[3-phase operating plan]

How to build an answer engine optimization program

A sequenced operating plan for B2B communications and marketing teams starting from zero: what to do first, what depends on what, and how to know whether any of it is working.

days 1–30
Baseline Audit

days 31–60
Content and Entity Fixes

days 61–90
Third-Party Validation and Measurement

Actual Agency · Computational Communications

November 2026

[the case for starting now]

The buyer's shortlist is being assembled by a machine

51%

of B2B software buyers now start research in an AI chatbot rather than a search engine, up from 29% eleven months earlier

G2 · 2026 Answer Economy survey

69%

of those buyers chose a different vendor than they originally planned, based on what the AI told them

G2 · 2026 Answer Economy survey

20%

of marketing teams have begun any form of AEO implementation, while 70% expect it to significantly affect their strategy within three years

Acquia & Researchscape · 500+ marketers

Most B2B communications teams are not losing deals because their product is inferior. They are losing them because an AI answered the category question before the buyer ever reached the company's website.

According to G2's 2026 Answer Economy survey, 51% of B2B software buyers now start research in an AI chatbot rather than a search engine, up from 29% just eleven months earlier. More striking: 69% of those buyers chose a different vendor than they originally planned, based on what the AI told them. The buyer's shortlist is being assembled by a machine, and most companies have no idea what that machine is saying about them.

The gap is wide. A 2026 survey found 96% of B2B companies are effectively invisible in AI discovery. Only 20% of marketing teams have begun any form of AEO implementation, despite 70% believing it will significantly affect their strategy within three years, according to an Acquia and Researchscape study of 500+ marketing professionals.

That gap is also an opportunity, at least temporarily. The teams that build a systematic AEO program now will establish a citation footprint before the space gets crowded. The teams that wait will spend considerably more effort catching up to companies that moved earlier.

This guide lays out a practical 3-phase program for a B2B communications or marketing team starting from zero. Not a vendor pitch. Not a list of tactics. A sequenced operating plan, organized by what to do first, what depends on what, and how to know whether any of it is working.

What AEO actually requires

Before the calendar, a quick framing. Answer engine optimization is not a content hack. It is not about stuffing FAQ sections with keywords or restructuring every blog post as a Q&A. Those things may help at the margins.

What AEO actually requires is that your company's narrative be legible, credible, and corroborated across the sources AI systems use to build their answers. Research from Princeton, Georgia Tech, and IIT Delhi (published at ACM KDD 2024) tested nine optimization methods across 10,000 queries and found that targeted techniques can improve a source's visibility in AI-generated answers by up to 40%. The strongest levers were adding statistics and adding quotations from credible sources, not structural formatting tricks.

That is a useful anchor for the work ahead: evidence-rich, externally corroborated content outperforms cleverly formatted thin content. The 3-phase plan below is built around that principle.

[phase 01 · days 1–30]

Baseline Audit

You cannot improve what you have not measured. The first month is entirely diagnostic. Resist the temptation to start producing content before you understand what the engines are saying about you today.

Run a prompt audit across the major engines

Start by systematically querying ChatGPT, Perplexity, Google AI Overviews, and Gemini with the questions your buyers actually ask. These fall into three categories:

- **Category questions**: "What are the leading [category] platforms for [use case]?" or "Which [category] vendors do analysts recommend?"
- **Comparison questions**: "How does [your company] compare to [competitor]?" or "What are the alternatives to [competitor]?"
- **Reputation questions**: "What do customers say about [your company]?" or "What is [your company] known for?"

Document every response verbatim. Note which sources the engines cite. Note whether your company appears, and if so, how it is described. Note what your competitors are being credited with that you are not.

This is the baseline. Everything in months two and three is an attempt to move it.

Map your citation sources

The engines are not generating opinions from nothing. They are synthesizing from sources. Meltwater's analysis of 9.5 million AI citations found that LinkedIn, Reddit, and YouTube together accounted for 47.5% of them, while company websites accounted for only 18.7%. That ratio matters for where you invest.

During the audit, identify which third-party sources are being cited when your category comes up. Industry publications, analyst reports, review platforms, LinkedIn profiles, Reddit threads, and news articles all feed the answer layer. Understanding which sources carry weight in your specific category is more useful than any generic best-practices list.

Audit your entity consistency

AI systems build their understanding of a company from signals across many surfaces. Inconsistent information creates ambiguity that models resolve conservatively, which usually means they either exclude you or describe you inaccurately.

Check these surfaces for consistency in how your company is named, described, and categorized:

surface	what to check
Your own website	Company description, category language, product names
LinkedIn company page	Tagline, about section, industry classification
Google Business Profile	Business category, description, website link
Wikipedia / Wikidata	Existence, accuracy, citations
Industry directories	Consistent naming and category
Press coverage	How journalists describe your category and positioning

Inconsistencies here are not cosmetic problems. They are signals that models use to decide how confidently to include you in an answer. Fix them before you build anything new on top of them.

Establish your measurement baseline

Before the month ends, record the following numbers. You will compare against them at day 60 and day 90.

- How many times does your company appear across your prompt audit responses?
- In how many category-level answers are you named?
- What sources are cited when you appear, and how do they compare to competitors?
- What is the sentiment and accuracy of descriptions?

You are not competing against an abstract standard. You are competing against specific companies for the same buyer attention in the same AI-generated answers. That comparison is the most operationally useful number in your baseline.

[phase 02 · days 31–60]

Content and Entity Fixes

The audit told you where you are invisible, inaccurate, or inconsistently described. Month two is about fixing those problems at the source. This is not a content sprint. It is a structural repair job.

Fix the entity layer first

Before producing a single new article, correct the inconsistencies you found in month one. This means:

- Aligning your company description, category language, and product naming across your website, LinkedIn, Google Business Profile, and any industry directories where you appear
- Ensuring your Wikipedia entry exists and is accurate, with citations to independent sources
- Updating stale press coverage where possible by issuing corrections or publishing new material that supersedes outdated descriptions

The reason to do this first is simple. New content you produce will be crawled and indexed against your existing entity signals. If those signals are inconsistent, the new content will inherit the ambiguity rather than resolve it.

Rebuild your owned content for AI legibility

AI systems favor content that front-loads a direct, self-contained answer. A page that opens with three paragraphs of company history before getting to the point is less likely to be cited than a page that opens with a clear, specific claim followed by evidence.

Audit your highest-traffic pages and your most important category pages. For each one, ask:

- Does the opening paragraph answer the question a buyer would ask about this topic?
- Does the page include specific statistics, with sources?
- Are there named experts or quotations from credible third parties?
- Is the content updated within the last twelve months?

That last point matters more than most teams realize. Ahrefs' analysis of 17 million AI citations found that AI-cited content is 25.7% fresher on average than non-cited content. AirOps' 2026 State of AI Search Report found that 83% of AI citations for commercial queries came from pages updated within the past twelve months.

The practical implication: a well-structured page from two years ago is likely being passed over in favor of a thinner page updated more recently. Freshness is a real signal. Build a refresh cadence into your content operations, not as a one-time project.

Develop a set of category-level answer pages

These are pages designed specifically to answer the questions your buyers ask AI systems. Not product pages. Not blog posts that bury the answer. Pages that open with a direct answer, cite evidence, and establish your company's position in the category.

Think about the questions from your prompt audit. For each high-value category question where you were absent or poorly represented, there should be a page on your site that answers it credibly. These pages become the owned-content foundation that earns citations over time.

Aim for five to ten pages in month two. Quality matters more than volume here. A single well-evidenced answer page will outperform ten thin ones.

Strengthen your LinkedIn presence

LinkedIn is the second most-cited source for B2B category queries across AI engines, according to Meltwater's analysis. About 75% of LinkedIn's AI citations come from individual member profiles, not company pages. That means your executives need active, substantive presences, not just profile completions. Publish thought leadership that addresses category-level questions, engage in industry conversations, and ensure executive bios reflect current expertise. This is a citation surface, not a vanity exercise.

[phase 03 · days 61–90]

Third-Party Validation and Measurement

Owned content is necessary but not sufficient. The data on this is consistent across multiple large-scale studies: AI systems favor third-party sources over brand-owned pages by a wide margin.

Muck Rack's analysis of over one million AI citations found that 82% came from earned media rather than brand-owned pages. A pilot study by Stacker and Scrunch found that distributing content through third-party news outlets lifted AI citation rates from roughly 8% (brand-owned content) to 34%, a 325% increase across five LLMs. You can build the best answer pages on your website and still be invisible if no independent sources are corroborating your claims.

Month three is about building that external corroboration.

Prioritize earned media with citation value

Not all press coverage carries equal weight with AI systems. A mention in a trade publication that AI models frequently cite is worth more than a dozen mentions in outlets that do not appear in the citation data you mapped in month one.

Go back to your citation source map. Identify the five to ten publications and platforms that appear most frequently when your category comes up in AI answers. Those are your earned media targets for the quarter. A focused effort to place substantive, quotable content in those specific outlets will do more for your AI visibility than a broad spray of press releases.

Practical formats that earn citations:

- **Data-driven contributed articles** in industry publications, with specific statistics attributed to your team's research
- **Executive interviews** that produce quotable, expert-attributed statements on category questions
- **Analyst briefings** that result in reports or commentary that cite your company's position
- **Customer case studies** published on third-party platforms or in trade publications

The common thread is specificity and attribution. AI systems are more likely to cite a source that contains a named expert making a specific, verifiable claim than a source that contains general observations.

Build a structured review presence

For B2B companies, review platforms are a meaningful citation source. They provide third-party, customer-attributed assessments of your product that AI systems treat as independent corroboration.

If you do not have a structured approach to gathering and publishing customer reviews on platforms relevant to your category, build one now. This is not about gaming review scores. It is about ensuring that the customer evidence AI systems draw on reflects your actual customer experience, not a thin or outdated sample.

Run your second prompt audit

At day 75 or so, repeat the prompt audit from month one using the same questions and the same engines. Compare the results against your baseline.

What you are looking for:

- Has your company appeared in category answers where it was previously absent?
- Are the sources being cited now including the new content and earned media you produced?
- Has the accuracy or completeness of descriptions improved?
- Has your citation rate relative to competitors shifted?

Do not expect dramatic changes. AI models update their knowledge at different cadences, and some changes will take longer to surface than others. What you are looking for is directional movement, not a complete transformation.

A note on timelines: ChatGPT's training data has a knowledge cutoff, but its web-browsing capability means recently published, well-indexed content can appear in answers relatively quickly. Perplexity updates more frequently from live web results. Google AI Overviews draw heavily from indexed content. The implication is that your owned content changes may show up in Perplexity faster than in ChatGPT, and your earned media placements may surface in Google AI Overviews before they affect other engines.

Establish your ongoing measurement cadence

By day 90, you should have a repeatable measurement system, not just a one-time audit. This means:

- A defined set of prompts that you run monthly across the major engines
- A tracking document that records appearances, sources cited, and description quality over time
- A process for flagging inaccurate or outdated descriptions when they appear
- A quarterly review of which citation sources are carrying the most weight in your category

The goal is not to achieve a perfect score on any single metric. The goal is to build a system that tells you whether your narrative is moving in the right direction across the surfaces where your buyers are forming their initial views of your market.

---

*That is what an AEO program actually is: **not a campaign**, but an ongoing intelligence and communications discipline.*

The 3-phase plan gets you to the starting line. The work after it is what builds the compounding advantage.

[case study · redacted]

An Established E-Commerce Brand in Custom Publishing

1.9×

owned-site citation rate in AI answers, Q2 to Q3 (0.55 to 1.02 per chat)

3×

video-platform citation rate, Q2 to Q3 (0.46 to 1.46 per chat)

+4.2 pts

ChatGPT visibility across the last two measurement windows (44.1% to 48.3%)

Client, competitor and publication names are withheld.

The brand. An established e-commerce brand in custom publishing. It already led its category in AI answers, but its lead over the closest competitor had narrowed from 8.5 points to 1.3 points in a single quarter, and the brand had little visibility into where it was strong or exposed.

what the program did

The work followed the same sequence as this guide.

Baseline.

The tracked prompt set was widened from 147 to more than 250 buyer questions, which exposed weak topics the original set had hidden: holiday gifting, cards and weddings. A citation source map showed that earned press, not the brand's own site, was the largest source of AI citations.

Content and entity fixes.

A monthly cornerstone guide, supporting articles rewritten from older posts, and batches of FAQs with structured data. Every page passed an editorial check before schema was generated, so a page whose FAQ answer failed to render was held back until it was fixed.

Third-party validation.

Pitching was aimed at a monthly list of the 50 publications AI engines weight most, alongside review-creator outreach, an executive thought-leadership feature, and steady community and professional-network content.

Measurement.

Prompt testing every two weeks and a quarterly report against the baseline.

what changed

measure	before	after	change	period
Owned-site citation rate (per chat)	0.55	1.02	▲ +0.47	Q2 to Q3 2026
Video-platform citation rate (per chat)	0.46	1.46	▲ +1.00	Q2 to Q3 2026
Community-forum citation rate (per chat)	0.72	0.92	▲ +0.20	Q2 to Q3 2026
Average position in AI answers (lower is better)	2.2	2.0	▼ −0.2	Q2 to Q3 2026
Win rate	16.9%	19.7%	▲ +2.8 pts	Q2 to Q3 2026
Visibility across engines	40.9%	43.1%	▲ +2.2 pts	Last two windows
ChatGPT visibility	44.1%	48.3%	▲ +4.2 pts	Last two windows
Closest competitor's ChatGPT visibility	56.3%	50.7%	▼ −5.6 pts	Last two windows

Changes are calculated from the rounded figures shown. For average position, a falling number is an improvement.

what stood out

Earned media carried the most weight.

The top 20 editorial pages citing the brand drew more than 47,200 citations, and pages from a single consumer publication out-cited the brand's entire website.

Depth beat volume.

One cornerstone guide earned 3.2 times the citations of a typical supporting page (1,614 against 498).

Targeted content moved specific topics fastest.

The top-gaining topic rose 20.6 points in Q3. In the latest cycle one prompt rose from 29% to 64% and another from 32% to 46%.

Source: third-party AI-visibility tracking across ChatGPT, Google AI Overviews and Perplexity, comparing Q2 2026 (April to June) with Q3 2026 (July to September) and the two most recent two-week windows to November 4, 2026. These results coincided with the program; they do not prove it caused them.

[the plan on one page]

The 3-Phase Plan at a Glance

The sequencing matters. Fixing entity inconsistencies before producing new content prevents the new content from inheriting old ambiguities. Building owned content before pursuing earned media gives journalists and analysts something credible to reference. Running the second audit before the program ends tells you whether the investment is working and where to focus next.

01 / weeks 1-4

Baseline Audit

Understand current AI representation

key outputs

- Prompt audit
- Citation source map
- Entity consistency review
- Measurement baseline

02 / weeks 5-8

Content and Entity Fixes

Repair inconsistencies, build owned content

key outputs

- Entity alignment
- Refreshed pages
- 5-10 answer pages
- LinkedIn activity

03 / weeks 9-12

Third-Party Validation

Earn external corroboration

key outputs

- Earned media placements
- Review presence
- Second prompt audit
- Measurement cadence

A team of two to three people, with clear ownership across content, communications, and data, can execute this program without external support. Whether you need outside help depends less on the plan and more on whether your team has the capacity to hold the measurement cadence and the editorial discipline to produce content that is actually worth citing.

For communications and marketing leaders who want to understand the broader strategic context for this work, the age of agentic awareness is worth reading. The 3-phase plan above is the operational layer. That piece addresses the strategic question underneath it: why AI systems have become the intermediary that matters most in the pre-sales environment, and what that means for how communications functions need to be organized.

If you want help running the baseline audit or building the measurement system, we are available to talk.

[questions, answered]

FAQs

What is answer engine optimization?

Answer engine optimization is the practice of making your company easier for AI systems to understand, trust, and cite when they answer buyer questions. It is less about formatting tricks and more about clear entity signals, credible content, and third-party validation.

How do you start an AEO program?

Start with a baseline audit of how AI systems describe your company and competitors, then fix entity inconsistencies, publish answer-focused content, earn third-party validation, and measure changes over time. The three phases should focus on visibility, credibility, and repeatable measurement.

What should a B2B team measure in AEO?

Measure whether your company appears in category answers, which sources AI systems cite, and how accurately your company is described. Track those results across a defined set of prompts so you can compare month to month and against competitors.

Why do third-party sources matter for AEO?

Third-party sources matter because AI systems often rely on independent validation when building answers. Trade coverage, reviews, analyst commentary, and credible publications can reinforce your owned content and improve the odds that your company gets cited.

How long does it take to see AEO results?

Some changes can show up quickly, especially in systems that browse live web content, but most programs need multiple audit and update cycles before the impact becomes clear. A 3-phase plan is enough to establish the foundation, not finish the job.

Need a hand running the baseline audit?

We can help you run the baseline audit or build the measurement system.

Talk to Actual →

actualagency.ai/contact