



The Definitive AEO Playbook

Table of Contents

3	Introduction
4	The AEO playbook: A step-by-step framework
5	Step 1: Build a prompt index
8	Step 2: Define what success looks like
10	Step 3: Optimize your website
14	Step 4: Optimize third-party mentions
17	Step 5: Monitor your AI search performance
19	How to execute the AEO playbook
20	The four pillars of AEO maturity
21	The five stages of AEO maturity
24	Operationalizing AEO across the organization
26	The AEO playbook in review

The way people decide what to buy changed faster than most marketing organizations were able to adapt. Today, your customers ask an AI model, and largely take the answer at face value. As a result, the traditional keyword-centric playbook for search visibility doesn't map to the highly personalized, conversational prompts your audience uses to discover solutions today.

Achieving success requires rethinking how you measure visibility, structure your content, and organize your teams. This playbook breaks down the four core components of enterprise AEO: understanding the shift itself, measuring and winning AI visibility, benchmarking your organization's readiness, and operationalizing it enterprise-wide.

The AEO playbook: A step-by-step framework

The most common mistake marketing leaders make when assessing AI visibility is treating it as a single binary metric. Brand coverage varies based on persona and intent stage. A simplified *are we visible or not* view is misleading and fails to capture the nuanced ways AI personalizes information.

Step 1: Build a prompt index

The foundation of any successful AEO playbook is the **prompt index**. It's a curated, validated database of the exact conversational queries and complex questions your target audience asks AI engines.

Every metric you'll report on, every content decision you'll implement, pulls its accuracy from this one input. Build it loosely, and you'll spend a year optimizing against questions your audience isn't actually asking.


Often, an AEO index is organized into two layers: topics and prompts. Topics are the broad themes your business is looking to own and the content pillars you'd defend in a planning meeting. Prompts are the hundreds of specific questions per topic that reveal how real people actually ask about them. That distinction matters because AI engines evaluate topical authority, not keyword coverage, so the index has to mirror how models organize answers instead of how your site organizes pages.

Most indexes get assembled one of two ways: an LLM is asked to brainstorm plausible questions, or a vendor supplies a panel-derived list. Brainstormed prompts carry no demand signal and no awareness of what your site can actually answer. Panel data is small, generalized, and describes a market average rather than your buyers.

Most prompt panel datasets are pulled by creating automated bots that simulate human user behavior and prompt LLMs at scale, AKA scraping. While it's a straightforward solution, scraping **represents less than 1% of the real query landscape**, against the roughly 2.5 billion prompts a day on ChatGPT alone. Beyond the sample size, the personas used for these prompt panels often don't resemble the personas enterprise brands target.

Prompt panels can also be unstable. These datasets can be cut off with little notice when models update or privacy regulations change.

The test for a real prompt index is traceability. Every prompt should map back to something verifiable—a specific persona, a specific moment in the buying journey, content you own, and measurable search volume. If a prompt can't be traced to all four, it's a guess wearing a data label.



This is the standard worth holding any vendor to, including us. It's the reason **Conductor** builds your prompts from a map of your own content, so every suggestion reflects the topics you actually have authority on, rather than scoring you against a generic index, and grounds them in more than a decade of search demand data instead of panels.

A reliable prompt index requires three critical inputs:

- **Real search demand data:** Accurate SEO data is a window into real user intent, which makes it essential to unlocking great prompts. Root your prompts in actual user behavior, leveraging traditional SEO data, question-based modifiers, and long-tail query variations to understand what users are actually looking for.
- **Your own site and content footprint:** Analyze the topics and entities your website currently commands authority over to map existing strengths and identify immediate optimization opportunities.
- **Your competitive landscape:** Evaluate where competitors are winning citations to uncover gaps in your own coverage and prioritize high-value topics.

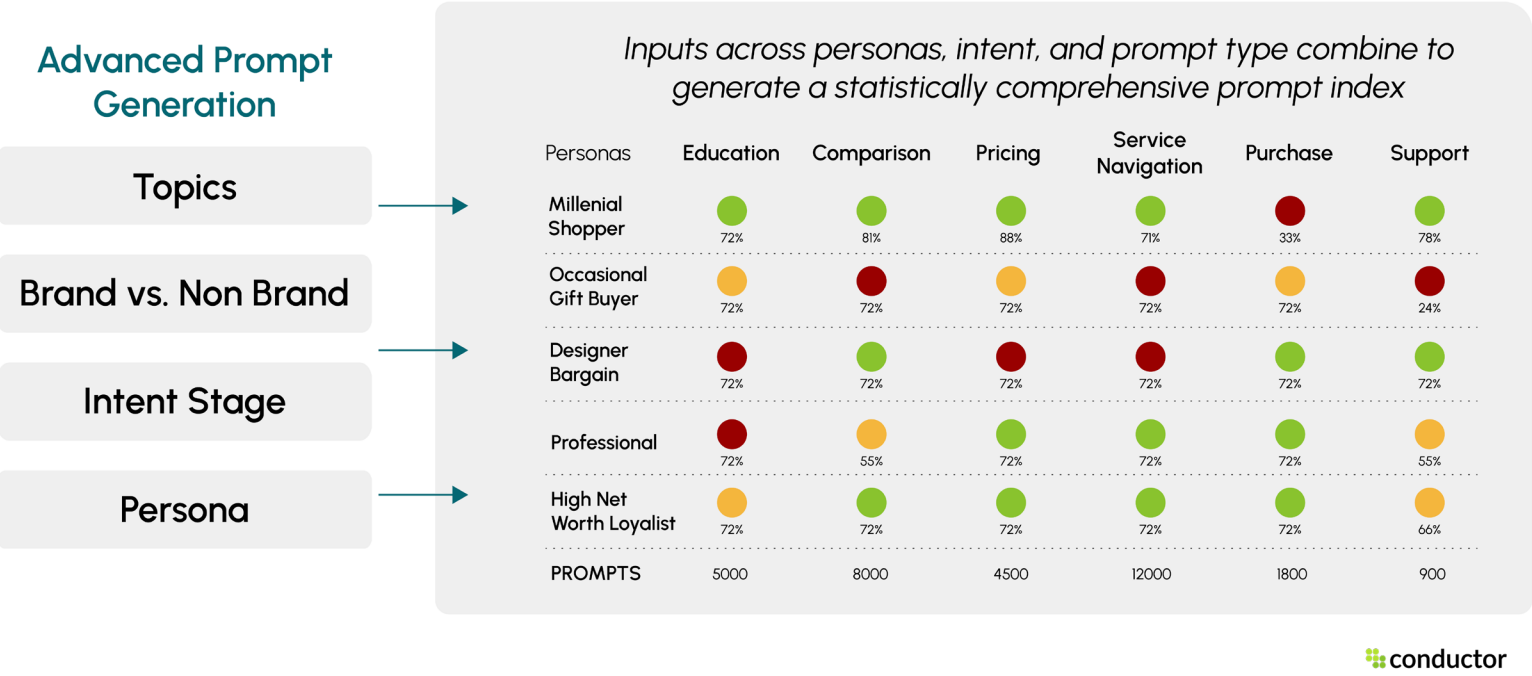
Once you gather these inputs, every topic must be expanded across three key dimensions. Expanding each topic across persona, intent stage, and brand versus non-brand is what makes the number actionable.

It's the difference between:

- We appear in 68% of AI answers
- Developers find us, CFOs don't; we own the education stage and disappear entirely at purchase.

The first tells you nothing you can act on. The second tells you exactly where to put your resources.

One benchmark to hold your team to: unbranded prompts should account for roughly 75% of your index. You should already be winning visibility for branded searches. Unbranded coverage tells you whether you're being discovered by people who don't know you exist yet.



Validation is critical here. Establishing measurement around these specific dimensions from the start, rather than bolting them on later, ensures your data accurately reflects how different users experience AI search.

See how Conductor builds your prompt index

Grounded in real demand data, your content footprint, and your competitive landscape—not panels or guesswork.

Request a demo

Step 2: Define what success looks like


Most teams inherit their definition of success from SEO: a ranking, a share, a number that goes up.

AI search doesn't reward traditional search rankings. It rewards being the answer.

The outcome worth targeting is a disproportionately positive recommendation: when an AI engine names or cites your brand, it does so favorably and in the context you'd choose.

That outcome isn't produced by any single asset. It's pieced together from everything an LLM can find about your brand. Your website is one source among several, and it's the only one you fully control.

LLMs pull the rest of their knowledge from third-party sites like analyst and trade coverage, G2 and Trustpilot, Reddit threads, and Wikipedia entries.



conductor

If your citation or mention sentiment is unfavorable, publishing more content alone won't fix it; the levers sit with other teams across the organization, like PR and product marketing.

Treating AEO as a content-only initiative caps the ceiling of the program before it starts.

The AI mention spectrum

Despite how many teams report on them, not all mentions and citations are worth the same.

Being named as the recommended solution, appearing as one of six alternatives, and being referenced without a link are three different outcomes that often get collapsed into one number in AEO tools with a narrow scope. The **AI mention spectrum** separates them, from primary linked citation down to unlinked passive reference.

Your program can hold a strong overall visibility score for two straight quarters while the quality of the mentions slowly degrades or skews neutral instead of positive.

Mention quality is just as important to measure as mention volume, and both need to be paired with sentiment analysis. Volume tells you whether AI engines see you. Quality tells you whether being seen is worth anything. Sentiment tells you which direction it's working in. A prominent, confidently cited negative mention does more damage to your brand and visibility than not appearing at all.

A CMO and a CRO can type identical words and get different answers, because the models infer context from profile, history, and phrasing. So the most common way teams evaluate AEO—running a few prompts and evaluating visibility based on the output—produces an anecdote, not a measurement.

Scale and repetition are what make the number defensible: a large, segmented prompt set, run consistently over time. Individual answers stay unpredictable; the distribution doesn't.

For example, say that overage runs at 100% on brand and service navigation prompts and under 20% on education prompts across nearly every persona. That's not variance, but a brand found by people who already know its name and is largely absent while buyers build a shortlist.




Step 3: Optimize your website

Of the two sources feeding an AI recommendation, your website is the one you have direct control over.

Optimization breaks into three areas:

- The **technical health** that makes citation possible.
- The **new content** that makes you eligible.
- The **content optimizations** that keep you competitive.



 conductor

1. Improve website health

If [AI crawlers can't reach or render a page](#), it won't be visible, and all of your optimization efforts will be for nothing. The page is either accessible or invisible.

Teams often create more content on top of sites with unresolved structural problems, then read flat AI visibility as a content quality issue, and commission more. Resolve the foundation first.

The good news is that issues can often be identified and resolved. For example:

- Content that only renders through JavaScript
- Missing or incorrect schema
- Canonical tags pointing at the wrong version of a page
- Server response times slow enough that crawlers abandon the request
- Orphaned pages with nothing linking to them
- Content sitting behind cookie walls, gates, or interstitials

None of these are strategic problems. Nobody owns them, and they don't surface in a content review, so it's critical to make sure your team is monitoring your site's technical infrastructure around the clock, in case issues arise.

A new consideration in AI search is what to allow AI crawlers to access. Managing your robots.txt so legitimate AI bots can reach your most valuable content while proprietary material stays protected is a web and IT conversation rather than a content one, and it's a conversation most marketing organizations have never formally had.

2. Create new content

Next, you need to prioritize creating new content. Every prompt you're competing for needs at least one owned page built to answer it directly, and with a unique, expert perspective on the topic. Generating AI content on a topic that sounds like everyone else's won't move the needle. You need to leverage your own authority and unique POV.

For example, brands should think about creating the following content:

- Create comprehensive "buyer's guides" to help users make purchase decisions
- Adding chunkable content like FAQs and pro tips to your content
- Create detailed product pages with optimized structure
- Create competitor comparison content

You can't be cited for an answer you never published, and optimization on adjacent content is no substitute for a page that addresses the prompt head-on.

3. Optimize your existing content

Finally, optimizing your existing pages is still critical to your visibility. This ensures that your most important content is fresh and has the most up-to-date insights and data behind it.

The first key to optimizing your content is extractability: clear headings, the direct answer stated up front, structured data, and content organized the way a model would want to lift it.

AI engines favor recently updated sources, which means a page that was authoritative in January can quietly stop being cited by June without anything about it changing. Updating statistics, examples, links, and claims isn't housekeeping—it's what keeps a citation you've already earned.

The useful shift here is where the diagnosis comes from. Your citation data should tell you which action to take in terms of creating net new content or optimizing existing content.

What you see in AI citations	What to do
Reddit, Wikipedia, generic sources	Build / Optimize
Low-quality or thin sites	Build / Optimize
Your competitors	Build / Optimize
Zero citations at all	Improve Site Health
Your own page already cited	Refresh


Every prompt deserves at least one strong owned page.
You can't be cited if you don't exist.



If a competitor holds the citation, you have a specific content gap and a benchmark to beat with new content. If Reddit, Wikipedia, or a generic aggregator holds it, no one owns the category, which is the highest-return opportunity on the board and the one most teams never look for. If your own page already holds it, the job is to optimize and ensure it's accurate as time passes.

Keep in mind that zero citations anywhere on a prompt you've created content around and want to be visible for is a technical signal that something is keeping your content from being properly crawled.

That turns content operations from a one-off publishing calendar into ongoing maintenance with weekly analysis, monthly refinement, and continuous updates.



It's also the point where AI earns its place in the workflow. Purpose-built AEO content generation tools, like [Conductor Creator](#), compress research, drafting, optimization, and refresh cycles and combine them with real-time search intelligence to make a more active cadence scalable at enterprise volume.

Recap:

What should you actually do

- **Confirm AI crawlers can reach and render your priority pages.** Check JavaScript rendering, robots.txt, and CDN-level bot rules before funding any new content.
- **Assign ownership of technical health.** Most of the time, these issues persist because nobody owns them, not because they're difficult to resolve.
- **Publish at least one strong owned page for every prompt you're competing for.** Eligibility comes before optimization.
- **Let citation data assign the work.** Competitor citations reveal content gaps. A generic source cited means an open category. Your own page cited means optimize and keep it current. Nothing cited means a technical problem.
- **Put refresh on a schedule, not a backlog.** How current and accurate your content is has become a visibility input, and it needs a standing resource rather than a campaign budget.

Turn your citation data into optimized AI visibility

Conductor tells you who holds the citation, where the gaps are, and what to do next, so your team optimizes the right pages, not just the most visible ones.

[Request a demo](#)

Step 4: Optimize third-party mentions

When a model assembles an answer about your category, it draws on sources that have no stake in the outcome, in addition to your own authored work. AI engines like to leverage these third-party sources because they're inherently more neutral and unbiased about your brand than your own published content is likely to be.


Third-party sentiment is the more difficult aspect of your site visibility to optimize, but it's also where AI recommendations are usually decided.

While with owned content, you earn visibility by building it, you can't build authority in third-party sources as easily. You can ship a page in a sprint. You can't ship a G2 review, an analyst mention, or an organic Reddit thread recommendation, and any attempt to manufacture them carries real **brand and compliance risk**.

The optimization work for third-party sites is divided into three areas:

1. Earned media and trade coverage
2. Reviews and ratings
3. Community participation

All three are influenced rather than controlled, which changes how they're planned, resourced, and measured.



conductor

Start where engines already look

What you can do is look at the review platforms, trade publications, and forums where AI engines are actually sourcing answers in your category, and concentrate your effort on those specific properties. The right solution can show you where AI is pulling information from and whether it's from a competitor or you.


Pro Tip:

It's key to measure your AI visibility against both your direct and indirect competitors to get a sense of how your visibility compares to the top names in and out of your industry.

You may not see some sources, like Wikipedia or Reddit, as competitors, but in terms of AI visibility, they are. Those sources, and others like them, are competing with brands like Nike for citations and mentions just as much as Adidas is. AI will pull what it deems to be the best information from the best source from across the web, not just your industry.

Most brands instead spread third-party investment across the outlets they've historically valued, which may have no overlap with the ones AI engines cite.

These levers sit with PR, customer marketing, and community teams that rarely see AI visibility data. Getting that data in front of them often costs fewer resources than other content strategies and tends to move the number faster.



conductor

Earn placements in sources that get cited

Tier 1 press, trade publications, and analyst coverage build the entity record LLMs rely on to establish category authority. The more consistently your brand appears in credible independent coverage, the more confident an engine will be in placing you in an answer.

Historically, this has been slow, manual, and hard to connect to any visibility outcome.

That's the gap Conductor's [partnership with Noble](#) closes. Conductor identifies where your brand is missing from high-citation sources and where competitors are winning them; Noble handles the activation, identifying which external properties matter most, automating outreach and negotiation, and managing placements end to end.

Manage reviews as a live signal

Review platforms are among the most consistently cited sources in AI answers about products and services, and they work differently from the ratings dashboards most teams monitor.

AI engines synthesize the substance of reviews, not the average rating. A strong overall rating carrying recent detailed complaints about onboarding, support responsiveness, or pricing can still produce an unfavorable summary, because the model is prioritizing what people wrote rather than solely relying on how they scored the product.

That makes three things critical for your visibility:

- **Recency** because models prioritize fresh input
- **Specificity** because detailed reviews give engines more to quote
- **Your response** because visible, substantive replies to criticism become part of the record that models read

It pays to reach out to customers and run campaigns to drive reviews to third-party sites, and it should really be a standing motion inside customer marketing rather than a quarterly push before renewal season.

Participate in community conversations without selling

Reddit, industry forums, and Wikipedia are the least controllable sources and often the most cited, which makes them the highest-risk and highest-return area on this list.

The rules of engagement are narrow, and getting them wrong is worse than not showing up:

- **Answer questions; don't pitch your product.** The contributions that earn visibility are the ones that help someone with a problem, including when the honest answer isn't your product. Communities identify promotional intent quickly and punish it harshly.
- **Disclose affiliation every time.** Undisclosed brand participation gets removed, and on Reddit, it can get a domain banned outright. Transparency isn't just compliance; a named expert who says who they work for is more credible than an anonymous recommendation that is later uncovered to be an employee.
- **Send people, not brand accounts.** Subject-matter experts posting under their own names build durable authority. Corporate accounts dropping links build nothing.
- **Provide facts, don't manage narrative.** Politely fixing an inaccuracy with a source can add to a conversation. Arguing with a negative opinion won't, and it tends to generate more of the sentiment you were trying to fix.
- **Leave Wikipedia to editors.** Editing your own entry violates conflict-of-interest guidelines and typically gets reverted. The productive path is ensuring credible independent sources exist for editors to cite, and using talk pages to flag factual errors with references. You can also work with Wikipedia editors to ensure that content on pages about you is accurate. Wikipedia wants to provide readers with correct and verifiable information, so they often work with brands on ensuring information is up to date and accurate.

This is a sustained effort that needs a dedicated owner, not a one-off campaign, which is why most organizations struggle with it.

Recap: What should you actually do


- Identify which third-party sources AI engines cite in your category—then compare that list against where you currently invest.
- Get citation data in front of PR, customer marketing, and community. These teams control the levers and usually can't see the scoreboard.
- Close citation gaps systematically rather than through one-off outreach, and measure placements by their effect on visibility rather than by volume.
- Create a plan and run campaigns to drive user reviews, and monitor review substance, not just the star average.
- Assign a named owner for community participation—with clear rules on disclosure, and a mandate to help rather than sell.

Step 5: Measure and monitor your AI search performance

Every step so far is critical, but none of them is a one-and-done process. Models retrain, competitors publish new content, your own pages age out of favor, and a technical change ships without anyone connecting it to AI visibility drops.

So this step is really about ensuring measurement and monitoring are a continuous process. Here's a 4-step process to implement to ensure your efforts are informed by the latest visibility data, and you're able to report on true visibility ROI:

1. **Track:** What activities are we logging across content and third-party sources?
2. **Diagnose:** What's moving and where are we stuck?
3. **Decide:** Which pages are winning and need a refresh? Which ones aren't winning and need new content?
4. **Act:** How do we improve our work cross-functionally and implement solutions faster?



TRACK

What activities are we logging across content and third-party sources?

DIAGNOSE

What's moving and where are we stuck?

DECIDE

Which pages are winning and need a refresh? Which aren't winning and need new content?

ACT

How do we unblock cross-functionally and go faster?

conductor

A team completing this cycle monthly compounds its position against a competitor running it quarterly, and given how durable citation patterns become once established, that gap widens rather than closes.

Coverage erodes over weeks as competitors gain ground or content ages, and a regular review cycle catches it in time to respond.

Technical issues produce instant visibility drops. An accidental disallow directive, a broken schema implementation, or a rendering change that hides your primary content from AI crawlers doesn't produce a gradual decline. They remove you from answers entirely. That requires **real-time alerting** on site health, not a line item in a monthly report.

It's also worth knowing that this loop tends to break at the handoff from a visibility insight to a person on another team with the ability to do something about it. If the diagnosis lands with content, PR, or web development as a request with no priority attached, the optimization loop has essentially stalled.

The brands winning AI search are the ones watching it continuously

Conductor monitors your AI search performance in real time—alerting your team to citation drops, technical failures, and competitive shifts before they compound into revenue problems.

[Request a demo](#)

How to execute the AEO playbook

Executing this playbook requires an executive perspective to align the right people, establish the right processes, and deploy the right technology.

To close the gaps and successfully prioritize AEO strategies, you need to evaluate where your organization stands today and map a deliberate strategy toward maturity.

Keep in mind that a mature SEO program doesn't directly translate into AEO maturity. Solid technical foundations and a proven enterprise content strategy give you a head start, but they typically don't get you to long-term maturity.

SEO maturity is built around rankings, crawlability, and backlinks. AEO maturity is built around citability, entity clarity, and whether your teams can act on shared visibility data and insights.

The four pillars of AEO maturity

	STAGE 1 Reactive	STAGE 2 Operational	STAGE 3 Strategic	STAGE 4 Authority-First	STAGE 5 Agentic
People	SEO Owner	SEO & Content Aligned	Cross-functional AEO Team	Enterprise Governance	AI-guided Teams
Process	Reactive Fixes	Repeatable Workflows	AEO Embedded in Planning	Authority Managed Across Channels	Automated Optimization
Tech	SEO Tools	Centralized Data	Unified Intelligence	Enterprise Visibility Intelligence Layer	AI Agents & Orchestration
Metrics	Traffic & Rankings	Visibility Tracked	Citations & Share of Voice	Authority & Business Impact	Predictive Forecasting & Optimization



The maturity of your AEO program is carried by these four pillars:


- **People:** Who owns it
- **Process:** How the work gets done
- **Tech:** What technology supports it
- **Metrics:** What gets measured

An organization can license an enterprise platform, build sophisticated dashboards, and still generate Stage 2 outcomes, because AEO remains one person’s secondary responsibility, and nothing in the workflow forces a decision.

The right tool can accelerate a strategy, but it can’t create one from scratch.

Expect unevenness as you measure, as metrics are likely to vary by business unit or topic area rather than resolving into one tidy score. The distance between your strongest and weakest pillar is your roadmap; closing that gap usually returns more than advancing the pillar you’re already good at.

The five stages of AEO maturity



The five stages of AEO maturity

Stage 1 — Reactive

- **Goal:** Understand how to get visible.
- **The tell:** Executives are asking about AI search, and no one can answer them.

- **People:** Individual contributors or agencies managing SEO and content independently, with no dedicated AEO ownership
- **Process:** Reactive, ad hoc responses to AI search changes with no documented workflows or editorial standards
- **Metrics:** No formal AI visibility tracking; executives are starting to ask questions teams can't yet answer
- **Tooling:** Disconnected point solutions with no unified view of AI or traditional search performance

Stage 2 — Operational

- **Goal:** Centralize and organize the effort.
- **The tell:** You track AI visibility, but only content and SEO care about it.

- **People:** Core SEO and content teams aligned on AEO basics, with emerging ownership and shared editorial standards
- **Process:** Documented content optimization workflows in place; basic AI search policies established but not yet cross-functional
- **Metrics:** AI search visibility is being tracked at the channel level, but not yet tied to broader business KPIs
- **Tooling:** An AEO platform like Conductor is in use for core teams, providing unified visibility across AI and traditional search

Stage 3 — Strategic

- **Goal:** Drive business value cross-functionally.
- **The tell:** AI visibility shows up in revenue conversations; coordination is still manual.

- **People:** AEO spans marketing/brand, PR, product, and web teams; a dedicated leader or function owns the program with executive visibility
- **Process:** Cross-functional workflows connect content, technical, and brand teams around a shared AEO strategy and editorial calendar
- **Metrics:** AI visibility tied to revenue metrics; shared dashboards serve multiple departments with consistent reporting
- **Tooling:** AEO platform integrated directly with CMS, CRM, and BI tools, providing a unified source of truth across functions

Stage 4 — Authority-First

- **Goal:** Be the definitive answer everywhere.
- **The tell:** The mandate comes from the C-suite; workflows are still more manual than they should be.

- **People:** C-suite mandate drives company-wide adoption; every customer-facing team has AEO embedded in their day-to-day responsibilities
- **Process:** Automated cross-team reporting and real-time alerting replace manual coordination; brand messaging is consistently governed at scale
- **Metrics:** Real-time executive dashboards track share of voice, sentiment, and citation share across all AI engines
- **Tooling:** AEO platform serves as the single source of truth for the entire enterprise, feeding live intelligence into every function

Stage 5 — Agentic

- **Goal:** Orchestrate and automate the whole process.
- **The tell:** Agents handle routine execution; your senior people spend their time on governance.

- **People:** Senior strategists and brand governance leads oversee the program; AI agents handle routine execution tasks
- **Process:** Agentic loops automatically monitor, surface, and act on AEO insights; human practitioners focus on compliance and strategic pivots
- **Metrics:** AI-native funnel analytics prove full-funnel ROI; automated reporting surfaces anomalies and opportunities in real time
- **Tooling:** AEO platform powers AI-native workflows, feeding real-time intelligence directly into agentic pipelines across the enterprise

Stages 1 and 2 are content-centric—AEO done from the bottom up by practitioners who found the problem before the organization did. This is where most organizations fall right now.

Stages 3 through 5 are based on authority and are the AI answer, driven from the top down. That's not an incremental step up the ladder. It's a change in who's holding the wheel.

This is why the Stage 2 to Stage 3 transition is the one most organizations struggle with. Everything before it can be accomplished by a motivated team. Nothing after it can be done without executive buy-in and sponsorship.

What each transition actually requires:

- **Stage 1 → 2:** Establish clear ownership of the AEO function, connect visibility tracking, and align SEO and content on a shared editorial plan.
- **Stage 2 → 3:** Earn executive sponsorship, extend AEO beyond the search and content teams, and integrate AEO technology with the broader enterprise stack.
- **Stage 3 → 4:** Achieve company-wide adoption of AI search intelligence, build a real-time reporting infrastructure, and convert KPI tracking into a top-down mandate.
- **Stage 4 → 5:** Implement agentic workflows to reduce manual overhead without loosening brand governance or compliance.

Stage 1 to Stage 2 moves relatively quickly, because it's largely an ownership and tooling problem you can solve inside your own org.

Stage 2 to Stage 3 depends on executive engagement and cross-functional cooperation that you can't unilaterally assign, which is why most organizations don't make this leap.


Stages 4 and 5 require sustained investment across all four pillars before you can start automating.

Again, our experience with hundreds of enterprise customers shows that most companies are in Stage 1 or 2 today. Translation: the majority of your market is stuck at the point where an executive decision is needed.

Operationalizing AEO across the organization

Transforming your website into an AI-ready growth engine requires breaking down the traditional silos that separate SEO, content, and technical web teams.

To successfully operationalize AEO across a complex enterprise, organizations have to adopt a unified approach to intelligence and execution.



Six things have to be true for a program to run at enterprise scale.

- **Assign ownership by citation source type.** Owned content and site structure sit with content, SEO/AEO, and web development. Technical health sits with web and IT. Earned media and analyst coverage sit with PR and comms. Reviews sit with customer marketing and customer success. Community presence sits with product marketing and named subject-matter experts. Measurement and reporting sit with marketing ops. Without that map and regular collaboration between these teams, AEO insights land as unassigned requests, and visibility is lost regardless of tooling.
- **Monitor AI and traditional visibility together, not in separate systems.** Teams need one view to weigh trade-offs and connect performance to business impact rather than defending two disconnected scorecards.
- **Generate and optimize content built to be cited, not just ranked.** The structural requirements differ, and they're most efficiently applied while content is being written rather than retrofitted afterward.
- **Activate third-party sources deliberately.** Identifying which external properties AI engines actually cite, closing those gaps systematically, and measuring the visibility effect has to be a managed motion with a named owner, not the byproduct of whatever PR happened to land that quarter.
- **Monitor site health and AI crawlability in real time.** Technical failures will remove you from answers until someone notices.
- **Make AEO intelligence available to any agent, LLM, or BI tool your teams already use.** Intelligence that lives only in a dedicated platform reaches only the people who log into it. Getting it into the systems where PR, product, and web teams already work is what turns cross-functional ownership from an org chart into a habit.

Conductor is the only end-to-end, enterprise AEO platform built on the industry's most complete data engine to power this exact transformation.

By combining unparalleled AEO and SEO intelligence, scaled AI content generation, and real-time website monitoring into one secure platform, Conductor replaces disconnected tools with a single source of truth.

The AEO playbook in review

Building a successful and scalable AEO strategy comes down to building an index you can defend, defining success as a favorable recommendation rather than a ranking, fixing the surfaces you own, influencing the ones you don't, and keeping the loop turning.

What will separate the brands that own AI answers from the ones that disappear into them is whether the organization is built to execute that strategy. That's an executive question rather than a practitioner one, and it's exactly why most companies' AEO maturity is stalled while the citation patterns in their category quietly solidify around someone else's expertise. Visibility in AI search compounds for whoever establishes it first. Honestly assess where your AEO strategy stands, and where you have the most opportunity to grow, then start optimizing.

