



variant1

What Your AI Audit Is Missing

How to connect AI visibility to your ideal customers before you plan 2027

For marketing leaders at B2B SMEs

Includes a sample AI visibility audit

Contents

Who this book is for	3
Part 1 The AI audit	4
1 What an AI audit is, and what each part is for	5
2 How to use the audit in your plan	10
3 Audit mistakes to avoid	13
Part 2 Research	15
4 8 connected areas	16
5 Market segmentation and targeting	18
6 Buying process and buyer journey	20
7 Who AI names in your place	23
8 How the research gets done	26
Part 3 ICP	28
9 The ideal customer company	29
10 The buying committee, buyer profiles and personas	31
11 From ICP to the question set	33
12 From ICP to marketing and sales	35
13 Your 2027 plan	37
What a sample AI visibility audit looks like	40
Sources	45

Who this book is for

For marketing leaders at B2B SMEs who want to understand what an AI visibility audit tells them, what it misses, and what to do next.

An audit can show whether AI engines name your company. But are you testing the questions your priority buyers ask? And which gaps matter enough to act on? Without market and buyer research, a visibility score alone cannot answer those questions.

Preparing your marketing for AI should start with understanding your market and buyers. That understanding determines which questions to test, which findings matter and where to invest. This book explains how to connect an AI visibility audit to that research and use the findings to make better marketing decisions.

The book has three connected parts. First, the AI audit: what it reveals and what it cannot tell you. Next, the research needed to judge which findings matter. Finally, the ICP: how that research defines your priority customers and buyers, shapes the audit questions and guides your marketing plan.

- **The AI audit.** What it measures, how to read it, and how it feeds the plan.
- **Research.** The 8 connected areas that establish which markets, buyers, buying situations and questions matter, so you can judge the relevance of audit findings.
- **The ICP.** A closer look at customer fit, the buying committee, buyer profiles and personas, turning the research into audit questions and marketing priorities.

Our client work starts with market and buyer research, including defining or refining the ICP. The audit comes first in this book because most teams start there. In practice the audit runs on a starter question set drawn from existing customer evidence, the research and the ICP refine that set, and the audit is re-run on the refined version.



Kevin Hayes
Managing Partner

Leads marketing and execution. Twenty years of growth systems, media budgets and sales and marketing capacity for global and regional brands.

kevin@variant1.com

linkedin.com/in/kevinhayesca



Mika Nurminen
Managing Partner

Leads research and strategy. Intelligence and commercial strategy at global enterprises, and go-to-market work for Fortune 500 brands, startups and PE-backed firms.

mika@variant1.com

linkedin.com/in/mnurminen

Part 1

The AI audit

What an audit measures, what it cannot decide, and how research prevents the wrong marketing investment.

- 1 What an AI audit is, and what each part is for
- 2 How to use the audit in your plan
- 3 Audit mistakes to avoid

What an AI audit is, and what each part is for

6 checks, 3 numbers, the sources traced, and a sample of what comes back for one firm.

An audit can show that your company is absent from an AI answer. It cannot decide whether that question matters to your growth, which buyers you should pursue, or whether a missing mention deserves investment. Those decisions require market and buyer research. Otherwise, marketing can spend its budget improving visibility among buyers the business does not want.

A standard Variant1 baseline tests 25 buyer questions once on three selected engines: 75 answers. It is an exploratory snapshot, not a stable ranking. Existing customer evidence informs the starter set. The fictional Fenwick Hale Advisory example runs throughout this book.

The 6 checks

The six checks organize the findings. The sample uses transparent pass rates, explained on page 8. Relevance to your target buyers comes before the overall score. These checks form the ANSWER framework used in our audit.

- **Appearance.** Does the engine name you on an open category question? Examples: “Which consulting firms help manufacturers improve operations?” and “Best operations consultants for a manufacturer in the Pacific Northwest.” Inspect the cited sources and your eligibility for them. Missing coverage is a lead to investigate, not proof of why the engine omitted you.
- **Nomination.** Does the engine put you forward for a buying situation? Example: “We are a manufacturer with shrinking margins. Who should we hire?” Check whether the situation matches your priority customers before treating absence as a problem.

- **Showdown.** How are you described against a named competitor? Example: “Fenwick Hale vs Corvane Partners for an operations review?” Verify the claims, then check whether the distinction matters to buyers and whether credible evidence supports it.
- **Worth.** Can the engine explain your pricing and value? Sample: “How much does a mid-market operations engagement cost?” One issue the check surfaces is a figure for competitors and none for you.
- **Expertise.** Does the engine describe your services and customer focus accurately? Example: “What services does Fenwick Hale Advisory offer?” Check cited pages and profiles against current records. An incorrect answer does not establish where the error originated.
- **Reputation.** Is the credibility assessment supported by verifiable evidence? Example: “What do clients say about Fenwick Hale Advisory?” Invented clients or results deserve attention because buyers could rely on them or challenge them.

What the audit cannot decide

Who is worth winning. A generic category question can include buyers outside your service model, geography or economics. Market attractiveness and your ability to compete determine which gaps deserve attention.

What will persuade them. A comparison answer does not establish the buyer's real decision criteria, internal objections or trusted sources. Interviews and deal evidence do.


What marketing should change. The answer could call for clearer positioning, stronger third-party evidence, a different target segment, or no action on that question. Research determines the choice.

The risk is acting on the wrong diagnosis: a better score without better prospects, while a gap affecting valuable buyers remains unresolved.

Why act during this planning cycle?

In Semrush's 2026 survey of 519 US B2B professionals who use AI at work, 92% said AI had shaped their vendor shortlist.[1] This is not a measure of all buyers, but it is a reason to check your own. Buyer interviews, credible customer evidence and independent coverage take time to develop. Starting now gives marketing evidence for its next budget and leadership discussion.

What a full audit looks like: the eight parts of the report



Executive summary

The overall score, the three headline numbers and the three priorities.

Scorecard


Share of voice, citation rate and accuracy, with the six category scores.

Where you show up

Every question on every engine: named or absent.

Answers buyers get

Verbatim answers, with who was named alongside or instead of you.




Six category pages

What each check found, the sample questions, the fix.




Sources traced

Every citation, grouped by type.



Accuracy errors

Each wrong claim about you, with its cited source where there is one.



90-day plan

Recommendations ranked by expected effect and effort.

The 3 numbers to report

Across the whole run, three figures summarize the result.

- **Share of voice.** Your mentions as a share of all supplier mentions across the discovery answers, counting each supplier once per answer.
- **Citation rate.** The share of answers in which your own website is cited as a source. This is different from being named, since an engine can name you from a directory without citing your site, and engines differ in how often they cite anything.
- **Accuracy rate.** The share of the engine's factual claims about you that are correct, checked against your own records, and where a claim is wrong, whether the error flatters or damages.

Read the score as a description of this test, not a verdict on the business. Fenwick Hale is rarely named in the discovery answers, while most checked claims about it are correct. Research must establish which omissions affect commercially attractive buyers.

Sample scorecard: Fenwick Hale Advisory, a 150-person management consulting firm.



Illustrative scoring only: each category is the percentage of answers meeting its stated test, rounded. The index weights are 25%, 20%, 15%, 15%, 15%, 10%, in the order shown. It is not an industry benchmark or a revenue forecast.

The answer a buyer gets

The example below is one of 33 discovery answers. Before commissioning a fix, identify the buyer, the decision and the evidence that makes it worth acting on.

How the sample scores are calculated

Appearance: named in 2 of 27 answers, 7%. Nomination: recommended in 1 of 6, 17%. Showdown: a credible option in 11 of 12 comparisons, 92%. Worth: a supported firm-specific price in 0 of 9, 0%. Expertise: services and focus accurate in 11 of 12, 92%. Reputation: credible without an invented reference in 8 of 9, 89%.

The weighted index is 42 out of 100, calculated from the unrounded rates. These tests make this fictional example reproducible. A commissioned audit should agree and disclose its own rubric, weights and treatment of missing or ambiguous answers before scoring.

Share of voice, citation rate and claim accuracy use different denominators. None establishes the number of buyers reached or the value of opportunities created.

Engine B • Nomination question


“We’re a 400-person manufacturer with shrinking margins. Which consulting firms should we talk to about an operations review?”

For a mid-sized manufacturer facing margin pressure, a few firms come up repeatedly. **Corvane Partners** focuses on operational turnarounds for manufacturers in roughly the 200 to 1,000 employee range and has published case studies with documented margin improvements. **Tallis Reed Consulting** is known for lean and supply chain work and is recommended by at least one industry association. **Hollis Kerr Group** offers fixed-fee operations reviews, typically in the \$60,000 to \$120,000 range according to its website, which may suit a company that wants a defined scope. I couldn’t find reliable pricing for the other two. It may also be worth asking peers in your industry association, since several of these firms are active there.

The firm is absent and three competitors are named with details about them, one with a price. Sample answer, invented.

Where the engines get their evidence

The audit records every source the engines cite, noting that some engines cite more than others and some answers cite nothing. In the sample, the 75 answers produced 214 citations: 171 to third-party pages such as directories, roundups and review sites, 29 to the firm's own site, and 14 to competitor pages. Most of the cited evidence comes from third-party pages, and the list shows which ones the company is missing from, which is a set of leads to investigate rather than proof of why the engine answered as it did.



The 214 pages cited across the 75 sample answers, by type.

Each part of the report feeds a decision in targeting, positioning, content, PR and directory budget, sales readiness and reporting, which the next chapter covers.

Get your audit before the 2027 plan is drafted

25 buyer questions, three engines, every answer scored and every citation recorded, with the findings in time to inform the plan. Book at variant1.com.

How to use the audit in your plan

The report arrives with a score and a list of recommendations. Here is how to turn it into decisions, and what to put next to it before you act.

The purpose of the audit is to improve a marketing decision. Start with the commercial question: which customers are we trying to win, what is stopping them, and what evidence would justify changing the plan? The worked example below illustrates the difference.

Worked example: the obvious fix is not the priority

In this fictional extension of the Fenwick Hale sample, the audit finds weak visibility for broad management-consulting questions and no published price. A score-led response would commission generic category content and publish a price range. Neither action has yet been justified by buyer evidence.

Suppose account and delivery records show the firm earns repeatable margins from manufacturers facing a specific operations problem, while broad strategy projects are costly to serve. The first research decision is to prioritize the former and stop treating every category omission as equally important.

Research changes the interpretation

Interviews with target buyers, including firms that chose another supplier, then reveal a different barrier: they need evidence that improvements can be delivered without disrupting production. They use industry peers, an association and case studies to check that claim. AI is one source among those checks.

The resulting priority is a substantiated customer case and relevant independent coverage, supported by pages explaining delivery and implementation risk. A generic content expansion is deferred. Pricing information is considered separately against buyer needs and the firm's commercial model. These are invented findings, not a client case or a promise of results.

Trace each finding to a decision

The example changes the plan because the research establishes account fit and buying criteria. Use the audit to test those priorities. The outputs below feed the detailed marketing decisions in Part 3.


- **Share of voice and the grid.** Separate priority-buyer questions from generic questions. Investigate omissions in the former before assigning content work.
- **The cited sources.** Compare the directories, roundups, review platforms and trade publications the engines cite with the ones your target buyers use before putting any of them in the PR or directory budget.
- **The comparison answers.** Use them for positioning and sales battlecards: the basis on which the engine distinguishes you from each competitor, and the claims sales should be ready to correct.
- **The cost answers.** Check whether the engine has a figure for you and, if not, whose figure it's using, then decide what pricing information you're willing to publish.
- **The accuracy errors.** Correct the pages and profiles you control, in order of damage, and request corrections where the source is a third party.
- **The 3 numbers.** Record them as the baseline for the visibility measures that sit under the commercial ones.

Sort the recommendations by the buyer they affect

Record each proposed action against a target buyer, a decision barrier, supporting evidence and a commercial objective. Keep uncertain recommendations as tests. A severe factual error can still justify immediate correction while the wider research continues.

What the audit measures

What only the CRM and the ICP can tell you



The audit observes the first two links under test conditions. The rest need CRM and buyer evidence.

3 things to put next to the audit before you act

Your buyer's questions. If the audit's question set came from category keywords, it may still overlap with what your buyers ask, but that has to be checked. Put your own question set next to the audit's and compare, and treat the findings on questions your buyers wouldn't ask as lower priority.

Your real competitor set. Establish who buyers actually considered, including doing the work internally or not buying. Compare that evidence with the engine's list. An unfamiliar AI-named firm is a research lead. Chapter 7 explains how to assess it without displacing established competitors.

Your CRM and buyer evidence. Agree what qualifies an opportunity and how conversion is measured. Add a discovery question about how the buyer found and assessed suppliers, including AI where relevant. Treat self-reported attribution cautiously. A visibility change and a pipeline change occurring together do not establish causation.

Timing

Running the audit before annual planning lets its findings inform the initial budget. An audit run later in the year can still inform targeting, content, positioning and reporting as those come up for review, so the timing is a benefit rather than a deadline. The final chapter sets out the order of work.

Where human judgment matters

Our audit is not fully automated. Human specialists shape the question set around the company, its buyers and its market. They check factual claims against available evidence, interpret findings in the business context and determine which issues deserve priority.

Technology supports the process, while people remain responsible for the assessment and recommendations. Recorded answers and citations are kept separate from their evaluation, so a finding can be reviewed against the same captured evidence rather than a newly generated answer.

The framework gives the review a consistent structure. Human judgment connects what an engine returned to the buyers and decisions that matter. Market and buyer research supplies the context for that judgment.

Audit mistakes to avoid

7 ways an audit loses its usefulness, and the fix for each.

The mistakes below show how an audit can lead to the wrong allocation of time or budget. They are illustrations rather than client cases. Research provides the context, while a documented testing method makes the observations interpretable.

The report becomes the to-do list

The team completes forty reasonable fixes but cannot explain which target buyers benefit. Require a buyer, a decision barrier and an intended outcome for each priority. An easy technical fix can be worthwhile, but ease alone should not determine the marketing plan.

The questions weren't your buyer's

A category-based question set can miss important buying situations. In the worked example, broad consulting visibility would divert effort from manufacturers seeking operational improvements. Check the set against interviews, sales evidence and the chosen target, rather than assuming that either keywords or a persona prove what buyers ask.

The website gets rebuilt when the evidence pointed elsewhere

The team rewrites service pages without reviewing the sources cited in relevant answers or the evidence buyers trust. Investigate both before choosing the work. Missing directory coverage is a hypothesis to test, and a relevant website improvement can still be justified independently.

Nobody checks what changed in sales

Share of voice goes up while inquiry quality goes down, and nobody notices for two quarters because the audit was never connected to the CRM. The decline may have nothing to do with the visibility work, but without agreed measures nobody can tell. Agree the measures with sales before the first fix ships.

The audit stays in marketing

An audit finding has an owner only when a team can act on it. Give sales the claims to verify, content the buyer questions to address, and PR the credible sources to assess. Each brief should include the relevant buyer evidence and the decision it supports.

Every named firm is treated as a competitor

A familiar firm may matter in deals despite being absent from the audit. An unfamiliar firm may dominate generic answers while serving a different customer. Compare the two sets against buyer interviews and wins and losses before changing competitive priorities.

One answer is treated as a verdict

In SparkToro and Gumshoe.ai's January 2026 study, 600 volunteers ran 12 prompts a total of 2,961 times across ChatGPT, Claude and Google AI. The authors found substantial variation in lists, order and list length. Some brands nevertheless appeared frequently for a given intent.[9] This vendor-linked study is evidence against treating one answer as a stable ranking, not a universal estimate of how many repeats every audit needs.

One run per question maps breadth, not response stability. Before acting on a costly omission, repeat the relevant prompts in fresh sessions on the same engines. Record date, model or product, search mode, location and personalization settings. For trend reporting, keep a fixed benchmark set and the same testing protocol, with repeated runs each wave. Report ranges and persistent patterns, not a one-point movement as improvement. New research questions belong in a separately versioned exploratory set.

Part 2

Research

The 8 connected areas that establish where to compete, which buyers to prioritize and what drives their decisions. The audit adds evidence to that wider research.

- 4 8 connected areas
- 5 Market segmentation and targeting
- 6 Buying process and buyer journey
- 7 Who AI names in your place
- 8 How the research gets done

8 connected areas

Variant1 organizes market and buyer research around 8 areas. How deep to go in each depends on the decision it supports and the evidence already in hand.

Research should establish where the business can grow, which customers it can serve profitably and what drives their decisions. An AI audit adds a view of how selected engines represent that market. It cannot replace the evidence needed to choose a target, a position or an investment.

The eight areas organize that wider investigation. They are not eight studies every SME must commission. Begin with the decision at risk, use reliable evidence already available, and investigate the gaps that could change it. Recent, validated segmentation may need little work. A move into a new market needs evidence beyond the current customer base.

Area	What it establishes	The AI question
Market segmentation	How the market divides into groups with different needs, characteristics and buying situations.	Does AI change how this segment finds suppliers?
Targeting	Which segments deserve priority, based on attractiveness and ability to compete, with the trade-offs explicit.	Where are we already visible to the engines, and where are we missing?
Ideal customer company	The account characteristics associated with strong fit and attractive business: needs, capabilities, service requirements, economics.	When asked who we serve, does the engine's description match the ICP?
Buying committee	Who initiates, influences, evaluates, uses, approves or can block the purchase, with their authority and relationships.	Which committee members now ask AI, and what do they ask?
Individual buyer profiles	The responsibilities, needs and buying authority of each role you need to reach.	Is AI an information source or a decision aid for this role?
Buyer personas	The priorities, motivations, objections, decision criteria and trusted sources of each participant.	What does this person type into an engine, and what comes back?
Buying process	How the organization evaluates suppliers, approves spending and completes procurement.	Does AI change how suppliers are screened or proposals assessed?
Buyer journey	How people recognize a need, explore options and develop preferences before they buy.	Where does AI sit among their sources, and how do they verify it?

Segmentation and targeting decide which accounts matter. The ideal customer company, the buying committee, the buyer profiles and the personas describe those accounts and the people in them, and together they form the ICP, which Part 3 covers. The buying process and the buyer journey describe how those people move from a problem to a purchase, and where AI now sits in that movement.

What AI adds to each area

We add an AI question to every area, such as whether AI changes how this segment finds suppliers, whether AI is an information source or a decision aid for this role, and how buyers verify what AI tells them, and the research is free to find that the answer in some areas is “not much.” The audit also adds evidence the research didn't have before: you can test what selected engines return to buyer-like questions, while interviews are still needed to learn what buyers asked and saw. The chapters that follow take the areas in the order we usually work through them with clients: the market and the targets, the process and the journey, the competitors, and how the work gets done.

Market segmentation and targeting

Whether your segments still explain differences in demand, which deserve priority, and why the trade-offs have to be explicit.

A segmentation built around industry and company size a few years ago can be tidy and no longer useful. In many markets what separates buyers who convert from buyers who don't is a trigger, a buying situation, or the way they research suppliers, and segments drawn on industry and size alone won't show it. Whether that's true for your market is something the CRM and the interviews can test.

Segmentation: do the groups still differ?

Keep segment boundaries where they explain commercially meaningful differences in needs, buying behaviour, economics or how you can serve the accounts. Similar messages alone do not make two segments identical. Split or combine groups only when doing so improves a targeting or marketing decision.

Use wins, losses and service costs to identify patterns, then test them beyond your existing customers. Industry statistics, competitor offerings and prospective-buyer interviews help reveal demand and segments the CRM does not contain.[2]

If different scenario prompts return similar supplier lists, investigate the wording and repeat the test. That result describes the tested answers. It does not establish that the segments are alike, or that engines never distinguish them.

Targeting: which segments deserve priority?

Assess each segment on two things.

- **Attractiveness.** The segment's size, growth, margin, and the cost of serving it.
- **Ability to compete.** Fit, evidence, references, sales results in the segment and, as one signal among those, presence in the tested prompts written for that segment.

Make the trade-offs explicit, so marketing knows where to concentrate and what it's giving up. The output is a short ranking with the reasons attached, agreed with sales, that the rest of the research and the 2027 plan are built on.

Where the audit fits

Compare matched questions for each priority segment under the same testing protocol. Equal question counts help, but wording, intent and engine settings also matter. A strong overall result can conceal weak presence for a valuable segment. Report the segment findings alongside their sample sizes and buyer evidence.

A segment opportunity is bigger than its AI score

Estimate the reachable account base, demand, buying triggers, competitive alternatives and economics before choosing where to concentrate. U.S. Census County Business Patterns and Statistics Canada business counts provide useful starting points by industry, size and geography.[3,8] They count establishments or locations, not necessarily unique target companies or demand for your offer.

Combine those sources with current industry evidence, company research and interviews outside the installed base. State the assumptions behind the opportunity estimate. A segment with weak AI visibility can still be attractive, while a segment where you appear frequently can be commercially unsuitable.

Buying process and buyer journey

How the organization evaluates suppliers and approves spending, how the people inside it move from a problem to a preference, and where AI now sits in both.

The buying process

The buying process is the organization's side: how it evaluates suppliers, approves spending and completes procurement. It has stages, gates and documents, and it's often the part of the journey marketing knows least about, because much of it happens after the buyer has stopped engaging with marketing. Map it for the priority segment: who has to approve what, at what threshold, with what paperwork, and how long each gate takes.

Then check whether AI changes any of it. Look at the tasks that used to take a person time, such as screening a long list of suppliers down to a short one, comparing proposals, checking references, or summarizing a vendor's materials for a committee. Where a buyer hands one of those tasks to a general-purpose AI engine, the engine's description of your company can enter the process at that gate, and the audit's comparison and cost answers show what that description might be.

The buyer journey

The journey is the people's side: how they recognize a need, explore options and develop preferences. In Semrush's 2026 survey of US B2B professionals who use AI at work, 92% said AI had shaped their vendor shortlist, and when AI mentioned a vendor, 71% visited its website and 38% checked reviews.^[1] The survey covers AI users rather than all buyers, and it doesn't fix the order in which website visits, referrals and sales calls happen, but for the buyers it describes the engines are part of the process rather than a sideline.

Map the journey by decision rather than by channel. From recognizing a problem to approving a purchase, a buyer decides whether the problem is worth acting on, which approaches to consider, which suppliers to look at, which to shortlist, which to talk to, and whether to sign. For each decision, record who makes it, what they need to know, which sources they trust and where AI sits among them, what they'd type into an AI engine at that moment, how they verify AI's answers, and which page they land on afterwards.

3 stages, 6 kinds of question

For this audit, discovery, evaluation and validation provide a practical grouping of questions. Buyers can revisit these decisions, work in parallel or enter through a referral. Refine the map from evidence rather than forcing every purchase into a linear funnel.

Stage	What the buyer is deciding	Check	Sample question
Discovery	Which suppliers to look at	Appearance	"Which firms help mid-sized manufacturers improve operations?"
		Nomination	"We're a 400-person manufacturer with shrinking margins. Who should we hire?"
Evaluation	Which to shortlist and talk to	Showdown	"Fenwick Hale vs Corvane Partners for an operations review?"
		Worth	"How much does a mid-market operations engagement cost?"
Validation	Whether this one can be trusted	Expertise	"What services does Fenwick Hale Advisory offer?"
		Reputation	"What do clients say about Fenwick Hale Advisory?"

- **Discovery.** Check which sources appear in the relevant answers and whether buyers trust them. Directories, trade coverage and comparison articles are candidates to assess. Their citation does not establish why a supplier was included or omitted.
- **Evaluation.** Buyers compare, and if you haven't published a price or what sets you apart, the engine may use a competitor's figures, hedge, or leave you out.
- **Validation.** Buyers check, and stakeholders who weren't involved in the research ask their own questions.

The sample grid shows a contrast: Fenwick Hale appears in only 3 of 33 discovery answers but in most answers to questions naming it. That can be difficult to detect through inbound inquiries alone, because those buyers have already found the company. Interviews beyond the inbound pipeline help investigate who may be missing.

	Engine A	Engine B	Engine C
Which consulting firms help mid-sized companies improve operations?	Absent	Absent	Absent
We're a 400-person manufacturer with shrinking margins. Who should we hire?	Absent	Absent	Absent
Fenwick Hale vs Corvane Partners for an operations review?	Named	Named	Named
How much does a mid-market operations engagement cost?	Absent	Absent	Absent
What services does Fenwick Hale Advisory offer?	Named	Named	Named
What do clients say about Fenwick Hale Advisory?	Named	Absent	Named

Six of the 25 sample questions, one per check. Each cell is one question on one engine.

Why discovery gaps can remain hidden

A validation problem can damage a live opportunity. A discovery problem can exclude a supplier before any inquiry reaches its CRM. Neither is automatically more costly. Use buyer evidence, the value of affected opportunities and the current funnel to decide which deserves priority. The audit identifies a possible exposure, not a measured loss of sales.

Ask recent buyers what they asked

Ask recent customers, lost prospects and target buyers who never approached you how they researched suppliers. Did they use AI? If so, what did they ask, what did they do next, and what did they verify elsewhere? Invite examples or saved answers where available. Start with a manageable qualitative sample across relevant roles and situations. Ten interviews is a starting point, not a representative finding, and recall needs checking against other evidence.

Who AI names in your place

The firms an AI engine names can differ from the competitors your sales team sees. Research establishes which differences matter to buyers.

Competitive intelligence has long covered what competitors do and how the market perceives them, and there's now one more question for marketing: when a buyer asks an AI engine about your category, who gets named?

Two competitor sets

Sales teams benchmark against the firms they lose deals to, see at trade shows, or hear about from prospects, which we call the sales set. AI engines produce a different set when asked an open question about a category, and the firms in it tend to be well represented in the sources the engines cite. The set can include firms the sales team never meets, because they appear in some engine responses even if they rarely appear in your CRM, and whether buyers encounter them is a separate question for the interviews.

Comparing the two sets produces two decisions. A firm in the AI set that isn't in the sales set belongs on the competitive watch list, and a firm in the sales set that isn't in the AI set is still a competitor, with the AI absence noted as one more fact about it.


Why engines name who they name

Ahrefs studied about 75,000 established brands and found strong associations between AI visibility and brand mentions, including on YouTube. Web-mention correlations were roughly 0.66 to 0.71.[5] Muck Rack classified 84% of more than 25 million cited links across 17 industries as earned media.[6] These are different findings: one is a correlation study, the other a classification of citations. Neither proves how to improve a particular SME's visibility. Together, they justify investigating credible third-party presence alongside your own content. Prioritize sources relevant to your buyers and category, rather than pursuing mentions wherever they are available.

Share of voice, from the audit

The discovery answers show the firms named under the chosen test conditions. In this fictional sample, three firms appear in most of the 33 answers and Fenwick Hale appears three times. Repeat priority prompts to assess which names persist. A changing order does not necessarily mean the underlying competitor set changed.

Share of the 33 discovery answers naming each firm (sample audit, three engines)



Sources cited for Corvane Partners

Trade publication roundup
"Top operations consultancies"

Industry association directory
Member listing with case studies

Review platform
14 client reviews

Software partner case study
Customer story on the vendor's site

Corvane's own website
Service and pricing pages

Four of the five are third-party pages.

Fenwick Hale appears in 3 of 33 discovery answers, or 9%. Its share of the 100 supplier mentions is 3%. The six plotted firms account for 96 mentions. Other suppliers account for four. These measures describe the sample, not actual buyer shortlists or revenue. Track a fixed prompt set and disclose any changes to the supplier universe.

Trace the sources

For each firm that keeps appearing, the audit records the sources the engines cited alongside it: a directory, a best-of roundup in a trade publication, a review platform, an association member list, a partner's case study, a journalist who covers the category. Check the same sources for your own company. Where a competitor is listed and you aren't, that's a specific item to assess for the PR and directory budget, weighed on whether your buyers use the source, what it costs, and whether you're eligible.

Read the comparisons the buyer sees

For each relevant competitor, compare the engine's claims with its actual offer, customer focus, pricing evidence and delivery model. Then compare those distinctions with the decision criteria from interviews. A precise AI description can still be wrong, and a visible difference may not matter to your buyer.

Watch the set over time

Monitor firms that recur in relevant answers alongside sales intelligence and market developments. A new name is a reason to investigate its offer and customer fit, not evidence that it is winning your deals. Record changes in cited sources as observations and test whether buyer exposure or commercial relevance has changed.

Assess more than the sources an engine cites

Competitive research should also establish what the rival actually offers, where it competes, which customers it serves and what evidence supports its claims. Review pricing, capabilities, delivery requirements, customer references and market activity. A citation list covers only the material surfaced in that test.

Then examine how the rival presents the offer across its website, campaigns, content and independent coverage. Differences between the service it sells and the benefits it emphasizes can reveal a positioning question worth testing with buyers. The output is a decision about how to compete, rather than a list of sites to copy.

How the research gets done

Market evidence, customer economics and buyer research, connected to a decision.

The research needs both an inside and an outside view. Existing accounts reveal what has worked. External evidence tests where else demand exists and whether the pattern can be repeated. Choose the sources according to the decision, and record their coverage and limitations.[2]

- **Market evidence.** Industry statistics, demand indicators and company lists establish the scale and shape of an opportunity. Check definitions, dates and geographic coverage.[3,8]
- **Account and deal records.** Compare margin, retention, service costs, wins and losses. Distinguish a profitable pattern from a few unusually successful accounts.
- **Buyer interviews.** Include customers, lost prospects and prospective buyers outside your base, across relevant roles. Ask about the decision, alternatives and information sources before probing AI use.
- **Competitor research.** Examine offers, target customers, pricing evidence, delivery, references and channels. The sales list and AI list are inputs, not the full market.
- **Digital evidence.** Search queries, website behaviour and campaign data help test language and observed demand. Their users and coverage can differ from the whole market.
- **The audit.** Record tested answers and citations. They show AI outputs, not market size or buyer behaviour.

Schedule, owner and record

Agree the decisions the research must support and a realistic scope. For an SME, begin with one priority segment and the gaps most likely to change the plan. Assign an owner, a completion date and access to the required records and buyers. Expand only when the next decision requires it. Keep findings, source dates, limitations and decisions in a shared record.

Separate evidence from interpretation

For each finding, record what was observed, what it might mean and what would change the conclusion. An interview reports one buyer's experience. CRM data covers recorded opportunities. A market estimate depends on definitions and assumptions. An audit records test outputs. None should silently stand in for the others.

Check important conclusions across sources. If buyers say implementation risk matters, look for it in lost-deal notes, committee requirements and competitors' proof. If sources disagree, investigate the difference rather than averaging it away. Personas should reflect evidence from people and decisions, with untested assumptions labelled as hypotheses.

Turn findings into a decision record

Record the target buyer, the issue, supporting evidence, proposed action, owner, cost and review date. Note what is deferred and why. Review fast-moving visibility findings more frequently than structural segmentation. Keep a fixed audit benchmark for comparisons, and a separate exploratory question set for new buyer situations. The output is a defensible marketing choice, with an explicit test of what happens next.

Part 3

ICP

The ideal customer profile, the buyer questions that come out of it, and what changes in marketing, sales and the 2027 plan once the audit is built on them.

- 9 The ideal customer company
- 10 The buying committee, buyer profiles and personas
- 11 From ICP to the question set
- 12 From ICP to marketing and sales
- 13 Your 2027 plan

The ideal customer company

Define which accounts you can win, serve profitably and retain, then test whether the market offers enough of them to support growth.

We've worked with brands that hadn't looked at their ideal customer profile in years. Often the profile was still roughly right about who the customer was, since the pain, the trigger and the people in the room were the same, but it said nothing about how the buyer now finds suppliers, and for buyers who ask an AI engine first, the engine's answer picks which suppliers to name and how to describe them.

The ICP supports account selection, positioning and investment priorities. Buyer questions are one practical output, alongside qualification criteria, decision barriers, evidence requirements and the sources marketing should reach. The audit tests one part of that picture.

What makes a customer ideal

Company size, industry and region are firmographics, useful for building lists and sometimes non-negotiable eligibility criteria, but they don't on their own say what makes a customer ideal. An ideal customer is one your company can win, serve profitably, keep, and find more of. Define the account characteristics associated with a strong fit and commercially attractive business, and include needs, capabilities, service requirements and economics alongside size and industry.

Start with the accounts you already have. Separate the customers who became profitable, repeatable business from the ones that were expensive to serve, set both next to the opportunities you lost, and look for what distinguishes the first group:

- **Need.** What problem were they trying to solve, and how urgent was it?
- **Fit.** How closely did what you offer match what they bought, and what did you have to adapt?
- **Buying trigger.** What happened inside the company that started the search? A new leader, a missed target, a failed supplier, a regulatory change, a growth plan.
- **Capabilities and service requirements.** What did the customer need to have in place to get value, and what did it take to deliver and keep them?
- **Economics.** Deal size, margin, time to close, retention, expansion.

Size and industry help locate the pattern. Test it against prospective accounts as well as existing customers before treating it as a growth opportunity. The result identifies accounts to pursue, accounts to disqualify and existing relationships to develop, with the commercial reasons recorded.

Is your ICP doing the job?

Can your current profile answer these without someone having to go and find out?

6 questions your ICP should answer

- 1 What event inside the customer's company starts the search?
- 2 Who in the buying committee can stop the deal, and what would make them do it?
- 3 Which three competitors does this customer compare you with, as opposed to the ones you benchmark against?
- 4 Which accounts does the profile rule out, and why?
- 5 What does this customer type into a search box or an AI engine on the day they decide to look?
- 6 When was the profile last checked against the last twelve months of wins and losses?

The unanswered questions are the research agenda for the next two chapters.

What the audit adds

Where an engine describes your customer focus differently from the intended target, check the public claims and the underlying customer evidence. An inaccurate AI answer does not establish that your ICP is wrong. The investigation may reveal an outdated profile, weak public evidence or an engine error.

The buying committee, buyer profiles and personas

Who initiates, influences, evaluates, uses, approves or blocks the purchase, what each of them needs, and whether AI is a source or a decision aid for each.

The buying committee

The person who reads your content is rarely the only person who decides. Identify who initiates, influences, evaluates, uses, approves or can block the purchase, and establish their authority and relationships rather than assuming the visible contact controls the decision. Edelman and LinkedIn's 2025 survey of 1,934 US business executives found that 55% of “hidden” decision-makers, the people in finance, operations, procurement and similar functions who influence a purchase without being its main contact, use thought leadership as part of how they vet vendors.[4]

Finance and procurement teams have always conducted their own supplier checks, and AI makes that research easier and faster, often without the sales team knowing what questions were asked. A finance director can ask an AI engine whether a supplier is credible, what it typically costs, who else does the same thing and what clients say, and get an answer during the meeting, and a fluent answer is not the same as a verified one, which is a point sales can make.

Buyer profiles

For each role on the committee, describe the responsibilities, needs and buying authority of the person you need to reach. Profiles identify the relevant participants inside a target account, a map of who holds which part of the decision rather than a persona, and the practical test is whether sales could use the profile to work out, for a new account, who to find and in what order.

Role	Needs to be sure of	What they ask AI
Problem owner	That the problem is worth acting on and someone can fix it	"Who helps companies like ours with this?"
Evaluator	That the approach will work in practice for a company like theirs	"How do these two firms compare for a mid-sized manufacturer?"
Budget owner	What it costs, what it returns, and the risk	"What does this kind of engagement cost, and is it worth it?"
Blocker	That the supplier is credible and won't embarrass them	"Is this supplier reliable? What do clients say?"
End user	That it will fit how they work day to day	"What does working with them actually involve?"

Buyer personas

Personas go a level deeper: the priorities, motivations, objections, decision criteria and trusted sources of each participant. For each one, establish whether AI is an information source, a decision aid, or neither. A problem owner might use AI to understand the category before they know what to ask for, an evaluator to screen a long list, a budget owner to check a number in a meeting, and a reviewer to look for risks and unmet requirements. The questions in the table are starting hypotheses, and the interviews establish what each role asked.

Personas built from interviews are worth the effort. Personas generated from your website and a few public sources risk describing the customer your material implies rather than the one who bought, so check every claim in a persona against something a real buyer said or did.

What the audit adds

The audit's validation questions, what the firm does, who it serves, what clients say, are the kind of questions committee members ask. In the sample, the one damaging error, an invented client reference, surfaced on a reputation question, and the expertise questions, which the sample firm scored well on, are closer to what an evaluator asks. Reading the audit by role as well as by category shows which kinds of question the engines answer well about you and which they don't, which the interviews can then match to the people asking them.

From ICP to the question set

Turn validated customer and buyer evidence into a relevant question set. The number of questions follows the scope, rather than defining it.

This is the step that connects Part 3 to the audit in Part 1. For the priority customer and each role on the committee, write the questions they would ask at three points: when they first recognize the problem, when they're comparing options, and when they're checking a supplier before committing. Write them the way the buyer would type them. Some buyers use category terms, such as “which firms offer operational excellence consulting,” and others describe the situation, such as “we’re a 400-person manufacturer and our margins are shrinking, who should we talk to,” so the set should include both.

Draft 15 to 20 questions, then select and extend to 25 for the baseline, balancing discovery, comparison, cost and validation. Broader scope may need more.

- Category questions with no supplier named, in the buyer's language
- Scenario questions that describe the trigger and the situation, one per priority segment
- Comparison questions that name the competitors the buyer would consider, including the ones the engines keep naming
- Cost questions, in the terms the budget owner uses
- Questions about your company specifically: what it does, who it serves, what clients say

Comparing the two question sets

Compare the research-led set with the audit's questions. Retain questions supported by buyer evidence, refine weak assumptions and add missing situations. Category terms can be relevant too. Record why each question belongs, which buyer and decision it represents, and whether it is part of the fixed benchmark or exploratory work.

For each priority question, identify the evidence a useful answer would need: the buying situation, relevant capabilities, delivery requirements and verifiable proof. Check where that evidence exists, on your own site and in credible third-party sources. Then test the engines. Missing evidence is an issue to investigate. It does not guarantee that publishing a page will secure a recommendation.

A living document

Maintain two question sets. Keep a fixed benchmark for trend comparisons, using the same protocol and repeated runs each wave. Develop an exploratory set as buyers, competitors and priorities change. Version every revision. When a benchmark must change, run old and new sets together where feasible and mark the break in the series. A new question mix is not evidence of improved performance.

ICP research is where we start with a client

Variant1 builds the profile from your own accounts, interviews and lost deals: the company characteristics, the committee, the personas and the question set, dated and ready to run as an audit. If yours hasn't been checked against the last year of wins and losses, that's the first conversation to have. variant1.com

From ICP to marketing and sales

What changes in ads, content, social, the website and PR, what sales takes from it, and the evidence that justifies each.

Research determines the market, customer and decision barrier worth addressing. The audit helps identify how selected engines represent them. Use the table as a set of hypotheses to investigate, not automatic spending instructions. Approve work when buyer relevance, credible evidence and a proportionate expected benefit support it.

Check	Issue to investigate	Evidence to assess
Appearance, Nomination	Absence from the directories, roundups, associations and trade coverage the engines cite for your category. No pages that describe, in the buyer's words, the situations your best customers were in.	Share of voice well below the firms you compete with, and source tracing showing those firms listed where you aren't.
Showdown	No comparison material that states, in specifics, how you differ from each firm the engines keep naming.	Head-to-head answers that describe the competitor in specifics and you in generalities.
Worth	No stated price range or account of what the buyer gets, where publishing one suits the pricing model.	Cost answers that give figures for competitors and none for you, or a borrowed figure.
Expertise	No current facts page (what you do, who you serve, where, how big, since when), inconsistent facts across profiles, or old pages still live.	Accuracy errors traced to pages or profiles you control.
Reputation	Reviews on platforms the engines cite. Named clients and substantiated results where you can. Earned coverage, which made up most of the links cited in one large sample. [6]	Credibility answers citing thin or invented evidence, or no third-party evidence at all.

Marketing

Advertising. The question set adds buyer language to the keyword and audience research the ad account already has. Scenario questions in the buyer's words are a starting point for ad copy, and the competitors the engines keep naming are candidates for comparison campaigns, subject to the usual checks on relevance, policy and budget. Presence in AI answers says nothing on its own about paid search budget in a segment, which stays with the channel's conversion and incrementality data.

Content and social. The named-or-absent grid is one input to the content plan alongside opportunity size, search behaviour and conversion data. Two candidates for early work are the questions you're absent from in the segments you want most, and the pages buyers land on from the questions where you're already named, since those answers get checked. Demonstrated expertise can do some of the work brand recognition would otherwise do, and in Edelman and LinkedIn's survey, 53% of decision-makers agreed that if thought leadership is high quality, brand recognition matters less.[4]

Website. The site is where the buyer lands to check what the engine said, and it's where third parties get the facts they repeat, so the facts have to be current and consistent: what you do, who you serve, where, how big, since when. Google's own guidance is that standard search practice still applies to its AI features.[7] Keep the description of the company consistent across the site, directory listings, partner case studies and what sales says, because inconsistent versions give the engines more than one to choose from.

PR and directories. The cited sources are a list of outreach candidates. Review the directories, roundups, review platforms and trade publications the engines cite for your category, prioritize them by audience fit, credibility, cost and eligibility, and make sure the facts page they would link to exists before pursuing listings that link.

Sales

Battlecards. The comparison answers are what a buyer who asked a similar question may have read before the first call. Give sales the engine's description of each competitor, the basis on which it distinguishes you, and the claims about competitors and about you that are wrong, so they can correct them if a buyer repeats them.

Discovery questions. The scenario questions from the ICP, validated in the interviews, are what buyers in the priority segment were asking before they reached you, and sales can open first calls with those questions.

What to correct first. Rank errors by likely buyer harm, exposure and ease of correction. An invented client reference can be serious, but the right order depends on the specific mistake and the buyers affected.

Agreed measures. Track qualified opportunities, conversion and acquisition cost alongside the fixed visibility benchmark. Use buyer interviews to investigate changes. The aim is to explain what marketing learned and changed, without attributing sales movements to AI work on timing alone.

Your 2027 plan

What the audit, the research and the ICP change in each section of the plan, in the order to do it, and the briefing your CEO will ask for.

Start with the decision and the evidence already available. An initial audit can expose questions for investigation, while research establishes the customers, buying criteria and opportunities that deserve investment. Scope the work to the team's capacity. The plan should state what is known, what is being tested and what will wait.

Plan section	What changes	The evidence it rests on
Targeting	Priority segments ranked, and the accounts ruled out.	Market potential, buyer needs, account economics, wins and losses.
Messaging	Consistent, verifiable facts about the company, and comparison claims in specifics.	Showdown and Expertise answers, and the personas.
Content and PR	Scenario pages for the questions you're absent from, and listings in the sources the engines cite where eligible.	The named-or-absent grid and the cited sources.
Channel budget	Comparison campaigns considered against the firms the engines name, decided on channel data.	Share of voice by segment, the two competitor sets, and the channel's own results.
Sales enablement	Battlecards, discovery questions, the errors to correct first.	Comparison answers, and accuracy errors ranked by damage.
Measures	Commercial measures reported above visibility measures, with date, engines and denominators.	The baseline run and the schedule.

The order of work

Before planning: define the commercial question. Review existing market and customer evidence and select a priority segment. An initial audit can reveal how engines describe that market. The 75-answer snapshot maps breadth. Repeat consequential prompts before making costly decisions, and record the test conditions.

During planning: close the evidence gaps. Test segment attractiveness externally, compare account economics, interview relevant buyers and examine alternatives. Refine the ICP, decision criteria and question set. Research may change the target or position, not merely the audit prompts.

In the plan: the decisions. Targeting, messaging, content priorities, channel budget, PR and directory presence, sales enablement, and the measures, each with the audit or research finding that supports it, so the plan can be defended when someone asks why.

Through the year: test and review. Repeat the fixed benchmark consistently. Keep exploratory questions separate. Refresh buyer and market evidence when decisions require it.

When your CEO asks what you're doing about AI

Have a one-page assessment ready, with numbers from your own run and their limits stated. It covers what's changing for your buyers, what you're already doing, where the gaps remain and who's being named alongside you, what you recommend next with owners and costs, and how you'll measure progress.

The one-page briefing

What is changing

What the survey evidence and your own interviews say about buyers using AI, with one number from your own run.

What we are already doing

The ICP review, the journey map, the baseline audit, the fixes under way, and what we have learned.

Where the gaps remain

Where we are absent or misdescribed, which competitors are named alongside us and what sources are cited, and the questions still open.

What we recommend next

Priorities in order, each with an owner, a cost and what it is expected to produce.

How we will measure progress

The commercial outcomes leadership should track, and the indicators that move first.

Commercial measures: qualified opportunities, conversion, acquisition cost, where buyers first heard of us.
Visibility measures: share of voice, citation rate, accuracy rate, with date and engines.

Use the briefing to explain the decision, not just the scores: which buyers matter, what evidence changed the plan, what has already been done and what remains uncertain. Report the fixed visibility benchmark beside commercial measures. Investigate divergence before claiming impact, and show what would cause you to change course.

Research and credible customer proof take time. Start now so marketing can explain its priorities and evidence at the next budget discussion, before leadership asks.

Run the first cycle with Variant1

We run the audit, the research and the ICP work with your team, produce the briefing and the plan inputs, and leave you with a schedule you can keep. Book at variant1.com.

What a sample AI visibility audit looks like

Fenwick Hale Advisory, a 150-person management consulting firm

Four pages from a sample report. The firm, its competitors, the answers and the scores are invented to show the shape of a real report.

This fictional sample tests 25 questions once on each of three engines: 75 answers. Nine category questions and two scenarios produce 33 discovery answers. Four comparison, three cost, four expertise and three reputation questions produce the other 42.

Category scores are the percentage of answers meeting each stated test. The rubric and index calculation are disclosed on page 8. The index is illustrative, not an industry benchmark, a probability of selection or a forecast of sales. The next four pages show the summary, one answer, the grid and one category page.

42

Illustrative index,
out of 100.

3%

Share of voice: 3 of 100
supplier mentions across
the 33 discovery answers

17%

Citation rate: the firm's site
cited in 13 of 75 answers

93%

Accuracy: 41 of 44 claims
about the firm correct

Executive summary

42


out of 100
Sample index

The engines describe Fenwick Hale Advisory accurately once it is named. They rarely name it unprompted.

Most of the lost points are in discovery. Across the 33 answers to open and scenario questions, the firm was named three times, all on the same engine, while Corvane Partners, Tallis Reed Consulting and Hollis Kerr Group were named on all three engines.

Where the firm is named, the engines are accurate: 41 of the 44 claims they make about it are correct, and in head-to-head comparisons it comes across as a credible mid-market choice. The one damaging error is an invented client reference on one engine, which cites a third-party page that repeats it.

Category scores



Top three priorities

1. Correct the invented client reference

Request correction from the cited publisher and verify the firm's own records.

2. Assess relevant third-party sources

Check buyer fit, eligibility and evidence before committing budget.

3. Decide whether pricing disclosure would help

Use buyer needs and the commercial model, not a missing number alone.

The answer buyers get


One of the 33 discovery answers in the sample. The firm is absent and three competitors are named with details.

Engine B • Nomination question

“We’re a 400-person manufacturer with shrinking margins. Which consulting firms should we talk to about an operations review?”

For a mid-sized manufacturer facing margin pressure, a few firms come up repeatedly. **Corvane Partners** focuses on operational turnarounds for manufacturers in roughly the 200 to 1,000 employee range and has published case studies with documented margin improvements. **Tallis Reed Consulting** is known for lean and supply chain work and is recommended by at least one industry association. **Hollis Kerr Group** offers fixed-fee operations reviews, typically in the \$60,000 to \$120,000 range according to its website, which may suit a company that wants a defined scope. I couldn’t find reliable pricing for the other two. It may also be worth asking peers in your industry association, since several of these firms are active there.

Firms named across the 33 discovery answers



The six firms above account for 96 mentions. Other suppliers account for 4, bringing the total to 100. Fenwick Hale's share of voice is 3%. It appears only on one engine, in region-specific questions. These are invented observations from this snapshot, not stable buyer exposure.

Where Fenwick Hale shows up


Six of the 25 questions. Each cell is one question on one engine.

	Engine A	Engine B	Engine C
Which consulting firms help mid-sized companies improve operations?	Absent	Absent	Absent
We're a 400-person manufacturer with shrinking margins. Who should we hire?	Absent	Absent	Absent
Fenwick Hale vs Corvane Partners for an operations review?	Named	Named	Named
How much does a mid-market operations engagement cost?	Absent	Absent	Absent
What services does Fenwick Hale Advisory offer?	Named	Named	Named
What do clients say about Fenwick Hale Advisory?	Named	Absent	Named
Is Fenwick Hale a good fit for a family-owned manufacturing?	Absent	Absent	Absent
Which firms are best for a mid-market operations review in Chicago?	Absent	Absent	Named
Which firms are best for a mid-market operations review in Chicago?	Absent	Named	Absent
Who should a 300-person distributor hire to cut freight costs?	Absent	Named	Named

The full report lists all 25 questions.

The firm is named in most answers to questions that include its name, and in 3 of the 33 discovery answers.

Where the engines' sources come from



The 75 answers produced 214 citations: 171 to third-party pages, 29 to the firm's own site across 13 answers (the 17% citation rate), and 14 to competitor pages.

Category page: Appearance

One of the six category pages in the full report.

7

out of 100
25% index weight

Appearance is the percentage of open category answers naming the firm. Two mentions in 27 answers produce 7%, rounded. The sample gives this check a 25% weight to emphasize discovery. That is an illustrative choice, not a finding about every buyer journey.

Across the 27 answers to the nine open category questions, the firm was named twice, both on the same engine and both in answers to questions that mentioned its region. On the other two engines it did not appear. The engines named Corvane Partners in 24 of the 27 answers, citing a trade roundup, an association directory and a review platform the firm is absent from.

Sample questions: “Which consulting firms help mid-sized manufacturers improve operations?” “Best operations consultants for a 300-person manufacturer in the Pacific Northwest” “Who does supply chain consulting for mid-market companies?” “Top management consulting firms for family-owned manufacturers”

Nomination	Showdown	Worth	Expertise	Reputation
17%	92%	0%	92%	89%
Sample pass rate	Sample pass rate	Sample pass rate	Sample pass rate	Sample pass rate

The other five category pages are in the full report.

Top recommendation

Assess the three sources the engines cite for this category

The firm is missing from three sources cited in the sample. Check eligibility, buyer relevance and the accuracy of existing coverage before pursuing inclusion. The research may identify a more valuable specialist source. A listing is an action to evaluate, not a guaranteed route into AI answers.

Affects: Appearance, Nomination. Effort and eligibility need checking. Repeat the relevant prompts before and after the work.

Get your audit

All six categories, every citation recorded, accuracy errors traced where a source exists, and a prioritized plan. Built from your buyer questions. Book at variant1.com.

Sources

1. Semrush. How AI tools shape the B2B buying process. July 8, 2026. US survey, fieldwork March-April 2026. Findings for 519 AI users.
2. U.S. Small Business Administration. Market research and competitive analysis. Guidance on market demand, direct research and competitors.
3. Statistics Canada. Canadian business counts, December 2025. Released February 13, 2026. Counts are business locations.
4. Edelman and LinkedIn. 2025 B2B Thought Leadership Impact Report. Survey of 1,934 US business decision-makers.
5. Ahrefs. Top Brand Visibility Factors in ChatGPT, AI Mode, and AI Overviews (75k Brands Studied). December 12, 2025. Correlations, not causes.
6. Muck Rack. What Is AI Reading? May 7, 2026 findings. More than 25 million cited links across 17 industries. Citation shares are not influence shares.
7. Google Search Central. AI features and your website. Official guidance for AI Overviews and AI Mode.
8. U.S. Census Bureau. County Business Patterns. Establishment counts, employment and payroll by industry, size and geography.
9. Rand Fishkin and Patrick O'Donnell. AIs are highly inconsistent when recommending brands or products. SparkToro, January 27, 2026. Vendor-linked research, 2,961 runs.

Sources 1-9 link to their publications. The worked example, sample firms and audit results are fictional. The eight-area research framework and planning recommendations are Variant1's approach.

variant1

Find out what AI says about you.

Variant1 runs AI visibility audits and builds market intelligence systems for B2B companies in Canada and the US.

kevin@variant1.com

mika@variant1.com

www.variant1.com

Vancouver | Toronto