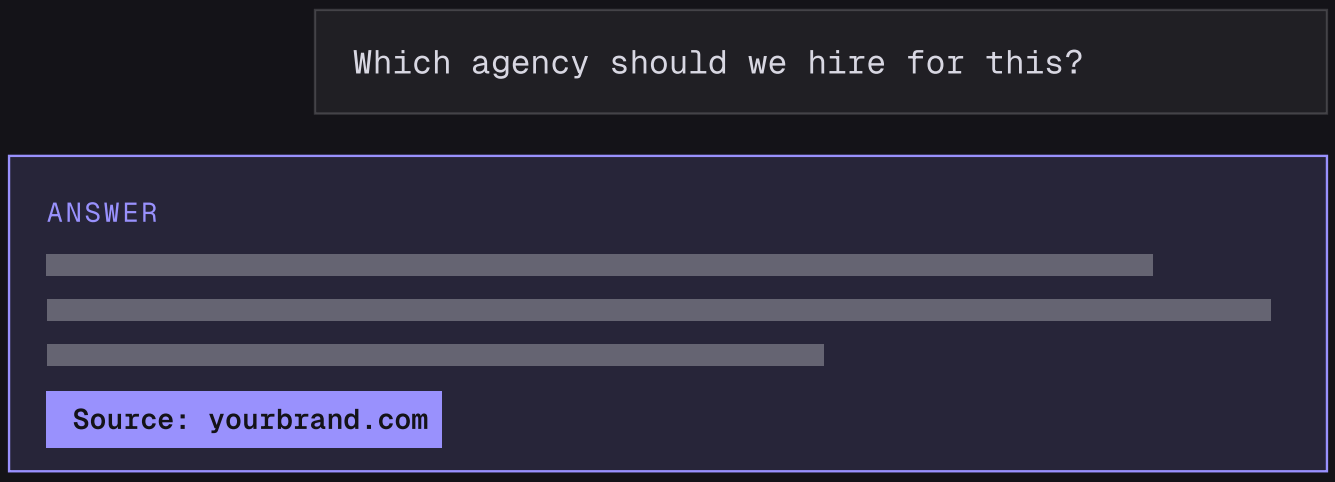
KOOKIES

GUIDE

The AEO Readiness Guide

6 CHAPTERS



SEO AND AEO

The AEO Readiness Guide

An AEO-ready site lets AI tools reach its pages, states facts in plain answer-first copy, labels them with accurate structured data and covers the questions buyers ask. Most of this is good SEO done with more care. Nobody can guarantee a citation, so this guide tells you what to fix, how to check it and what to measure.

Who it is for. Marketing leads, founders and developers who run a marketing site and want it ready for AI search.

01 What AEO is and what changes

02 Answer-first copy

03 Schema per page type

04 llms.txt and AI crawlers

05 FAQ, comparison and glossary pages

06 Measuring it, and what nobody can promise

Read online at www.kookies.io/guides/aeo-readiness. Updated 2026-11-07.

CHAPTER 1 OF 6

What AEO is and what changes

01

What stays the same, what changes, and the four levels of readiness this guide follows.

The short answer

AEO, answer engine optimisation, is the work of making your pages easy for AI answer tools to reach, read, quote and describe correctly. Most of it is good SEO done with more care. What changes is the unit of competition: you are no longer only trying to win a click on a whole page. You also want single passages on that page to stand up on their own.

If you want the plain definition and the history of the term, read what is AEO. This guide starts after that point. It is a working plan: what to check, what to rewrite, what to generate, and how to tell whether any of it is done.

What stays the same

Start with the most useful sentence a vendor has written on this subject. Google says there are "no additional requirements" to appear in AI Overviews or AI Mode, and that no special optimisation is necessary. To be shown as a supporting link, a page must be indexed and eligible to appear in Search with a snippet. Google also says that indexing and serving are not guaranteed.

So the old foundations still carry the weight:

- **Crawling.** Robots rules and your host or CDN let the crawler in.
- **Indexing.** The page is in the index and allowed to show a snippet.
- **Internal links.** People and crawlers can find the page from other pages.
- **Text.** The important content is written as text, not locked inside an image or a video.
- **Page experience.** The page works well for the person who lands on it.

If those are broken, fix them before you read further. See our Webflow SEO checklist if your site runs on Webflow. A clever answer on a page nobody can crawl does nothing.

What changes in the work

The tools differ, and each company publishes only part of how its product works. But Google describes one mechanism that changes how you should plan content. It calls this query fan-out: AI Overviews and AI Mode may run several related searches across subtopics to build one response. A single buyer question can become many smaller ones. That is why a page full of clear sub-answers has more chances to be useful than one long, vague page. That last step is our reading, and no vendor promises it.

HABIT	SEARCH-FIRST VERSION	ANSWER-READY VERSION
Opening	A scene-setting intro, then the point	The answer first, then the detail
Headings	Clever or brand-led	The question a buyer would ask
Facts	Spread across paragraphs, images and PDFs	Stated in one plain sentence on the page
Trust	Mostly left to the brand	Named author, dates, sources and checkable details
Consistency	Each page words the offer its own way	One source for prices, services and names
Measuring	Rankings and clicks	Those, plus access checks, sampled answers and referrals

Our comparison of the terms lives in the glossary: AEO and GEO. For this guide, use whichever word your team likes. The work is the same list.

Three kinds of AI access

Many teams treat "AI bots" as one thing and either allow them all or block them all. The vendors' own docs split them by job, and this is the distinction that matters most for your decisions.

1. **Training.** Crawlers that collect content that may be used to train models. GPTBot and ClaudeBot are described this way by OpenAI and Anthropic.
2. **Search.** Crawlers that help a site appear in an AI tool's search answers. OAI-SearchBot, Claude-SearchBot and PerplexityBot do this job.
3. **User fetches.** A visit made because a person asked a question or gave a link. ChatGPT-User, Claude-User and Perplexity-User do this job.

Blocking the first group is a policy choice about training. Blocking the second can remove you from an answer tool's results. Chapter 4 has the full table and the robots.txt choices.

Who should do this first

Use this rule to decide how much effort to spend.

- **Do all of it** if buyers research before they talk to you: B2B, SaaS, agencies, professional services and considered purchases. They ask for options, comparisons and prices.
- **Do levels 1 and 2 below** if you are a small site with few pages. An hour on access and an afternoon on your main page copy is proportionate.
- **Hold off on new content** if your core pages still cannot say what you do, who it is for and what it costs. Fix those first.
- **Think again about access** if your content is your product, such as paid research or licensed data. The decision is commercial, and you may want advice on it.

The four levels of readiness

The rest of this guide follows four levels. Finish each one before you spend much on the next.

LEVEL	QUESTION IT ANSWERS	WHAT YOU DO	CHAPTER
1. Reachable	Can the tools get to the page?	Check robots.txt, firewall rules, indexing and snippet settings	4
2. Readable	Is the content clear to a machine and a person?	Answer-first copy, text in the HTML, clear headings	2
3. Described	Are the facts labelled and consistent?	JSON-LD per page type, one source for facts, an llms.txt if you want one	3 and 4
4. Quotable	Do you have the pages that answer buyer questions?	FAQ, comparison and glossary pages, then review over time	5 and 6

Being mentioned on other sites sits alongside all four. Our post on getting cited by ChatGPT and Perplexity covers that side, so we do not repeat it here.

How we run it on this site

This site is built the way the guide describes. Many pages open with a short answer box. JSON-LD is generated from the same data files that render the page. The build writes llms.txt from the live pages. Our robots.txt allows AI crawlers by name. We chose that because we want to be found. We make no claim here about what that setup has earned us, and you should be wary of anyone who claims a fixed result from these steps.

If you would like a team to do the work, see our SEO, AEO and GEO service. Otherwise, carry on to the copy chapter.

TAKEAWAYS

- AEO is mostly good SEO, with single passages that stand on their own.
- Google says its AI features have no extra requirements beyond normal Search eligibility.
- AI access has three jobs: training, search and user fetches. Decide on each one separately.
- Work in order: reachable, readable, described, quotable.

CHAPTER 2 OF 6

Answer-first copy

02

The lift test, a simple shape for every section, and before and after rewrites.

The short answer

Answer-first copy puts the answer in the first sentence or two of a page or section, then adds the conditions and the detail. Each answer should still make sense if someone lifts it out of the page, so it names its subject and states the fact. No vendor publishes a formula for being quoted. This is a writing habit that makes a page easier for people and machines to use, and that is all we claim for it.

The lift test

Copy one sentence from your page into an empty document. Read it with no other context. Then ask four questions.

1. **Does it name the subject?** "It is fast" fails. "Astro sends little JavaScript by default" passes.
2. **Does it answer a question someone would ask?** If you cannot write the question above it, it is probably a slogan.
3. **Does it contain a fact?** A price, a time, a condition, a name or a clear yes or no.
4. **Is it true everywhere you say it?** If the pricing page and the FAQ disagree, a tool may repeat the wrong one.

Run the test on the first sentence under every main heading. Anything that fails goes on the rewrite list.

A simple shape: answer, condition, detail

Most sections can follow the same three steps.

- **Answer.** One or two sentences that settle the question.
- **Condition.** When the answer changes: size, plan, region, type of buyer.
- **Detail.** Steps, examples, a table and links to go deeper.

A worked example, using our own prices: "A Kookies landing page costs $500. A website of up to 5 pages starts at $2,500, and larger sites start at $3,000. You get a fixed quote before we start." Answer, then condition, then detail. Each sentence is usable alone.

Before and after

These are sample rewrites to show the pattern. Square brackets mark facts you must supply, and only true ones. If you do not have the fact, the rewrite is blocked until you find it.

PAGE	BEFORE	AFTER
Home hero	Elevating brands through digital excellence.	[Company] builds marketing websites in [platform] for [type of customer].
Pricing	Flexible packages tailored to your unique needs. Get in touch.	Plans start at [price] a month. Most teams of [size] pick [plan]. Enterprise is quoted.
Service	We handle migrations with care and attention to detail.	We move [old platform] sites to [new platform]. Every old URL gets a 301 redirect to its new page, so links and bookmarks keep working.
Comparison	Why we are the best choice.	Pick [us] if you need [strength]. Pick [rival] if you need [its strength].
Timeline FAQ	Every project is different, so timelines vary.	Most projects take [time]. The usual delay is [cause]. A site with [condition] takes longer.
Support FAQ	We are always here for you.	Support runs [days and hours] by [channel]. We reply within [time].

Notice what happened in each row. The rewrite is more specific, and it can be wrong. That is why it is useful. A sentence that cannot be checked gives nobody anything to quote. For more on this, our AEO post has a full page example.

Rewriting headings

A heading tells the reader and the machine what the next passage answers. Keep it close to the words buyers use.

BEFORE	AFTER
Our approach	How a Webflow migration works, step by step
Investment	How much a website costs
Peace of mind	What happens after launch
Why us	When to hire an agency instead of a freelancer

You do not have to make every heading a question. A brand page can keep some voice. The rule is that a reader who skims only the headings should still learn what the page covers. To find the buyer's words, read your sales call notes, your inbox and your support tickets before you open a keyword tool.

Where the answers should go

- **Top of the page.** One short block: what this is, who it is for, what it costs or how long it takes.
- **Top of each section.** The first sentence after a heading carries the answer.
- **In the FAQ.** One question, one direct answer. Chapter 5 has the rules.
- **In tables.** Prices, plans and comparisons work well as simple HTML tables, with a sentence above that states the conclusion.

Keep all of this in real text. Google's guidance lists keeping important content in text form as a fundamental. Quick check: open your page, view the source and search for the sentence. If it only appears after scripts run, or it exists only in an image, move it into the HTML. Vendors do not all say how their crawlers treat scripts, so HTML is the safe choice.

What to leave out

- **Throat clearing.** Delete any first sentence that says how important the topic is.
- **Invented numbers.** A made-up statistic is easier to quote and worse for you when someone checks it.
- **Questions nobody asks.** Stuffing a page with questions makes it worse to read.
- **Text written only for machines.** Do not hide text from readers. It helps nobody, and structured data rules ask for the opposite: mark up what people can see.
- **Overclaims.** "The best agency" is opinion. "Webflow Premium Partner" can be checked, if it is true for you.

A rewrite session in six steps

1. Pick five pages: home, pricing, your top service, your top comparison, your FAQ.
2. Write the three questions each page must answer, in the buyer's words.
3. Check whether each answer appears in the first two sentences of a section. Mark it yes or no.
4. Rewrite the nos with the answer, condition, detail shape.
5. Gather the facts in one source, such as a data file or a CMS collection, so a price changes in one place only.
6. Have someone who knows the product read each new sentence and sign it off as true.

On this site, prices live in one data file, and the pricing page opens with an answer in plain sentences. That is the model to copy, whatever platform you use. Need help with the words? See our strategy and copywriting service.

TAKEAWAYS

- ✓ Put the answer in the first one or two sentences of each page and section.
- ✓ Every key sentence should name its subject and survive being lifted out of the page.
- ✓ Specific claims can be wrong, which is why they are useful. Only state true ones.
- ✓ Keep answers in HTML text, and keep facts in one source.

CHAPTER 3 OF 6

Schema per page type

03

A page type to schema table, the FAQPage change, decision rules and testing.

The short answer

Give each page type the structured data that truthfully describes it, in JSON-LD, generated from the same data as the page. Start with Organization and BreadcrumbList, then add Article, Service, Product or FAQPage where the page really is one. Be clear about what this buys you. Google says there is no special schema you need to add for AI Overviews or AI Mode, and that structured data should match the visible text. Schema labels your facts. It does not earn citations by itself.

For syntax, examples and how to add it in Webflow or Astro, read our schema markup guide. This chapter is the decision table.

Page type to schema: the table

The last column shows whether the feature was listed in Google's structured data gallery when we read it. This changes, so check the current docs. "No" does not make the type wrong. It means you should expect no special Google display from it.

PAGE TYPE	SCHEMA TYPE	FILL IN AT LEAST	IN GOOGLE'S GALLERY
Every page	Organization (on the home page at least), BreadcrumbList below home	Name, URL, logo, contact, social profiles; each crumb's name and URL	Yes, both
Home	Organization, WebSite	Site name and URL	Organization yes, WebSite no
Service	Service	name, serviceType, provider, areaServed, offers if public	No
Physical product	Product	Name, price, availability, real ratings	Yes
Software product	SoftwareApplication	Name, description, rating if real	Yes (Software app)
Blog post or guide	Article or BlogPosting	headline, image, datePublished, dateModified, author	Yes (Article)
Case study	Article	Same as above	Yes (Article)
Glossary term	DefinedTerm	name, description, inDefinedTermSet	No
Page with a visible FAQ	FAQPage	Each question and its visible answer	No longer

PAGE TYPE	SCHEMA TYPE	FILL IN AT LEAST	IN GOOGLE'S GALLERY
Physical location	LocalBusiness	Address, hours, contact	Yes (Local business)
Team or founder profile	ProfilePage with a Person	Name, role, links	Yes (Profile page)
Comparison page	BreadcrumbList, plus FAQPage if it has a visible FAQ	Skip Product markup for rivals' products	See above

Build yours with the schema generator, then check it with the tests below.

FAQPage: what changed

Google's documentation says FAQ rich results no longer appear in Google Search, and it has removed the documentation for that feature. Plenty of older advice still treats FAQ markup as a way to win extra space on the results page. Do not plan on that.

Should you remove it? Our rule:

- **Keep it** if it is generated from your visible FAQ and costs you nothing to maintain. FAQPage is still a type on schema.org, and it states your questions and answers clearly. We have not read any vendor saying its AI product uses it, so treat it as tidy labelling, nothing more.
- **Remove it** if it is hand-written, out of date or does not match the page. Wrong markup is worse than none.
- **Do not add it** to win a display in Google. That display is gone.

The visible FAQ is what matters. Chapter 5 covers it.

Decision rules

1. **Is it visible?** Google's guidelines say not to mark up content that readers cannot see. If it is not on the page, it does not go in the markup.
2. **Is it true of this page?** Markup must represent the page truthfully. A service page is not a Product because Product has richer results.
3. **Is it complete?** Items missing required properties are not eligible for rich results. Article has no required properties in Google's docs, but add the recommended ones that apply.
4. **Who owns it?** If nobody updates it when the page changes, generate it or leave it out.
5. **One main type.** A page can hold several items. Include a main type that reflects what the page is about, and connect related items with `@id`.

Article and author details

Most marketing sites get this slightly wrong, and it is easy to fix. Google's Article docs recommend `headline`, `image`, `datePublished`, `dateModified` and `author`. For authors, the docs ask you to:

- include every author shown on the page

- give each author a separate entry
- put only the name in `author.name`, with no job title or company in that field
- use `Person` for people and `Organization` for organisations
- add a `url` or `sameAs` with a valid address

Only change `dateModified` when you really changed the content. A date that moves every build tells readers nothing.

Generate it from one source

Hand-written schema drifts away from the page. Build it from the data that renders the page instead.

- **Coded sites.** A shared function takes the page data and prints JSON-LD. On this site, `src/lib/seo.ts` does this.
- **Webflow.** Site-wide Organization code goes in site settings. Per-page and CMS schema goes in page custom code, using CMS fields. Need a hand? See our Webflow SEO service.
- **Sanity or any headless CMS.** Keep the fields in the content model and print them at build time.

Test it in three places

Google names three tools, and they catch different problems.

TOOL	WHEN	WHAT IT TELLS YOU
Rich Results Test	While building	Whether the markup is valid, and sometimes a preview
URL Inspection	After publishing	Whether Google found the structured data on the page
Rich result status reports	Ongoing	Whether a template or serving change broke the markup later

These tools catch technical errors, not quality problems. Google says quality needs a human check, so also read the markup against the page. Structured data issues can also bring a manual action. Google says that costs a page its rich result eligibility, and that it does not affect ranking. You can see any in the Manual Actions report.

What this site uses

For a real reference: every page built with our shared helper includes an Organization node, plus BreadcrumbList below the home page. Service pages add Service, blog posts add BlogPosting with a named author, case studies add Article, glossary terms add DefinedTerm, and pages with a visible FAQ add FAQPage from the same data as the FAQ. You can see the terms in our JSON-LD and schema markup glossary entries.

TAKEAWAYS

- Mark up only what is visible and true of the page.
- Google says no special schema is needed for its AI features.
- Google no longer shows FAQ rich results. Keep FAQPage only if it is accurate and cheap to maintain.
- Generate JSON-LD from page data, then test it in the Rich Results Test and URL Inspection.

CHAPTER 4 OF 6

llms.txt and AI crawlers

04

Who crawls for what, your robots.txt policy, firewall checks and when llms.txt is worth it.

The short answer

Crawler access comes first and llms.txt comes a long way second. Decide which AI crawlers may visit, set that in robots.txt, then confirm your host or firewall is not blocking them anyway. After that, an llms.txt file is a cheap, optional extra. Google says you do not need AI text files or new machine-readable files for its AI features. The llms.txt site itself describes the format as a proposal.

Who is crawling, and for what

This table comes from each vendor's own crawler page. Names and behaviour change, so check the current docs before you act.

NAME	COMPANY	JOB	WHAT BLOCKING DOES
GPTBot	OpenAI	Content that may be used to train models	Signals the content should not be used in training
OAI-SearchBot	OpenAI	Helps sites appear in ChatGPT search	Not shown in ChatGPT search answers, though navigational links can still appear
ChatGPT-User	OpenAI	Actions started by a user	OpenAI says robots.txt rules may not apply
ClaudeBot	Anthropic	Content that may be used to train models	Future material should be excluded from training
Claude-SearchBot	Anthropic	Improves search results	May reduce visibility in search responses
Claude-User	Anthropic	Fetches pages for a user's question	Claude cannot retrieve the site for that question
PerplexityBot	Perplexity	Surfaces and links sites in its results, not for training foundation models	Perplexity recommends allowing it
Perplexity-User	Perplexity	Visits a page when a user asks	Perplexity says it generally ignores robots.txt
Google-Extended	Google	A control token for Gemini training and grounding, not a separate crawler	Does not affect inclusion or ranking in Google Search

Two details are easy to miss. OpenAI says each of its settings is independent, so you can allow search and refuse training. And Google-Extended is not the control for what Search shows. To limit that, Google points to `nosnippet`, `data-nosnippet`, `max-snippet` or `noindex`.

Choosing your robots.txt policy

There are three sensible positions. Pick one on purpose.

POLICY	SUITS	SET IN ROBOTS.TXT
Allow all	Marketing sites that want reach	No blocks. Naming each crawler is optional, because `User-agent: *` with `Allow: /` covers those without their own rule.
Search yes, training no	Teams that want to appear in answers but keep content out of training where they can	Add `Disallow: /` under `User-agent: GPTBot`, `User-agent: ClaudeBot` and `User-agent: Google-Extended`.
Block most	Paid content, licensed data, members' material	Block search crawlers too, and put real access control on the content. Robots rules are a request.

This is a business choice, and for licensed content a legal one. We are not giving legal advice here, so ask your lawyer if your rights to the content are in question.

Some practical notes. Anthropic's rules apply per subdomain, so repeat them for each one. Anthropic supports the non-standard `Crawl-delay` if a crawler is loading your server. OpenAI says robots.txt changes can take about a day to affect search. Our own robots.txt uses the first policy.

robots.txt is not the only gate

A site can say yes in robots.txt while its firewall says no. Google's guidance names the CDN and hosting layer as places to allow crawling, and it is easy to miss because it never appears in your code.

- **Check your firewall and bot settings.** Look at your CDN's bot rules and security events for blocked AI user agents.
- **Match the name and the address.** Anyone can fake a user agent. OpenAI, Anthropic and Perplexity all publish IP lists, so allow rules should use both. Perplexity's docs give Cloudflare and AWS WAF examples.
- **Know the Cloudflare catch.** Perplexity notes that Bot Fight Mode cannot be skipped with WAF custom rules and may still challenge its bot.
- **Do not block by IP to opt out.** Anthropic warns against it, because a blocked bot cannot read your robots.txt, so the opt-out may not work reliably.
- **Keep lists fresh.** Perplexity says its ranges change and suggests automating updates.

Test what your site tells each crawler with the AI crawler checker. Read about the wider group in our AI crawler entry.

llms.txt: when it is worth it

The format, from llmstxt.org, is short: a Markdown file at `/llms.txt` with an H1 name, a blockquote summary and H2 sections of links with notes. An H2 called Optional marks links an agent can skip. The site describes the file as read on demand when an agent needs context for a user's request, which is a different job from robots.txt or a sitemap.

Our llms.txt guide has a full example, so here is the decision rule only.

- **Add it** if you can keep it in sync with little effort, ideally from your build or your CMS.
- **Skip it for now** if access, copy or schema are unfinished. They come first.
- **Prioritise it** if you have documentation or an API. A curated index read on demand fits that kind of content well, in our view.
- **Do not expect** a Google benefit. Google has said it is not needed for its AI features, and we have not read a vendor page that promises one.

Make a first version with the llms.txt generator. Then write the summary yourself. It should state what you sell, to whom, and your key facts in plain sentences.

How we generate ours

A hand-kept file goes stale. Ours does not, because a script runs after each build. It walks the built pages, skips anything marked noindex, reads each title and description, and writes grouped links: company pages first, then services, migrations, comparisons, industries, case studies, blog, locations and glossary. The summary at the top holds the prices and contact details. It also writes llms-full.txt with the full text of core pages, and a separate llms-integrations.txt so the main file is not swamped.

The spec also suggests Markdown copies of pages at matching URLs. Our build writes the full-text files instead. Pick whichever you can maintain.

Access checklist

- You have written down your policy: allow all, search yes and training no, or block most.
- robots.txt matches that policy and returns a 200 on every subdomain you care about.
- A sitemap is linked from robots.txt.
- CDN or host bot rules do not contradict robots.txt.
- Staging rules, such as a site-wide `Disallow` or `noindex`, did not reach production.
- Key pages are indexed and allowed a snippet. Check with URL Inspection.
- Important text appears in the HTML source.
- If you have an llms.txt, it opens, its links work and its prices match the site.
- Someone owns a review of this list after any platform or host change.

TAKEAWAYS

- ✓ Crawler access comes first. llms.txt is a cheap, optional extra.
- ✓ You can allow AI search and refuse training, because the vendors separate them.
- ✓ Check your CDN and firewall, because robots.txt is not the only gate.
- ✓ Google-Extended does not affect inclusion or ranking in Google Search.

CHAPTER 5 OF 6

FAQ, comparison and glossary pages

05

Match each question to a page type, then write FAQs, comparisons and glossary entries worth quoting.

The short answer

FAQ, comparison and glossary pages suit AI answers because each maps to a question people actually ask: how does this work, which one should I pick, what does this mean. Nobody can promise any of them will be quoted. But if you only build three kinds of content for answer tools, make it these three, written to the rules below. Use real questions, direct answers and honest conclusions.

Match the question to the page

WHAT THE BUYER ASKS	PAGE TO BUILD	FIRST SENTENCE SHOULD
What is X?	Glossary term	Define X in one sentence
How much, how long, can you, do you?	FAQ section or single FAQ page	Give the number or the yes or no
X or Y? X vs Y?	Comparison page	Say who should pick each one
How do I do X?	Guide or blog post	Sum up the method
Who should I hire for X?	Service page, plus reviews and case studies	Say what you do, for whom, at what price

Google's description of query fan-out helps explain why several small, clear pages can fit together: one big question may be split into subtopics. That is an inference on our part, and it does not mean more pages win. A thin page adds nothing to any of them.

FAQ: rules for answers worth reading

Collect questions from places where real people asked them: sales calls, your inbox, support tickets, proposals that stalled. Then apply these rules.

1. **Ask it the way buyers ask it.** "Can I edit the site myself after launch?" beats "Post-launch content autonomy".
2. **Start the answer with the answer.** Yes, no, a number or a name.
3. **Keep it short.** Two to four sentences, with a link to the full page if there is more to say.
4. **Include the awkward ones.** Refunds, what is not included, when you are the wrong choice. Buyers do ask them.
5. **One source.** Prices and terms should come from the same place as the rest of the site.

Section or own page? Use a section on the relevant page by default. Give a question its own page only when all three are true: people ask it on its own, it needs more than a few sentences, and you can say something useful that is not on another page. Otherwise you are building duplicates.

We run both: FAQ sections on service pages, and a FAQ hub with single pages for bigger buyer questions.

Comparison pages: be fair or do not bother

A comparison page written as an advert has nothing in it worth repeating. A fair one can be the most useful page on your site. This structure works, and you can see it on our Webflow vs Astro page and the rest of the comparison hub.

1. **Verdict first.** "Pick A if you need this. Pick B if you need that." Two sentences.
2. **Short "best for" lines.** One for each option.
3. **One table.** The same rows for each option, covering what a buyer weighs: editing, speed, cost model, limits, who runs it.
4. **Where the other option wins.** Say it plainly. If you cannot name a case, the reader will not trust the page.
5. **A short FAQ.** Migration, cost and the common confusion.

Before and after, using sample wording:

BEFORE	AFTER
Unlike other tools, our platform gives you everything you need.	Pick [A] if your marketing team edits pages daily. Pick [B] if developers run the site and you need [specific feature].

Fact rules for comparisons:

- Read each rival's own docs for every claim about it. If you cannot confirm a fact, leave it out.
- Leave out competitor prices unless you can keep them current. They change, and a wrong one damages the rest of the page.
- Do not mark up a rival's product with your own ratings.
- Show when the page was last checked, and do the check.

Glossary pages: small and exact

A glossary works when the terms are ones your buyers meet while choosing what you sell. A list of generic marketing words does little for anyone.

- **Sentence one is the definition.** "A 301 redirect is a permanent redirect from one URL to another." Do not begin with history.
- **Sentence two says why it matters** to the reader's decision.
- **Then an example,** a common mistake, and links to the service or guide that goes further.
- **Link terms from your other pages.** Internal links are on Google's list of fundamentals, and they help readers too.

- **Skip terms you have nothing to add to.** If your entry could be anyone's, leave it out.

BEFORE	AFTER
In the ever-changing world of SEO, canonical tags have become a key consideration.	A canonical URL tells search engines which version of a page is the main one, when several URLs show the same content.

Ours is in the glossary, with an entry such as AI citation as a model. Each term carries DefinedTerm markup and links onward.

Keep them true over time

These pages go out of date faster than most, because they describe other people's products and your own prices. Set a simple routine:

- **Prices and terms:** pull from one data source, so there is nothing to update by hand.
- **Comparisons:** re-read the vendor docs on a schedule, and whenever a vendor makes a major change.
- **FAQ:** add a question each time sales hears a new one, and remove any nobody asks.
- **Dates:** change the shown date only when the content really changed.

What about other sites?

Your own pages are only part of what these tools read. Reviews, directories, partner pages and articles by others also describe you. Keep those facts in line with your site: same name, same services, same location. Our post on getting cited has the practical steps. The rule that fits this chapter: when someone else's page describes you wrongly, ask them to correct it before you write a new page of your own.

Content checklist

- Every question on the site came from a real buyer or a real search, and you can say which.
- Each FAQ answer starts with the answer.
- Each comparison has a verdict in the first two sentences and says where the other option wins.
- Each rival fact was read on the rival's own page.
- Each glossary entry opens with one sentence that defines the term.
- Each of these pages links to at least one page that helps a buyer act.
- Nothing is duplicated across pages just to target another phrase.
- Someone owns the review schedule.

Short of time to write these? Talk to us about our SEO and AEO service.

TAKEAWAYS

- Match the question to the page: definition, direct answer or fair verdict.
- Collect questions from sales calls and support, not only keyword tools.
- A comparison must say where the other option wins, or nobody will trust it.
- Give a question its own page only if it has something new to say.

CHAPTER 6 OF 6

Measuring it, and what nobody can promise

06

Three layers of measurement, a monthly prompt panel, honest limits and a one-page audit.

The short answer

Measure what you control every month: access, indexing, page checks and copy coverage. Sample what you do not control: whether AI tools name you and how accurately. Count what reaches you: referrals and leads who say they found you through an AI tool. Treat all of it as a direction of travel. Nobody can promise a citation, and no single score can tell you that you will get one.

Three layers of measurement

LAYER	WHAT TO TRACK	WHERE	HOW OFTEN
Inputs you control	Technical checks passed, pages with an answer block, valid schema, crawler access	AEO score checker, AI crawler checker, Rich Results Test	Monthly, and after any release
Search presence	Indexed pages, impressions and clicks for your question queries	Google Search Console	Monthly
Answers you can sample	Named or not, cited or not, accurate or not	Your own prompt panel, below	Monthly
Business signals	AI referrals, form answers, sales notes	Analytics and your lead sheet	Monthly

The score checker fetches one page and runs 12 technical checks, including title and description length, one h1, canonical link, JSON-LD, AI crawler rules, llms.txt and a sitemap. The score is the number passed. It is a checklist of signals, and it does not predict rankings or citations. Use it to find gaps, then stop looking at it.

What Google reports

Google says clicks from AI features count towards overall Search traffic and appear in the Performance report under the Web search type. In practice, watch your total clicks and your question-style queries together. Check Google's current docs for any newer breakdown before you build a report on this.

If a page should appear and does not, work through the basics: use URL Inspection to confirm indexing and what Googlebot sees, and check for `nosnippet` or `noindex`. If you change those, Google says recrawling can take from several days to several months.

Build a prompt panel

This is manual, and it is the most useful thing on this page. A prompt panel is a fixed list of buyer questions that you ask the same tools each month.

1. **Write 10 to 20 questions** from sales calls. Mix three kinds: about the category ("how do I move from WordPress to Webflow"), about options ("what should I look for in a Webflow agency") and about you ("what does [your company] charge").
2. **Choose your tools.** Pick the ones your buyers use, such as Google's AI features, ChatGPT, Claude and Perplexity.
3. **Log each answer in a sheet.** Columns: date, tool, question, were you named, were you linked, which URL, were the facts right, which sources were cited instead.
4. **Act on it.** A wrong fact goes to the copy or to a third-party correction. A rival's page that is often cited is worth reading to learn what it says clearly that you do not.

Be honest about the limits. Answers can change between runs and between people, so a single result proves little. Look for patterns over months, and do not celebrate or panic over one screenshot. See AI visibility and AI citation for the terms.

Count what reaches you

- **Referrals.** In your analytics, filter sources for domains such as chatgpt.com, perplexity.ai, claude.ai and gemini.google.com. Expect this to undercount, because a person may read an answer and then search for your name or type your address.
- **Ask people.** Add "How did you hear about us?" to your form with a free text box, and have sales note the answer on calls. If someone says "ChatGPT told me", that counts, even when analytics shows direct traffic.
- **Branded search.** Watch searches for your name in Search Console. It has many causes, so use it as context and not as proof.

Our post on marketing site tracking covers the setup behind these numbers.

What nobody can promise

Be direct about this with your team, and expect the same from anyone you hire.

- **A citation.** Each company decides how its product chooses sources. Google says indexing and serving are not guaranteed, even when requirements are met.
- **A rich result.** Google says valid markup does not guarantee one.
- **Accuracy.** A tool can get your facts wrong even on a clear page. See AI hallucination.
- **Lasting placement.** Products and their sources change.
- **A benefit from llms.txt.** It is a proposal, and Google says it is not needed.
- **Traffic.** An answer may satisfy the reader with no click at all.

Red flags when buying this work: guaranteed citations, a secret formula, a score presented as a ranking factor, pages generated in bulk with no human check, hidden text for AI, or fake reviews. Ask instead: what will you change, how will we check it, and what will you do if the sampled answers do not move?

The one-page audit

Print this, and mark each line yes, no or not relevant. Start with your five main pages.

Reachable

- A written policy for training, search and user crawlers exists.
- robots.txt follows it, and the CDN or firewall agrees.
- The main pages are indexed and allowed a snippet.
- Important text is in the HTML source.

Readable

- The page's first block says what it is, who it is for and the key fact.
- The first sentence under each main heading passes the lift test.
- Headings use the buyer's words.
- Prices, times and limits appear as text, and you are happy to be quoted on them.

Described

- Organization and BreadcrumbList are present and valid.
- Each page type has the schema from the table in chapter 3, and it matches the visible page.
- Articles name a real author, with honest dates.
- Facts come from one source, and third-party profiles agree with it.
- llms.txt is current, or you have decided not to have one.

Quotable

- A visible FAQ answers real questions, including the awkward ones.
- Comparison pages open with a fair verdict and rival facts you read at the source.
- Glossary entries open with a definition, and exist only where you can add something.

Measured

- A prompt panel is logged monthly, with a named owner.
- AI referrals and "how did you hear about us" answers are reviewed monthly.
- The audit runs again after any redesign or platform move.

What to fix first

Fix in this order, and stop when the next item is not worth your time.

1. Anything that blocks access or indexing.
2. Wrong facts, on your site or elsewhere.
3. Answers on your five main pages.
4. Schema on your templates.
5. Missing FAQ, comparison or glossary pages your buyers ask about.
6. llms.txt and other extras.

Then get on with running the site. If you would like us to run the audit or the fixes, talk to us or look at our SEO, AEO and GEO service.

TAKEAWAYS

- Measure inputs, sample answers and count what reaches you.
- Google reports AI feature clicks within the Web search type in Search Console.
- A prompt panel shows patterns over months. One result proves little.
- Be wary of anyone who guarantees citations.

QUESTIONS

Questions people ask.

Is AEO different from SEO?

It builds on SEO. Google says its AI features have no extra requirements beyond normal Search eligibility. The added care goes into direct answers, consistent facts and checking how AI tools describe you.

Can anyone guarantee that ChatGPT or Google will cite my site?

No. Each company decides how its product picks sources, and Google says indexing and serving are not guaranteed. You can make your pages easier to reach, read and quote, then measure what happens.

Do I need an llms.txt file?

No. Google says you do not need AI text files for its AI features, and llms.txt is a proposal. It is a cheap extra if you can keep it in sync, and it should come after access, copy and schema.

Should I block GPTBot and ClaudeBot?

That depends on your view of training. Both vendors describe these crawlers as collecting content that may be used for training, and search has separate crawlers. So you can block training and still allow OAI-SearchBot and Claude-SearchBot. If content rights are in question, ask your lawyer.

Is FAQ schema still worth adding?

Not for a Google display, since Google no longer shows FAQ rich results. Keep FAQPage if it is generated from your visible FAQ and easy to maintain. Remove it if it is out of date or does not match the page.

How do I know if AEO work is doing anything?

Check three things monthly: technical inputs, a fixed list of buyer questions asked in the AI tools your buyers use, and referrals plus what leads tell you. Look for patterns over months, because a single answer proves little.

Keep going

- What is AEO?
- How to get cited by ChatGPT and Perplexity
- llms.txt explained
- Schema markup guide
- GEO in the glossary
- SEO, AEO and GEO service

Check one page with the free AEO score checker

Kookies. www.kookies.io. hello@kookies.io