EXECUTIVE BRIEFING

KOHL DIGITAL

AI Visibility for B2B Companies

How to Become the Answer

What AI visibility is and how to manage it: the four jobs, a monthly measurement procedure your team can run by hand, a metrics table and how it relates to SEO.

Mike Kohl
Founder, Kohl Digital

mikekohl.com/guides/ai-visibility-for-b2b-companies

KEY TAKEAWAYS

- AI visibility is how often, how accurately and in what position assistants name you; manage it with four jobs: measure, fix, earn, track.
- Measure by hand: ask 15 to 20 buyer questions at least three times each in fresh sessions, and record naming, position, accuracy and cited sources.
- Fix wrong and inconsistent facts first, then deepen pages for three capabilities, then pursue the most-cited outside sources that omit you.

AI visibility is how often, how accurately and in what position AI assistants mention your company when buyers ask about your category. You manage it with four jobs: measure what AI says, fix the facts on your own site and profiles, earn a place in the sources AI cites, and track the result every month.

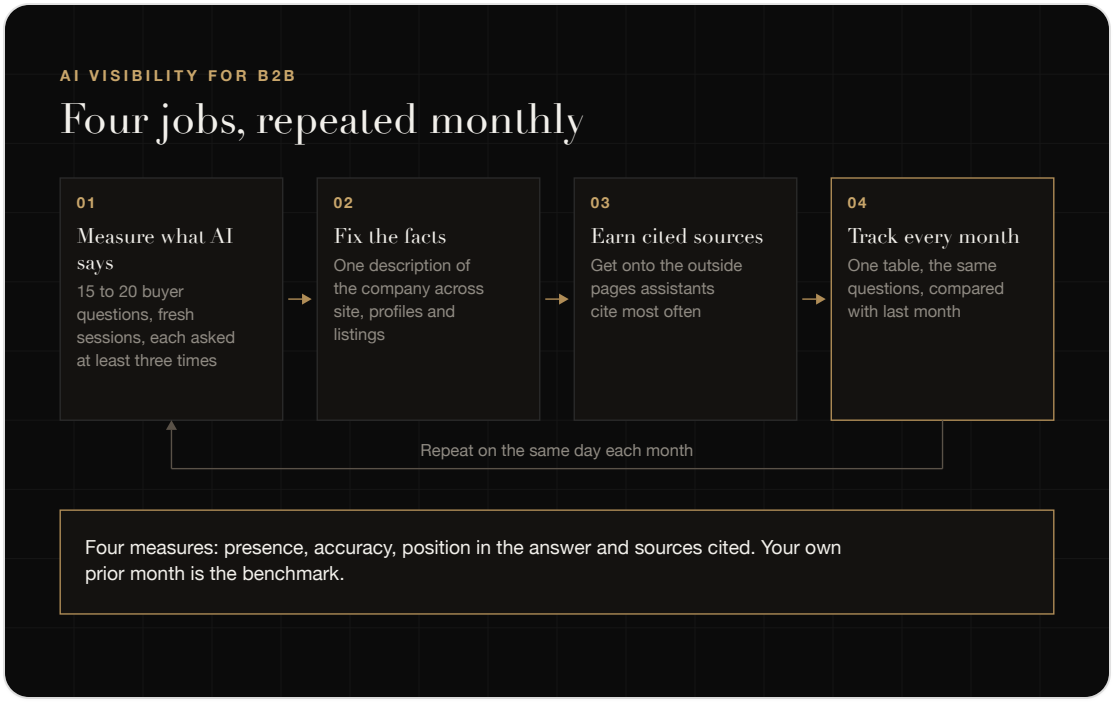

Run the same four jobs every month and compare against the month before.

This guide is for the CMO or head of marketing who owns pipeline quality and brand position. The number you own is the share of qualified buyers who

arrive already knowing who you are and what you do. AI answers now shape part of that, and they are measurable.

Why this matters now

Buyers ask AI assistants for vendor shortlists before they talk to sales. A buyer might type "who are the leading providers of [your category] for a company our size" or "how should a CFO evaluate [your category]," and read a short list of named companies with a sentence about each.

If you are on that list and described correctly, you start the conversation ahead. If you are missing, or described wrongly, you may never learn the deal existed. Two companion essays cover the evidence. This guide is the working manual, with the steps your team can follow.

- Read the essay on how AI now shapes the vendor shortlist
- Read the essay on the shift from clicks to answers

How assistants form an answer, in plain terms

Without claiming anything about the internals of a particular product, the general pattern is this. When an assistant answers a question about vendors, it draws on what it learned in training and, often, on sources it retrieves at the moment of the question. Many assistants show those sources as citations.

Three consequences follow, and they hold up whichever assistant a buyer uses.

- The assistant reads your public pages, your profiles and what other sites say about you. Inconsistency across them produces a muddled description.
- Sources that the assistant trusts and cites carry more weight than your own claims. Third-party pages that describe you shape the answer.
- The answer varies between runs and between askers. One screenshot proves little, so you measure patterns over repeated runs.

The four jobs

1. Measure what AI says

Find out how you appear today, across the assistants your buyers use, for the questions your buyers ask. The procedure is below. Without it, every later decision is a guess.

2. Fix the facts on your own site and profiles

Assistants quote what is easiest to read and most consistent. Check that your company name, descriptions, services, leaders and locations match across your website, LinkedIn, directory listings and any press or bio pages.

Then look at depth. One page that mentions twelve capabilities gives an assistant nothing to quote. A page per capability that states the problem, your method, the evidence and what an engagement involves gives it something to cite. Write direct, complete answers to the questions buyers actually ask.

3. Earn a place in the sources AI cites

When you run your measurement, record which sources the assistants cite. Those sources are the map. They might include industry publications, analyst-style roundups, directories, review sites, trade association pages and community discussions.

Your work is to earn a place on the pages that matter: a contributed article, an accurate profile, a data point others reference, a case study a publication will cite. Prioritize by how often a source appears across your question set.

4. Track every month

Run the same measurement every month and record the results in one table. Compare against the month before. Note what you changed, so you can connect changes to results.

A measurement procedure your team can run by hand

You can start this week with no special tools.

1. Write 15 to 20 questions a buyer would ask, in the buyer's words. Cover category discovery ("who leads"), evaluation ("how should I choose"), comparison ("X versus Y") and your own name ("what is [company] known for").
2. Choose the assistants. Use the two or three your buyers are most likely to use.

3. Open each assistant in a fresh session with no personal history, such as a private window or a new account.
4. Ask each question at least three times, because answers vary between runs.
5. Record each result in a spreadsheet: question, assistant, run number, whether your company was named, position in the answer, whether the description was accurate, and every source cited.
6. Ask a follow-up on each answer: "Where did that come from?" Record the response.
7. Summarize: in what share of runs were you named, where did you appear, what was wrong, which sources came up most.
8. Choose the three biggest gaps and assign owners.

Repeat on the same day each month, with the same questions. You may add new questions, and keep the original set fixed so the trend is comparable.

The metrics table

Four measures cover most of what you need.

Metric	What it tells you	How to calculate it
Presence	How often you are named	Runs where you appear, divided by total runs
Accuracy	Whether the description is right	Runs with a correct description, divided by runs where you appear
Position in the answer	Whether you lead or trail	Average position when named, or share of runs in the first three
Sources cited	Where the answer comes from	Count of citations by source, ranked, with your own pages marked

Add two supporting notes: sentiment of the description (positive, neutral, negative) and the competitors named alongside you. Your results are your own, so there is no outside benchmark to compare against. Your own prior month is the benchmark.

Reading the results: an illustrative example

The numbers below are hypothetical and exist only to show the arithmetic. Suppose your set has 18 questions, you ask each three times in two assistants, and that gives 108 runs. If your company is named in 27 of them, presence is 25 percent. If 18 of those 27 describe you correctly, accuracy is two thirds.

Now look at the pattern behind the totals. Perhaps you appear often on the "what is [company] known for" questions and rarely on category discovery questions. That points to a different fix than the reverse pattern. In the first case buyers who already know your name get the right picture, and the work is

to enter the category conversation. In the second case you are a known name with a thin description, and the work is on your own pages.

Then read the citations. If the same three outside sources appear in most of the discovery answers and none of them mentions you, those three pages are your first targets.

Set the cadence and the owners

A monthly program needs a rhythm that fits inside normal work.

- **Week 1 of the month:** run the measurement and fill in the table.
- **Week 2:** write the one-page summary, with the metrics, the biggest gaps and the sources cited.
- **Weeks 3 and 4:** do the work on the top three gaps.

Assign roles. One person owns the monthly report. The web owner fixes facts on your site and structured data. Whoever handles PR, partnerships or analyst relations owns the outside sources. The CMO reviews the table each month and decides where the next hour of effort goes.

What to fix first

Order the work by effort and effect.

1. **Wrong facts.** Fix every inaccurate description at the source you control: your site, your profiles, your listings. This is quick and removes the most damaging answers.

2. **Inconsistent descriptions.** Pick one approved description of the company and each offering, and use it everywhere.
3. **Thin pages.** Build or deepen the pages for the three capabilities where you most want to be the clear answer.
4. **Outside sources.** Pursue the top-cited outside pages that omit you, one at a time.
5. **New question sets.** Once the base is stable, add questions for adjacent topics and for specific buyer roles.

How AI visibility relates to SEO

The two overlap and each has its own job. Classic SEO aims to win a click from a results page. AI visibility aims to be named inside an answer, often without a click.

Much of the underlying work carries over. Clear pages, consistent facts, strong pages for each specialty and credible third-party mentions help both. What differs is measurement. Rankings and traffic tell you about SEO. Presence, accuracy, position and citations tell you about AI visibility. Run both, and put the AI measures next to your search data in the same monthly review.

Part of a larger program

AI visibility is one part of an AI program. The other part is internal: systems that change how work flows inside the company, such as support, renewals and quoting. The blueprint for the whole program has six moves, and its last move is "be the answer."

The discipline carries between the two. Facts that are consistent, sourced and owned make your public answers accurate and your internal systems reliable. The controls in engineering AI for production apply to anything you build, including the content systems that feed your public facts.

For the buying-committee version of this measurement, where different buyers get different answers, read the sibling guide.

- Read what AI tells each buyer about your company

What goes wrong

- **One screenshot as the evidence.** A single answer looks like proof and is a sample of one. Repeat runs.
- **Inconsistent facts.** The site says one thing, LinkedIn says another, and an old directory listing says a third. Assistants repeat the mess.
- **Content written for volume.** Fifty thin posts give an assistant nothing quotable. A smaller body of deep, specific pages does more.
- **Ignoring third-party sources.** The team improves its own site and never looks at which outside pages the assistants cite.
- **Changing the question set every month.** The trend becomes impossible to read. Keep a fixed core set.
- **No owner.** Marketing, web and PR each assume another team handles it. Name one person for the monthly report.
- **Reporting only presence.** Being named with the wrong description hurts. Track accuracy next to presence.
- **Treating it as a one-time project.** Answers change as sources change. The monthly rhythm is the program.

Questions leaders ask

Do we need special software to start?

No. A spreadsheet, a fixed question list and a few hours each month give you a working baseline. Tools can automate collection later, and the procedure above is what they would automate.

How is this different from what our SEO agency does?

Your agency likely tracks rankings and traffic. Ask them whether they measure how assistants describe you, which sources get cited and how that changes by buyer. If not, the work sits beside SEO and needs an owner.

How long until we see changes?

It depends on which sources change and how quickly assistants pick them up. Fixing facts on your own site and profiles is within your control, and third-party placements take longer. The monthly table shows movement as it happens.

Can we control what AI says about us?

You influence it through three levers: accurate, consistent facts, deep pages that answer buyer questions and credible third-party sources that cite you. Nobody dictates the answer, so measure it every month.

The Leverage Diagnostic takes about six minutes, scores the six places leverage hides and prices the gap with the math shown, and it includes how your company shows up to buyers. Start with the Leverage Diagnostic.

NEXT STEP

Find where leverage hides in your company.

The Leverage Diagnostic takes about six minutes, scores the six places leverage hides and prices the gap with the math shown.

mikekohl.com/diagnostic

Kohl Digital builds AI systems for mid-market and enterprise companies. Founder Mike Kohl architected the enterprise sales platform at Level 3 Communications and Lumen Technologies that went from $0 to $500 million in four years.

Book a call: mikekohl.com/contact