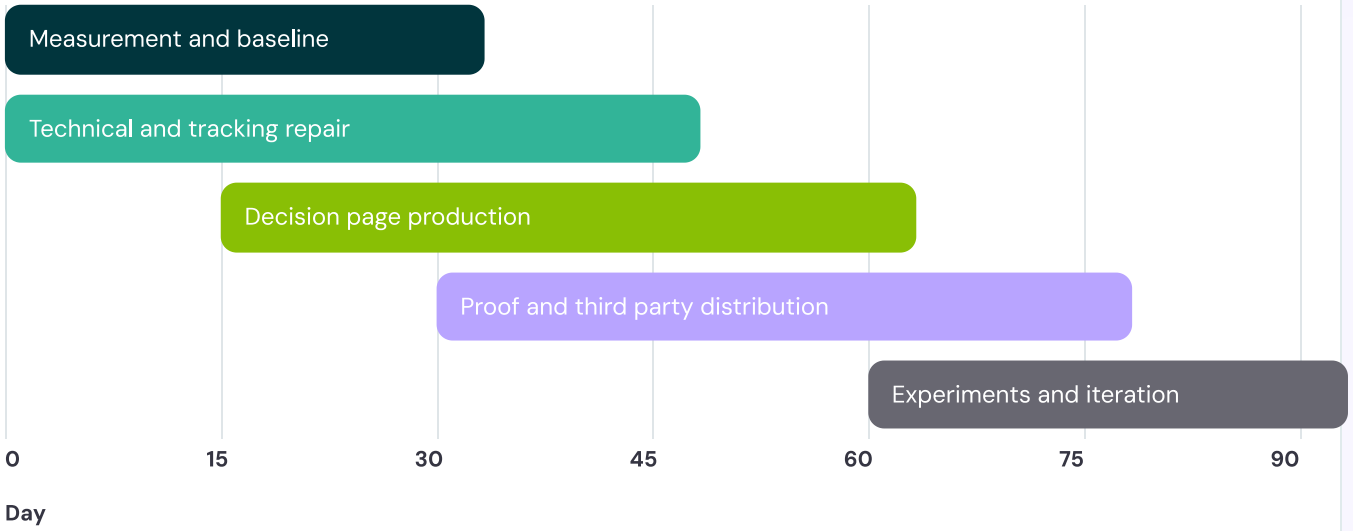


**Figure 13 Ninety day delivery sequence**

Munro recommended operating sequence based on the supplied personas and current evidence.

## Days 1 to 30 · Baseline and repair

- Agree the ICP, excluded demand, qualified conversion, opportunity stages, and attribution view with Sales and Finance.
- Audit analytics, consent, CRM routing, offline conversions, search reports, indexation, and template health.
- Build a 50-prompt panel across problem, category, shortlist, comparison, risk, and validation questions.
- Benchmark search share, AI mentions, citations, cited sources, answer role, accuracy, and competitor presence by engine.

## Days 31 to 60 · Evidence production

- Repair or consolidate the pages that already have search visibility but weak commercial relevance.
- Publish the minimum viable authority set around one category or problem cluster.
- Create a case-evidence pipeline with Sales, Customer Success, Product, and Legal.
- Improve review coverage and correct inconsistent descriptions on the third-party sources engines already use.

## Days 61 to 90 · Distribution and learning

- Pitch the original evidence to relevant publications, partners, associations, and experts.
- Run two controlled tests, such as evidence depth on matched pages or a comparison-page distribution test.
- Repeat the fixed prompt panel with multiple observations and report engine-level changes.
- Join visibility to accepted leads, opportunities, and pipeline; present the next-quarter investment decision.

# Ninety day exit evidence

| EVIDENCE AREA | MINIMUM STANDARD | OWNER |
|---|---|---|
| **Commercial baseline** | Qualified conversion, exclusions, stages, values, and attribution view agreed | Marketing, Sales, Finance |
| **Technical foundation** | Priority URLs healthy; analytics and CRM routing verified with test transactions | SEO, developer, operations |
| **Decision content** | One complete problem or category cluster published with proof and fit criteria | Product marketing and experts |
| **Distributed proof** | Priority reviews, profiles, partners, and factual inconsistencies addressed | Customer and partner marketing |
| **AI measurement** | Fixed 50-prompt panel repeated by engine with citations, role, and accuracy logged | SEO or insights lead |
| **Commercial signal** | Accepted leads and opportunities joined to search and AI evidence with caveats | Marketing operations |



# Experiments worth running.

*Use local tests to replace universal claims with evidence from your category and buyers.*

| HYPOTHESIS | TEST DESIGN | PRIMARY SIGNAL | GUARDRAIL |
|---|---|---|---|
| **Decision pages outperform generic thought leadership** | Match page groups by existing demand and links; publish one complete decision page per group. | Qualified conversions and AI citation rate | Organic clicks and indexation stability |
| **Original data earns distributed authority** | Publish one benchmark with a clear method; compare earned references with a standard guide. | Relevant linking domains and cited-source growth | Lead quality and factual corrections |
| **Corroboration improves shortlist inclusion** | Correct priority profiles and earn verified reviews in one segment while holding others constant. | Brand mention rate on shortlist prompts | Review authenticity and platform policy |
| **A stronger first answer improves extraction** | Rewrite openings on matched pages without artificial keyword repetition. | Passage citation and search engagement | Conversion rate and user satisfaction |
| **Technical cleanup releases hidden assets** | Repair internal links, redirects, and duplicate clusters by template cohort. | Indexed pages, crawl efficiency, and impressions | No loss on protected URLs |

# Experiment record

| FIELD | ENTRY |
| --- | --- |
| **Decision to inform** | _____ |
| **Hypothesis and mechanism** | _____ |
| **Control or comparison** | _____ |
| **Start and end date** | _____ |
| **Primary measure** | _____ |
| **Guardrail measure** | _____ |
| **Minimum useful change** | _____ |
| **Result and uncertainty** | _____ |
| **Scale stop or revise decision** | _____ |

Do not change the prompt panel, measurement tool, page set, and content at the same time. When a platform changes its interface or reporting, note the break in the series. A cleaner experiment with a smaller sample is often more valuable than a larger dashboard whose definition moves each week.

# The operating cadence.

*The work compounds when responsibility, evidence, and reporting are explicit.*

| CADENCE | REVIEW | OWNER | DECISION |
|---|---|---|---|
| **Weekly** | Crawl and index alerts, tracking failures, prompt anomalies, major answer inaccuracies | SEO and marketing operations | Repair urgent defects and annotate platform changes |
| **Monthly** | Search visibility, AI mentions and citations, content output, qualified actions, opportunity creation | Marketing Director | Reallocate work across technical, content, and distribution |
| **Quarterly** | Category position, competitor source set, case pipeline, original research, channel economics | Marketing and Sales leadership | Choose the next problem cluster and budget |
| **Twice yearly** | Information architecture, entity consistency, regional drift, attribution design | Executive sponsor and agency lead | Approve structural changes and governance updates |



# Responsibility map

| DELIVERABLE | ACCOUNTABLE | CONTRIBUTORS | APPROVAL |
|---|---|---|---|
| **Commercial baseline** | Marketing Director or Founder | Sales operations, Finance, agency | CEO or CRO where material |
| **Technical backlog** | SEO lead | Developer, analytics, CMS owner | Product or website owner |
| **Category and comparison pages** | Product marketing | Subject experts, Sales, SEO | Product owner and Legal if required |
| **Case evidence** | Customer marketing | Customer Success, Sales, analyst | Customer and Legal |
| **Prompt panel and AI reporting** | SEO or insights lead | Sales, product marketing, agency | Marketing Director |
| **Quarterly investment decision** | Marketing Director or Founder | Agency, Finance, Sales | Executive sponsor |

# A twelve month sequence

| QUARTER | PRIMARY OUTCOME | WHAT COMPOUNDS |
|---|---|---|
| **Q1** | Clean baseline and one complete authority cluster | Technical health, shared definitions, first decision pages |
| **Q2** | Distributed proof and first original benchmark | Reviews, links, citations, expert recognition |
| **Q3** | Second category cluster and regional or segment expansion | Internal linking, reusable case evidence, broader prompt coverage |
| **Q4** | Optimised portfolio and budget decision | Known page formats, measured source set, pipeline attribution history |

# The action scorecard.

*Use this page to choose the next move without turning visibility into a vanity programme.*

| QUESTION | RED | AMBER | GREEN |
|---|---|---|---|
| **Can we trust the data** | Forms or CRM routing broken | Known gaps documented | Test leads and exclusions verified |
| **Can engines find the pages** | Priority pages blocked or orphaned | Eligible but inconsistent | Crawl and index coverage stable |
| **Does the category make sense** | Internal language only | Some buyer language mapped | Problems, alternatives, and fit are clear |
| **Is the content decision grade** | Generic summaries | Useful but thin proof | Original evidence, tradeoffs, and next step |
| **Is proof distributed** | Owned claims only | A few relevant references | Consistent reviews, partners, media, and experts |
| **Is AI visibility measurable** | Ad hoc screenshots | Tool score without controls | Fixed panel, repeated runs, engine-level rates |
| **Does visibility reach pipeline** | Rankings only | Leads tracked but quality unclear | Accepted opportunities and revenue connected |

# The budget order

1. Repair measurement and technical eligibility.
2. Strengthen the highest-value decision journeys already near visibility.
3. Create original proof that competitors cannot copy cheaply.
4. Distribute that proof through credible, relevant third parties.
5. Run controlled AI visibility tests by engine and prompt family.
6. Scale only the work that improves qualified demand or a validated leading indicator.



THE QUESTION THAT MATTERS

***The most useful question is not whether SEO or GEO wins. It is whether the company is consistently present and credible at the exact points where an ideal buyer defines the problem, builds a shortlist, checks the evidence, and decides to act.***

# Sources and research notes.

*Research checked to 12 September 2026. Case-study outcomes remain provider-reported unless stated otherwise.*

## Our source material


A1 Munro Agency. SEO and AI Search How They Work Together. 17 pages. Created 22 May 2026.

A2 Microsoft Advertising. From Discovery to Influence A Guide to AEO and GEO. 16 pages. Supplied PDF.

A3 Rupert Morris. Mastering AI Visibility for B2B Services. 21 pages. Supplied PDF.

A4 Munro Agency. Scotwork AEO GEO. 27 pages. Created 10 February 2026.

A5 The Munro Agency. Buyer Personas. Working draft September 2026, based on client and prospect calls January to September 2026.


## Web research


R1 Semrush. AI search and manufacturing SEO What the data shows. 8 September 2026.

R2 Google Search Central. Latest documentation updates. updated 8 September 2026.

R3 Google Search Central. Introducing Search Generative AI performance reports in Search Console. 3 June 2026 with worldwide rollout note 31 August 2026.

R4 Hog the Web. What 366000 Google AI Search impressions reveal about SEO. 27 August 2026.

R5 Wang and colleagues. AI in Search Reduces Publisher Referrals Without Improving User Experience. 18 August 2026.

R6 Ahrefs. AI adjusted volume and the challenge of estimating AI Search demand. 17 August 2026.

R7 Chapekis and colleagues. Investigating Click Behaviors on Google Search Result Pages That Produce an AI Overview. 5 August 2026.

R8 Analyze AI. The State of AI Search in B2B. 4 August 2026.

R9 Broadcastwell Research. The 2026 State of Generative Engine Optimization. July 2026 with data collected 18 to 23 July.

R10 Olivier Martinez. Optimizing Visibility in Generative Engines A Critical Survey. 15 July 2026.

R11 Semrush. AI Overviews are expanding across commercial intent search. 2 July 2026.

R12 Google Search Central. Optimizing your website for generative AI features on Google Search. updated 10 July 2026.

R13 CiteWorks Studio. How ZipRecruiter increased AI share of voice by 72 percent. published 3 August and updated 9 September 2026.

R14 TrioSEO. B2B and SaaS SEO case studies. datasets ending 27 to 30 August 2026.

R15 VisibleIQ. How a 10M SaaS generated 1 point 1M in ARR from organic and AI search. 19 July and updated 12 August 2026.

R16 PureLinq. How a B2B financial brand turned original research into 1300 links. 2026.

R17 CiteWorks Studio. Mattress Advisory technical SEO case study. 21 August and updated 24 August 2026.

R18 Ahrefs. We tracked 1885 pages adding schema and AI citations barely moved. 11 May 2026.

R19 Bing Webmaster Blog. Introducing AI Performance in Bing Webmaster Tools. 10 February 2026.

R20 BrightEdge. Weekly AI Search Insights. accessed 12 September 2026.

R21 TechRadar. Why AI visibility now demands paid and organic GEO optimization. 7 September 2026 opinion.

R22 The HOTH. We ran AI search campaigns for a year and there was no AI hack. 3 September 2026.


# Limitations

- AI outputs vary by date, engine, model, location, account state, prompt wording, and retrieval mode.
- Several current studies are vendor-funded and use proprietary datasets that cannot be fully reproduced.
- The two strongest recent user-behaviour studies concern Google and US participants; they do not directly estimate B2B conversion behaviour.
- Recent case studies often combine technical SEO, content, digital PR, reviews, and measurement. They do not isolate the effect of an individual tactic.
- This guide provides a strategy and operating model, not a forecast of rankings, citations, leads, or revenue.



WHAT TO DO NEXT

# Start with the *baseline.*

One search visibility system, not two disconnected programmes. We will run the diagnostic on your stack, site, evidence footprint and reporting, and tell you plainly where the chain breaks — and what we would fix first.

Book a Pipeline Growth Plan. Bring your Sales lead.

Get your pipeline growth plan →

munro.agency

Munro. agency

Munro Agency · Glasgow, Scotland · Working with B2B software and services companies in the UK and US

© 2026.