



PixelGem

# SEARCH HAS CHANGED

**SEO, AEO & GEO for Small Businesses in the Age of AI**

*A practical UK guide to getting found, becoming the answer and earning trust in AI-powered search.*

**SEO**

Get found

**AEO**

Become the answer

**GEO**

Become source-worthy

**From PixelGem**

Based in Cheshire, helping small businesses locally and across the UK.



# SEARCH HAS CHANGED

---

## *SEO, AEO & GEO for Small Businesses in the Age of AI*

**Search Has Changed**

**SEO, AEO & GEO for Small Businesses in the Age of AI**

A practical UK guide to getting found, becoming the answer and earning trust in AI-powered search.

Published by PixelGem

[pixelgem.co.uk](http://pixelgem.co.uk)

First edition - September 2026

# COPYRIGHT & IMPORTANT NOTE

Copyright © 2026 PixelGem. All rights reserved.

This book is educational. Search engines, AI products, ranking systems, crawler behaviour and reporting tools change frequently. Nothing in this book guarantees a ranking, citation, recommendation, lead or sale.

The research in this book is an original synthesis of public guidance, product documentation and published research. The wording, examples, frameworks and recommendations are original to this book. Where a factual claim depends on an external source, the source is identified in the research notes and references section.

The terms **AEO** (Answer Engine Optimisation) and **GEO** (Generative Engine Optimisation) are useful industry labels, not universal technical standards. Different platforms work differently, and no single checklist can force an AI system to cite or recommend a business.

UK spelling and examples are used throughout.

Worked examples that are labelled illustrative are fictional composites based on common small-business situations. They are included to show how the principles can be applied and are not presented as PixelGem client results.

## A QUICK WORD FROM PIXELGEM

Search has always changed. What is different now is the speed of the change and the number of places where a customer can get an answer.

A few years ago, a typical journey was simple: someone searched Google, scanned a results page, clicked a website and made up their own mind. That journey still exists, but it now sits alongside AI Overviews, AI Mode, Bing and Copilot answers, ChatGPT search, voice assistants, map results and increasingly capable agents that can inspect websites on a user's behalf.

That can make ordinary business owners feel as though everything they learned about SEO has suddenly expired. It has not.

The strongest lesson from the current official guidance is actually reassuring: the fundamentals still matter. A useful website that is technically accessible, clear about what the business does, accurate about where it works, rich in first-hand evidence and easy for a customer to use is valuable to both classic search and AI-assisted discovery.

What has changed is the job description of that website.

It is no longer only trying to win a click. It may also be trying to supply a fact, support a comparison, answer a question, confirm a business detail, provide evidence or become one of several sources used to build an answer.

This book is about understanding that shift without getting dragged into hype. We will use SEO, AEO and GEO as three useful lenses on the same practical problem: **how do you make your business easy to find, easy to understand and easy to trust?**

If you run a small business, that is the bit that matters.

# HOW TO USE THIS BOOK

You do not need to read this with a developer sitting next to you. The technical sections explain enough to help you ask better questions and spot obvious problems without turning the book into a coding manual.

Each chapter has three jobs:

- explain what has actually changed;
- separate useful practice from fashionable nonsense;
- give you something practical to improve.

If your website is already performing well, use the book as a review. If it is struggling, work through the chapters in order and use the 90-day plan near the end.

***The rule running through the whole book:** do not create content for an imaginary algorithm at the expense of the customer standing in front of you. Search systems are becoming better at interpreting meaning. The safest long-term strategy is to become a better source, not a louder one.*

# CONTENTS

## Part One - The Search Shift

- 1. From Search Box to Answer Layer 8
- 2. SEO, AEO and GEO: Three Lenses, One Job 13
- 3. How AI Search Actually Finds and Uses Sources 18

## Part Two - Foundations That Still Matter

- 4. SEO Foundations in 2026 23
- 5. Technical Visibility: Crawl, Index, Render and Discover 28
- 6. Page Experience, Accessibility and Trust 33

## Part Three - Becoming the Answer and the Source

- 7. AEO: Answering the Question Clearly 38
- 8. GEO: Becoming Source-Worthy 46
- 9. Evidence, Expertise and Originality 52
- 10. Information Architecture for Humans and Machines 57
- 11. Structured Data Without the Hype 62

## Part Four - The Local Business Playbook

- 12. Local Search in an AI-Search World 68

|                                             |    |
|---------------------------------------------|----|
| 13. Location Pages Without Doorway Spam     | 77 |
| 14. Reviews, Reputation, Links and Mentions | 83 |
| 15. Service Pages That Actually Convert     | 89 |

## Part Five - Measurement, Control and Action

|                                                                    |     |
|--------------------------------------------------------------------|-----|
| 16. Measuring Visibility When Clicks Are No Longer the Whole Story | 99  |
| 17. Crawlers, Preview Controls and Content Permissions             | 107 |
| 18. Agent-Ready Websites                                           | 113 |
| 19. Common Myths, Mistakes and Expensive Distractions              | 120 |
| 20. Your 90-Day Search Visibility Plan                             | 127 |

## Appendices

|                                        |     |
|----------------------------------------|-----|
| A. Service Page Blueprint              | 137 |
| B. Customer Question Bank              | 140 |
| C. Technical and Local Audit Checklist | 142 |
| D. Glossary                            | 144 |
| E. Research Notes and Further Reading  | 147 |
| F. Practical Examples by Business Type | 156 |
| G. Monthly Search Visibility Review    | 166 |

## PART ONE - THE SEARCH SHIFT



### CHAPTER 1

# FROM SEARCH BOX TO ANSWER LAYER

For most of the web's history, search placed the burden of interpretation on the user. You typed a query, received a list of links and decided what deserved your attention.

That model is still alive, but a second layer has grown on top of it. Search products can now interpret longer questions, break a complex request into smaller searches, retrieve material from several places and synthesise a response before the user visits a website.

Google describes AI Overviews and AI Mode as search experiences that can use multiple related searches to develop a response.

Microsoft describes a similar shift in terms of \*grounding\*: connecting an AI system to current web information so the answer is based on evidence outside the model itself. OpenAI's search product can likewise retrieve public web material and cite sources.

For a small business owner, the important change is not the model name. It is the customer behaviour that follows.

A customer may still type:

- electrician Northwich
- dog groomer Winsford
- accountant near me

But they can also ask:

- "Who is a reliable electrician near Northwich who can replace an old consumer unit, and what should I ask before booking?"
- "Which dog groomers around Winsford take nervous dogs and are open on Saturdays?"
- "What should a small limited company expect to pay for an accountant, and what is normally included?"

Those questions contain far more intent. They reveal the problem, the location, the concerns and sometimes the buying criteria.

That means the website has a new opportunity. A page can be useful before the click by contributing an accurate fact or explanation, and it can be useful after the click by giving the person enough confidence to act.

## The old mental model

**Query -> ranked links -> website visit -> decision**

## The newer mental model

**Question -> retrieval -> generated or enriched answer -> sources and businesses -> website or direct action -> decision**

The newer model does not replace the older one. In fact, the infrastructure underneath still depends heavily on crawling, indexing, relevance and quality. The difference is that the search system may do more of the early reading.

That is why "rank number one" is no longer a complete strategy. It remains useful, but it cannot describe all the ways a business may be discovered.

## Fewer clicks does not mean no opportunity

Research on user behaviour shows why this matters. Pew Research Center analysed tracked Google searches from U.S. adults in 2025 and

found that people clicked traditional results less often when an AI summary appeared. The study also found that direct clicks on citations inside the summaries were rare. This was a U.S. study at a particular moment in Google's rollout, so it should not be treated as a universal conversion law. It does, however, illustrate the basic commercial problem: a business can be visible without receiving the same click pattern it would have received from a classic results page.

By 2026, Pew survey data also showed that a majority of U.S. adults said they had read AI summaries at the top of search results. Again, the exact number is less important to a UK small business than the direction of travel: generated answers are no longer a niche behaviour.

So the question becomes wider:

- Are people seeing your business name even if they do not click immediately?
- Is your website supplying facts that strengthen an answer?
- Does a later branded search happen because someone first encountered you in an AI response?
- Can somebody move from an answer to a map listing, phone call, booking or enquiry without taking the journey you expected?

## A practical example

Imagine two roofers covering the same Cheshire towns.

The first has a page that says:

*Quality roofing services across Cheshire. Contact us today for a free quote.*

The second explains the types of roofs it works on, the areas it genuinely covers, how emergency call-outs differ from planned work, what affects the price of a repair, how a survey works, what

photographs are taken, what happens if hidden damage is found and how long common jobs normally take.

Which page gives a search system more usable evidence? Which page gives a customer more confidence? Which business is easier to compare with the alternatives?

That is the heart of the shift.

***Do this next:** Write down the five longest, most specific questions customers have asked you in the past month. Those questions are often more useful than a list of generic keywords.*

## What this means on an ordinary Monday morning

It is easy to talk about AI search as though every customer has become a technology enthusiast overnight. They have not. Most people are still trying to get something done. They need a roofer, an accountant, a restaurant, a garage, a solicitor, a new website or somebody who can fix a problem without making their week worse.

The change is that the question can now be more complete before the customer ever reaches your site.

A person who once searched for garage near me might now ask whether a garage can diagnose an intermittent engine warning light, whether it works on a particular make, whether a courtesy car is available and roughly how long diagnostics take. The business with the clearest, most useful information has more chances to become part of that early research.

That does not mean writing a page for every sentence anybody could possibly type. It means looking at the real buying journey and making the useful parts explicit.

A good exercise is to take one recent enquiry and write down what the person genuinely needed to know before they felt comfortable

contacting you. Usually it is not just the service name. It might be price, suitability, timing, area, risk, evidence, availability or what happens next.

Those details are the raw material for modern search visibility.

Do not panic if click-through rates become harder to read. Someone may see your name in an answer, search for you later, open Maps or ring the number in a business listing. Attribution gets messier, but the commercial job is familiar: be visible at the right moment and give people enough confidence to choose you.

The sensible response is neither "AI changes nothing" nor "everything is different". Both miss the point. Search still depends on the web, but increasingly reads it on the customer's behalf before they decide where to go.

## CHAPTER 2

# SEO, AEO AND GEO: THREE LENSES, ONE JOB

The digital marketing world loves a new acronym. Sometimes the acronym describes a real change. Sometimes it simply gives an old idea a new jacket.

SEO, AEO and GEO are useful only if they help you make better decisions.

## SEO - Search Engine Optimisation

SEO is the established discipline of helping search engines discover, understand and rank web content while making that content useful to the people searching for it.

In practice, SEO covers things such as:

- crawlability and indexability;
- site structure and internal linking;
- relevance to real search intent;
- page quality and originality;
- links, mentions and reputation;
- local business signals;
- page experience and conversion paths.

SEO remains the foundation. Google's current guidance explicitly says its generative search features are rooted in core Search systems and that standard SEO practices remain relevant.

## AEO - Answer Engine Optimisation

AEO usually means **Answer Engine Optimisation**. It is not an official Google standard and there is no universal AEO score.

The practical idea is simple: make your information easy to use when someone asks a question.

AEO encourages you to think about:

- the exact question a customer is trying to answer;
- a clear answer early in the section;
- headings that reflect real intent;
- definitions, steps, comparisons and exceptions;
- concise factual wording followed by useful context;
- content that works well in snippets, assistants and voice-style answers.

AEO is not "writing for robots". Done properly, it is closer to good customer service on a page.

## GEO - Generative Engine Optimisation

GEO usually means **Generative Engine Optimisation**. The term gained prominence through academic work exploring visibility in generative search engines and has since been adopted more broadly by the search industry. Microsoft now uses the term publicly when discussing how content participates in AI-generated experiences.

GEO asks a slightly different question:

**If an AI system is going to build an answer from several sources, why would your page be useful evidence?**

That brings attention to:

- clarity and factual consistency;
- original information;
- first-hand experience;
- evidence that can support a claim;
- freshness;
- clear business and author identity;
- source-worthiness rather than keyword repetition.

## The easiest way to remember the difference

| Lens | Main question                                                   | Useful outcome                                                         |
|------|-----------------------------------------------------------------|------------------------------------------------------------------------|
| SEO  | Can I be found and ranked?                                      | Search visibility and qualified visits                                 |
| AEO  | Can my content answer the question clearly?                     | Useful answer visibility and better customer understanding             |
| GEO  | Can a generative system retrieve, trust and use my information? | Citations, mentions, assisted discovery and stronger source visibility |

The lines are deliberately blurry.

A well-written service page can help all three at once. A case study can earn a link, answer a buyer's question and provide original evidence for a generated response. An accurate Google Business Profile can support map visibility, classic search and location-aware AI answers.

## Why the acronyms can become dangerous

The problem starts when a business treats each acronym as a separate campaign.

You end up with an "SEO page", an "AEO page" and a "GEO page" all saying roughly the same thing in slightly different language. That creates duplication, weakens the user experience and can drift into scaled content created mainly to manipulate visibility.

Google's 2026 generative AI optimisation guidance is unusually clear on this point: businesses should prioritise normal SEO, valuable non-commodity content and satisfying user experiences rather than chasing special AEO/GEO hacks, unnecessary AI text files or mass-produced query variations.

So do not build three websites in one.

Build one excellent source.

***A useful shorthand:** SEO helps you get discovered. AEO helps you become a clear answer. GEO helps you become credible source material. The customer still decides whether you deserve the business.*

## Where the three disciplines meet in real work

In a small business, it is rarely useful to run three separate projects called SEO, AEO and GEO. The same piece of work can improve all three.

Imagine an electrician has a weak service page that says little more than "professional rewire across Cheshire". A useful improvement might involve explaining what a rewire normally includes, how the first inspection works, what can change the price, whether people usually need to move out, what areas are genuinely covered and what happens after the quote is accepted. Add photographs from real jobs, identify the qualifications that actually apply and make the enquiry route obvious.

That is traditional SEO because the page becomes more relevant, useful and internally connected. It is AEO because important customer questions now have clear answers. It is GEO because the page contains specific, attributable facts and first-hand evidence that a retrieval system can use with less guesswork.

The labels are therefore most useful as review questions:

1. **SEO:** can search systems discover the page, understand what it is about and connect it to the right audience?
2. **AEO:** does it answer the questions a real customer would ask before making a decision?
3. **GEO:** does it contain enough clear, supportable information to be worth using as a source?

If one answer is weak, improve the page. You do not need a separate "GEO department" to do it.

This is also why guaranteed AI recommendations should make you wary. A business can improve its eligibility, clarity and evidence. It cannot dictate which sources a platform will retrieve for every future question, which model will be used, what competing information will exist or how the final answer will be phrased.

A practical optimisation programme works on what the business controls: the quality of the site, the truthfulness of the information, the strength of the evidence, the consistency of the business data and the ease with which a person can take the next step.

## CHAPTER 3

# HOW AI SEARCH ACTUALLY FINDS AND USES SOURCES

You do not need to understand transformer architecture to improve a small-business website. You do, however, benefit from understanding the journey between a public web page and a generated answer.

A simplified version looks like this:

**Crawl -> Index -> Interpret -> Retrieve -> Ground -> Generate -> Cite or link -> User action**

Different products do this differently, and some stages can happen in different orders, but the model is useful because it shows where ordinary website work still matters.

## 1. Crawl

A search system first needs access to the page. Crawlers follow links, read sitemaps and request URLs. Firewalls, bot protection, authentication, bad robots.txt rules and accidental `noindex` directives can block discovery before content quality is even considered.

This is why a site can have excellent writing and still be invisible.

## 2. Index

Once a page is discovered, the search engine may process and store a representation of it in an index. Indexing is not guaranteed. Duplicate pages, thin pages, canonical issues, rendering problems and low-value content can all affect what gets stored or surfaced.

Google's AI features require a page to be indexed and eligible to appear in normal Search with a snippet. There is no separate "AI index application form" for AI Overviews or AI Mode.

### 3. Interpret

The system tries to understand what the page is about: the business, topic, entities, services, locations, claims and relationships.

Clear language helps. So does a coherent site structure. A page called "Services" that mixes roofing, web design, dog grooming and cryptocurrency would give both people and machines a hard time.

### 4. Retrieve

When a user asks a question, the system looks for material relevant to the request.

Google says its AI search experiences may use a technique it calls **query fan-out**, issuing several related searches across subtopics and data sources. That matters because a page does not necessarily need to match the user's exact sentence word-for-word to become relevant.

A detailed page about boiler replacement may be useful for a question about expected cost, disruption, timing, radiator compatibility or whether an old system needs upgrading.

This is one reason the old habit of making dozens of near-duplicate keyword pages is becoming less defensible.

### 5. Ground

Grounding connects a generated answer to external evidence.

Microsoft describes grounding as a critical layer that helps AI use current, authoritative information rather than relying only on what a model learned during training. In May 2026, Microsoft explained the difference neatly: traditional search helps a human decide what to read; grounding helps an AI system decide what it can responsibly say.

That distinction explains why evidence matters so much.

A page that says "we are the best" offers almost nothing that can support a factual answer. A page that explains a method, publishes a price range with caveats, documents a completed project or clearly states opening hours provides something concrete.

## 6. Generate

The model synthesises information into an answer. It may combine several sources, paraphrase them, compare them or decide that the evidence is not strong enough to make a claim.

This is where the idea of "ranking" becomes less tidy. A page might be retrieved but not cited. It might be cited but contribute little to the final answer. It might influence the answer without generating a click.

Recent academic work has started separating **citation selection** from **citation absorption**: being selected as a source is not necessarily the same as materially shaping the final answer. This research is early, but the distinction is useful for marketers because it discourages vanity metrics.

## 7. Cite, link or surface

Some products show source links or citations. Some show business entities, maps or cards. Others may send referral traffic.

Bing Webmaster Tools now reports AI citations for Microsoft AI experiences. Google Search Console introduced dedicated generative AI performance reporting in 2026 and, according to Google's update, expanded the visibility to websites worldwide by the end of August 2026.

OpenAI says referral links from ChatGPT search include the `utm_source=chatgpt.com` parameter, which makes those visits easier to identify in analytics.

The reporting is improving, but it is still not as mature as decades-old search analytics.

## 8. User action

The business outcome can happen after any of these stages.

The user may:

- click the website;
- search the brand later;
- open Maps;
- call directly;
- compare two businesses;
- ask a follow-up question;
- use an agent to inspect a booking flow.

That is why the final measurement chapter in this book tracks more than visits.

***Do this next:** Pick one important service page and ask: if a machine had to extract three reliable facts from this page, what would they be? If you struggle to answer, the page may be too vague.*

## Ranking a page and using a fact are not quite the same job

Microsoft made an important distinction in 2026 when it described the difference between traditional search indexing and grounding for generated answers. Classic search is largely trying to decide which pages a person should visit. Grounding has a slightly different responsibility: it needs information that can safely support what the system is about to say.

That distinction is useful even if you never use Bing Webmaster Tools.

A human looking at ten search results can compensate for a mediocre result. They can skip it, open another tab, compare two sources and

change the query. A generated answer compresses some of that work into the system. If the retrieved information is stale, ambiguous or contradictory, the problem can be carried into the answer.

For a business website, this raises the value of details that are easy to verify:

- a price range with a clear date or explanation of what affects it;
- opening hours that agree with the business listing;
- a service area stated consistently;
- a process that describes what actually happens;
- a qualification named accurately rather than implied;
- a case study with enough context to understand the result;
- an update date where freshness genuinely matters.

It also explains why "just write more" is weak advice. A page containing 2,500 words of vague marketing language may have less usable information than a 900-word page containing specific answers, real examples and clear boundaries.

Google's query fan-out approach adds another wrinkle. A complicated question can lead to several related retrievals rather than one exact-match search. That makes narrow keyword thinking less useful. A page does not need to repeat every possible wording of a question, but it does need enough depth and context to remain relevant when the system explores related parts of the topic.

Think of your page as a knowledgeable member of staff being interviewed. Can it answer the obvious follow-up questions without changing the subject? If it can, you are usually building something stronger for both people and retrieval systems.

## PART TWO - FOUNDATIONS THAT STILL MATTER



### CHAPTER 4

## SEO FOUNDATIONS IN 2026

The arrival of generative search has created a strange market for advice. Some people describe SEO as obsolete at exactly the moment search systems are leaning more heavily on crawling, indexing, relevance, quality and evidence.

The more useful conclusion is the opposite: **SEO is becoming the foundation layer for more kinds of discovery.**

### Relevance still matters

A page has to be about something.

For a local service business, relevance starts with clarity:

- What do you actually do?
- Who is the service for?
- Where is it available?
- What problem does it solve?
- What does the process involve?
- What affects the price or timescale?

This sounds obvious, yet many homepages make visitors decode slogans instead of telling them what the company sells.

"Solutions for tomorrow" might sound polished. "Emergency and planned commercial refrigeration repairs across Cheshire" tells both people and search systems far more.

## Intent matters more than keyword repetition

Keywords remain useful because they reveal language and demand. The mistake is believing that exact repetition creates relevance by itself.

Search systems have become much better at understanding related concepts. Google's own guidance warns against making separate content for every possible query variation merely to manipulate rankings or generative answers.

A strong page should cover the job properly rather than mechanically repeat the phrase.

## Internal links are still underrated

Internal links help users and crawlers discover related pages. They also communicate structure.

A pest control website might sensibly connect:

- wasp treatment -> emergency pest control;
- rat control -> proofing and prevention;
- commercial pest control -> food premises guidance;
- service pages -> genuine location pages;
- guides -> the service they support.

The result feels like a connected body of knowledge rather than a pile of landing pages.

## Links and reputation still count

Search engines still use links and broader signals of prominence and reputation. For local businesses, these can come from trade bodies, suppliers, local press, community organisations, professional directories, partners, event sponsorships and genuinely useful resources that other sites choose to reference.

Buying a package of thousands of irrelevant links is not the modern version of this. It is simply old spam wearing a new invoice.

## Page quality has become harder to fake

Google's people-first guidance asks site owners to consider whether content provides original information or analysis, demonstrates first-hand knowledge and leaves the reader satisfied.

This aligns surprisingly well with GEO thinking. A page that contains first-hand photographs, actual project details, real constraints and useful judgement is harder to commoditise than a generic article assembled from the same five search results everyone else read.

## Conversion remains part of SEO

A ranking that produces no useful action is not a business result.

Every important page should make the next step obvious. That might be:

- call;
- request a quote;
- book;
- send a photo;
- check availability;
- visit the shop;
- compare packages;
- read a case study.

The page should remove anxiety around that step rather than hide it behind a generic "Contact us" button.

## The five foundations to protect

1. **Access** - crawlers and users can reach the content.
2. **Clarity** - the page clearly satisfies a real need.
3. **Structure** - the site and page are logically organised.
4. **Evidence** - important claims are supported.
5. **Action** - the page helps a person decide what to do next.

***Field test:** If someone landed on a core service page with the company logo hidden, could they work out within thirty seconds what is sold, where it is available, what happens next and why the business deserves trust?*

## "Non-commodity" content in plain English

Google's 2026 generative AI guidance uses the idea of non-commodity content. The phrase sounds more complicated than it is. In practical terms, it means publishing something that could not have been produced just as well by somebody who has never done the work.

Suppose five local landscapers all publish "10 tips for a better garden". The articles may be perfectly readable, but there is a good chance they repeat the same general advice. Now compare that with a landscaper who publishes a short case study explaining why a particular Cheshire garden kept flooding, what the soil and levels revealed, what was changed, what the client chose not to do and how the garden looked six months later.

The second piece contains information rooted in experience. That is much harder to commoditise.

Small businesses often underestimate how much original material they already possess. It lives in the questions customers ask before accepting a quote, the mistakes you regularly have to correct, the photographs you take during a job, the awkward exceptions that do not fit a neat marketing slogan, the local conditions you have learned to work around and the judgement involved in deciding when **not** to recommend something.

You do not need to reveal trade secrets or customer information. The point is to turn genuine experience into useful public knowledge.

A simple test is this: **could a competent writer with no access to your business produce this page by reading the first five results**

**on Google?** If the answer is yes, ask what you can add that comes from your own work.

That might be a photograph, a measurement, a before-and-after explanation, a pricing example, a checklist you genuinely use, a local observation or a short note explaining why a common shortcut causes problems.

This type of material is good SEO for a very old-fashioned reason: it gives people a reason to prefer your page. It is also useful in AI-assisted search because it provides distinct evidence rather than another version of the same generic paragraph already available everywhere else.

## CHAPTER 5

# TECHNICAL VISIBILITY: CRAWL, INDEX, RENDER AND DISCOVER

Technical SEO can sound intimidating because the vocabulary arrives before the explanation. For a small business, the goal is simpler than the terminology: **make sure important public pages are reachable, understandable and not accidentally excluded.**

## Start with the status code

When a crawler requests a normal live page, it generally needs a successful response. Broken URLs, redirect loops, server errors and security challenges can prevent discovery.

You do not need to memorise every HTTP code. You do need to know whether your important pages reliably return the page you intended.

## Robots.txt is a gate, not a marketing strategy

The `robots.txt` file tells crawlers which parts of a site they may request. A bad rule can block an entire site or important folders.

Do not edit `robots.txt` simply because an SEO checklist told you to. First understand what is currently allowed and which crawlers matter to you.

Google's AI Search guidance says normal Google Search crawling rules apply to its generative search features. OpenAI says OAI-SearchBot should not be blocked if a publisher wants public content to be available for summaries and citations in ChatGPT search.

## Search crawling and model training are not the same permission

This distinction matters.

OpenAI documents OAI-SearchBot for search discovery and GPTBot for potential training use. A publisher can make different choices about those crawlers.

The wider lesson is that "AI bot" is too vague a category. Different user agents can have different purposes.

## **noindex means what it says**

A `noindex` directive tells a compliant search engine not to include the page in its search index. That can be correct for private, duplicate, thin or utility pages. It is disastrous when accidentally applied to the homepage after a redesign.

Development sites often use `noindex` for good reasons. The danger is forgetting to remove it when the site goes live.

## **Canonicals help manage duplicates**

A canonical tag signals the preferred version when several URLs contain substantially the same content.

Common duplicate patterns include:

- tracking parameters;
- HTTP and HTTPS variants;
- www and non-www versions;
- filtered product pages;
- printer versions;
- copied location pages.

Canonicals are not a licence to create unlimited duplication. They are a clean-up and consolidation tool.

## **Sitemaps help discovery**

An XML sitemap gives search engines a list of important URLs. It does not force indexing, but it is useful for discovery and monitoring.

A sensible small-business sitemap usually contains real pages you want indexed, not every test URL and parameter the CMS ever created.

## JavaScript can work, but complexity has a cost

Google can process JavaScript, but Google's own guidance notes that JavaScript sites can make SEO more complex. A small business should not choose a complicated rendering architecture merely because it sounds modern.

The key information should be accessible and render reliably. Menus, contact details, service information and primary content should not depend on fragile client-side behaviour.

## Important information should exist as text

If the only copy of your opening hours, price list or service description is embedded inside an image, you make the information harder to reuse and less accessible.

Use images to support the message. Put business-critical facts in real text.

## A practical technical checklist

- The important pages return successfully.
- HTTPS works consistently.
- There is one clear canonical version of the domain.
- Robots.txt is not blocking important crawlers unintentionally.
- Important pages are not accidentally noindexed.
- Internal links lead to core services and locations.
- XML sitemap contains the URLs you genuinely want discovered.
- Mobile visitors can use the site without fighting the layout.
- Forms and phone links work.
- Important text is visible without requiring an image or hidden interaction.
- Search Console and Bing Webmaster Tools are verified.

***Do this next:** Run the homepage and your top three service pages through Google Search Console's URL Inspection tool. Confirm that the selected canonical, index status and rendered content match what you expect.*

## A simple diagnostic order when a page is invisible

When a page is not appearing, people often jump straight to rewriting the copy. That can waste hours if the real problem is technical.

Work through visibility in a boring, predictable order.

**First: can the page be reached?** Open it in a private browser window. Check that it returns normally and does not require a login, cookie state, geographic exception or JavaScript challenge just to see the core content.

**Second: is crawling allowed?** Look at `robots.txt`, CDN or firewall rules and any bot-protection service. A site can have a perfectly permissive robots file and still block a crawler at the security layer.

**Third: is indexing allowed?** Check for `noindex`, canonical tags and accidental staging settings. A surprisingly common launch mistake is leaving a development setting in place that tells search engines not to index the site.

**Fourth: can the important content be rendered?** If the service description, prices or links only appear after a fragile script runs, test what search tools actually receive. JavaScript is not automatically bad, but unnecessary complexity creates more places for discovery to fail.

**Fifth: can the page be found through the site?** A page that exists only because somebody knows the URL is weakly connected. Important pages should normally be reachable through sensible internal links.

**Sixth: is it actually worth indexing?** A technically perfect page can still be thin, duplicate or unhelpful. Technical eligibility is only the beginning.

This order matters because it separates access problems from quality problems. Search Console URL Inspection and Bing Webmaster Tools can help diagnose different parts of the chain, but the underlying logic is the same.

For a small site, you do not need a weekly 200-point technical audit. You need confidence that your main service, location, proof and contact pages are reachable, indexable, understandable and connected. Fix those first. The long tail of technical perfection can wait.

## CHAPTER 6

# PAGE EXPERIENCE, ACCESSIBILITY AND TRUST

A search engine can send the perfect visitor to a terrible page.

If that visitor cannot read the text, close the cookie banner, use the form, understand the price or find the phone number, the marketing work has failed.

Page experience is therefore not decorative. It is where visibility becomes business.

## Mobile is the normal case

Many small-business websites are designed on a large monitor and judged on a phone. That is backwards.

Check the pages that matter on an ordinary mobile connection. Look for:

- headings that wrap sensibly;
- buttons large enough to tap;
- no horizontal scrolling;
- phone numbers that can be tapped;
- forms that do not demand unnecessary information;
- images that load at a sensible size;
- menus that remain usable;
- important text that is not hidden by sticky bars.

## Speed matters because patience is finite

You do not need to chase a perfect performance score at the expense of the site. You do need to remove obvious waste.

Large uncompressed images, unnecessary scripts, badly configured fonts, bloated themes and third-party widgets can make a simple site feel slow.

For a local service business, the practical test is not "did we get 100/100?" It is "can a customer reach the information and act without waiting or fighting the page?"

## Trust is often visual before it is factual

Before a visitor reads your detailed proof, they notice whether the page feels current and coherent.

Trust can be damaged by:

- broken images;
- old copyright dates;
- obvious template text;
- dead social links;
- inconsistent business names;
- placeholder testimonials;
- pop-ups covering the content;
- pricing that contradicts another page;
- a contact form that disappears after submission.

A polished design cannot compensate for weak evidence, but poor presentation can make strong evidence look doubtful.

## Reduce uncertainty at the point of action

If the goal is a quote request, explain what happens after the form is sent.

If the goal is a phone call, explain the hours or expected response.

If the goal is a booking, show what information is required and whether payment is taken.

If the job cannot be priced until a survey, say so and explain why.

This is conversion work, but it is also answer-engine work. The page is resolving the questions that stop a person from moving forward.

***Do this next:** Ask someone who has never used the site to complete one task on a phone: find the relevant service, understand the next step and contact the business. Do not help them. Watch where they hesitate.*

## Accessibility is part of usability, not a separate internet

Accessibility is sometimes treated as a compliance task that sits in a different folder from SEO and conversion. In practice, many of the same improvements make a website easier for everybody to use.

A properly labelled form field helps a screen-reader user, but it also makes the purpose of the field clearer to browser automation. A real button is easier to operate with a keyboard and easier for an agent to identify than a decorative element with a click handler. A sensible heading structure helps a reader skim, helps assistive technology navigate and gives automated systems a clearer map of the page.

This does not mean "add ARIA everywhere". Bad ARIA can make a page worse. The first choice should usually be good native HTML: real headings, real buttons, real form labels, real links and a logical document order.

For a local business, the commercial benefits are wonderfully unglamorous:

- fewer people abandon a form because the labels make no sense;
- fewer mobile users miss the phone number or call-to-action;
- keyboard users can complete the same journey as mouse users;
- text remains readable when somebody zooms the page;
- automated tools have a better chance of understanding what each control does.

There is also a trust point. A website that behaves predictably feels more professional. Customers notice when a form loses their input, a menu cannot be closed, the text jumps around during loading or a

button looks clickable but is not. They may never use the phrase "page experience", but they feel the friction.

The best test is to try the main journey yourself without ideal conditions. Use a phone. Zoom the text. Tab through the form. Ask somebody who did not build the site to find a service, understand the price position and send an enquiry.

If the journey works under slightly awkward conditions, you have probably improved more than accessibility alone.

### A small-business accessibility baseline

You do not need to turn a five-page local business website into an accessibility research project. You do need a sensible baseline.

WCAG 2.2 is the current W3C recommendation and it is useful even when a business is not trying to make a formal conformance claim. It gives you a way to think about ordinary problems that are easy to miss when you know the site too well.

Start with the main customer journeys rather than trying to audit every decorative detail at once.

Can somebody use the navigation and contact form with a keyboard? Can they see which control currently has focus? Can they enlarge the page without the layout becoming unusable? Are buttons big enough to tap without surgical precision on a phone? Are form errors explained in words rather than shown only by a red outline? Does the same help or contact route stay reasonably consistent across the site?

These sound like accessibility questions because they are. They are also basic usability questions.

A useful small-business baseline is:

- every important image has appropriate alternative text when the image conveys information;

- decorative images do not create pointless noise for screen-reader users;
- headings follow a meaningful order rather than being chosen only because one size looks nice;
- links say where they go instead of relying on a long row of "click here" labels;
- buttons are actual buttons when they perform actions;
- form fields have visible labels;
- keyboard focus is visible and is not hidden behind a sticky header or cookie banner;
- colour contrast is strong enough that text is not dependent on perfect eyesight or a perfect screen;
- touch targets are not packed tightly together;
- error messages say what went wrong and what the person can do next;
- important instructions are not available only inside an image;
- authentication and checkout steps do not rely on puzzles that unnecessarily exclude people.

The point is not to treat WCAG as an SEO checklist. That would miss the reason it exists. Accessibility is about people being able to use the web.

The strategic overlap is simply a bonus. Clear labels, meaningful controls, predictable page structure and understandable error states also give automated systems less ambiguity to deal with.

Accessible design also exposes weak processes. A twelve-field form, icon-only navigation or vague "Continue" button often looks less clever once you ask whether somebody can use it without extra context.

Accessibility work can improve conversion without changing a ranking. Automated scans can flag obvious issues, but they do not prove full accessibility or support legal claims. For most owner-run businesses, start by making key journeys clear, keyboard-usable, readable, labelled and predictable.

## PART THREE - BECOMING THE ANSWER AND THE SOURCE



### CHAPTER 7

## AEO: ANSWERING THE QUESTION CLEARLY

Answer Engine Optimisation sounds technical. In practice, it is often a discipline of being less vague.

Customers ask questions because they are trying to reduce uncertainty. They want to know whether a service fits their situation, how much it might cost, what happens next and what could go wrong.

A page that answers those questions well is useful in ordinary search, snippets, voice-style results and AI-assisted answers.

### Start with the real question, not the keyword

A keyword such as `boiler service Cheshire` tells you very little about the customer's concern.

The questions behind it might be:

- How often should a boiler be serviced?
- What does a service actually include?
- How long does it take?
- Do I need to be at home?
- Will the engineer turn the heating off?
- Can a service find a carbon monoxide problem?
- What is the difference between a service and a repair?
- Does the price include parts?

Those questions are useful because each one can be answered directly.

## Answer first, explain second

Many business websites hide the answer behind an introduction.

### **Weak:**

"At ABC Services we pride ourselves on offering a comprehensive range of solutions tailored to the needs of our valued clients."

### **Better:**

"A standard boiler service usually takes around 30 to 60 minutes, depending on the appliance and whether the engineer finds anything that needs further investigation."

The second version gives the visitor something immediately. You can then explain the process, caveats and booking details.

That answer-first pattern is especially useful for:

- price ranges;
- opening hours;
- service coverage;
- eligibility;
- timescales;
- definitions;
- comparisons;
- troubleshooting steps.

## Use structure when structure helps

Lists, tables and headings are not ranking tricks. They are reading tools.

A comparison between two service packages is usually clearer in a table. A six-step process is usually clearer as six steps. A page answering ten unrelated customer questions is easier to scan when each question has its own heading.

The purpose is not to chop every paragraph into tiny "AI chunks". Google's 2026 guidance explicitly warns against over-focusing on

special chunking tactics for generative search. The purpose is to make the page easy to follow.

## A useful answer has boundaries

Good answers include exceptions.

"We can normally attend the same day" is not the same as "we guarantee same-day attendance".

"Most jobs cost between £X and £Y" is not the same as "your job will cost £X".

"This treatment is suitable for many properties" is not the same as "this is definitely the right treatment for your property".

Boundaries make an answer more trustworthy because they show where judgement is still required.

## Build answer layers

A strong page often works in layers:

1. **Direct answer** - one or two sentences.
2. **Context** - what changes the answer.
3. **Evidence** - examples, photographs, measurements, credentials or source links.
4. **Process** - what the customer should expect.
5. **Action** - what to do next.

This structure works well because different visitors need different depths.

Someone in a hurry gets the answer. Someone comparing suppliers can read the evidence. Someone ready to buy gets the next step.

## AEO is not an FAQ factory

One of the worst reactions to answer engines is to add hundreds of low-value questions purely because a tool suggested them.

The result often looks like this:

- Do plumbers fix taps?
- Can plumbers fix taps?
- Will a plumber fix my tap?
- Is tap fixing something plumbers do?

That is not coverage. It is duplication.

Group overlapping questions into useful sections and answer the underlying need once.

## Where to find the best questions

Your best question research may already exist in:

- emails;
- WhatsApp conversations;
- quotation notes;
- call logs;
- live chat;
- reviews;
- sales meetings;
- support tickets;
- conversations with staff;
- objections that repeatedly delay a sale.

Keyword tools are useful, but real customer language is often better.

## Example: a local web designer

Instead of a page that only says "web design Cheshire", a useful page might answer:

- How much does a small-business website cost?
- Can you take over an existing site?
- Will I own the domain?
- What happens to my current Google rankings?
- Is hosting included?
- Can I update the site myself?
- What happens if the person who built my site disappears?
- How long does a rebuild take?

Those are buying questions. Answer them properly and the page becomes useful before and after the click.

***Do this next:** Choose one core service page. Add three customer questions that genuinely affect the buying decision. Put the direct answer in the first sentence of each section, then add the nuance.*

## The answer-first pattern without making every page sound robotic

There is a risk in AEO advice: every page starts to read like it was written by the same template.

**Question. Answer. Three bullets. Next question.**

That can be useful for a support page, but a business website still needs a human voice. The trick is to be clear without becoming mechanical.

Take a common question such as "How much does a new website cost?" A weak answer hides behind a sales phrase: "Every project is unique, contact us for a bespoke quote." It is technically true, but it does not help the reader.

A better answer might say that a straightforward brochure site can cost a few hundred pounds while larger builds involving ecommerce, booking, integrations, copywriting or complex migrations cost more. It can then explain the factors that move the price and give the business's own starting point if one exists.

That is an answer. It is also a conversation.

The same principle works in other sectors. A garage can explain that a warning light needs diagnosis before a repair price is meaningful. An accountant can explain what is normally included in a monthly package and what tends to be charged separately. A pest controller can distinguish a single treatment from a longer control programme. A builder can explain why access, ground conditions and specification change a quote.

Useful answers often include a boundary. "Usually", "from", "in most cases", "subject to inspection" and "depending on" are not signs of weakness when they reflect reality. They stop the content making promises the business cannot keep.

One good writing habit is to imagine the customer has interrupted you mid-sentence and said, "Yes, but what does that mean for me?" If the next paragraph answers that, the page is probably moving in the right direction.

### **A practical answer framework: direct, useful, reassuring, specific**

When a customer asks a question, there are usually several things happening underneath the words. They want the factual answer, but they may also be trying to judge cost, risk, effort and whether they can trust the person giving it.

That is why the shortest possible answer is not always the best answer. AEO should not turn a website into a collection of one-line

definitions. The aim is to answer quickly **and then help the person make sense of the answer.**

A useful pattern for small-business pages is:

### **1. Give the direct answer.**

Do not make the reader hunt through three paragraphs of company history before finding out whether you cover their area, whether a service can be booked online or whether a job usually needs an inspection first.

### **2. Explain what changes the answer.**

Real businesses deal in variables. A roofer may need to know access, roof type and the extent of damage. An accountant may need turnover, VAT status and payroll requirements. A web designer may need the number of pages, integrations and whether existing content can be reused. Explain those variables plainly.

### **3. Add the practical consequence.**

Tell the customer what the answer means for them. If an electrician cannot price a fault accurately without testing, explain what the diagnostic visit covers. If a booking deposit is required, say when it is taken and what happens if the date changes.

### **4. Deal with the obvious concern.**

The person reading may be thinking, "Will this turn into a much bigger bill?", "Will I be without my car all day?", "Do I need to prepare anything?" or "Am I tied into a contract?" If the concern comes up repeatedly in real conversations, it belongs on the page.

### **5. Make the next step clear.**

A good answer should naturally lead somewhere: check availability, send photographs, request a quote, book an assessment, call with a registration number, or read the more detailed guide.

For example, a garage answering "**Can you tell me why my engine light is on?**" could say:

*Sometimes, but the warning light on its own does not identify the failed part. We normally begin with a diagnostic scan and, where needed, a few checks to confirm the cause before recommending a repair. If the light is flashing or the car is running badly, contact us before continuing to drive it.*

That answer is useful because it does not pretend the diagnosis is simpler than it is. It gives the customer a realistic next step and a sensible safety boundary.

The same approach works for service pages, FAQs, blog posts and location pages. It also has a useful side effect: the writing sounds more like a competent person explaining the job and less like content created to satisfy a keyword tool.

### A note on FAQs after 2026

Google stopped showing the FAQ rich result in Search in May 2026. That does **not** make customer questions useless. It simply removes one old reason people chased FAQ markup. Keep questions when they genuinely help the reader, not because you expect a decorative search feature.

## CHAPTER 8

## GEO: BECOMING SOURCE-WORTHY

If AEO asks "can this page answer the question?", GEO asks a harder question:

**Would this page be a good source for an answer?**

That shift moves the focus away from wording tricks and towards evidence.

### GEO began as a research problem

The term Generative Engine Optimisation was formalised in academic research published in 2023. The researchers tested ways of changing content to improve its visibility in generative engine responses. Their controlled benchmark reported substantial gains for some strategies, but the results varied by domain.

That last point matters.

There is no universal formula such as "add three statistics and one quotation, then AI will cite you". Real systems change, use different retrieval methods and evaluate different types of evidence.

The useful legacy of the research is the question it raised: **visibility in a generated answer is a different outcome from ranking in a list of links.**

By 2026, Microsoft was using GEO as a practical industry term when discussing citation and grounding data in Bing Webmaster Tools.

### Source-worthiness is not the same as length

A long article can still be useless.

A short page can be exceptionally valuable if it contains information nobody else has.

Examples include:

- a transparent price table;
- original measurements;
- before-and-after project details;
- a local availability schedule;
- a clear explanation of a specialised process;
- first-hand testing;
- a documented case study;
- a dataset;
- a genuine expert judgement with reasoning;
- photographs that support the claim being made.

The common feature is not word count. It is information value.

## Evidence beats adjectives

Compare these two claims:

**Claim:** "We are highly experienced in heritage roof repairs."

**Evidence:** "In the last 12 months we completed six slate roof repairs on pre-1900 properties. This case study shows the original failure, the replacement slate used, how the flashing was rebuilt and photographs before, during and after the work."

The second version gives a customer more confidence and gives a retrieval system more concrete material.

## Entity clarity matters

An AI system trying to answer a question about a business needs to understand what that business is.

Your site should make it easy to identify:

- business name;
- trading name if different;
- location or genuine service area;
- primary services;
- contact information;
- people behind the business where appropriate;
- credentials that can be verified;
- links to relevant profiles;
- consistent branding and naming.

This is not about publishing every personal detail. It is about reducing ambiguity.

## Freshness matters differently in generated answers

A stale page in classic search is annoying. A stale fact used inside a generated answer can become misleading.

Microsoft has highlighted freshness as especially important for grounding and promotes IndexNow as one way publishers can notify participating search engines quickly when content changes.

For a small business, you do not need to obsess over publishing dates on every page. You do need to keep facts current:

- opening hours;
- prices;
- services;
- staff;
- service areas;
- product availability;
- legal or compliance information;
- offers and seasonal conditions.

## Citation is not the same as influence

A page can be cited and still contribute very little to an answer.

A 2026 academic paper proposed separating \*citation selection\* from \*citation absorption\*. In its dataset, pages that had stronger influence on answers tended to be more structured, semantically aligned and rich in extractable evidence such as definitions, numbers, comparisons and procedures.

This is interesting, but it remains early research. Do not turn it into another magic checklist.

The practical lesson is safer: **give systems something precise to use.**

## Make your pages easy to verify

If a claim matters, ask how a reader could verify it.

- Qualification -> name the qualification accurately.
- Membership -> link to the organisation if appropriate.
- Price -> explain what is included and what changes it.
- Case study -> show the work.
- Coverage -> list the real areas.
- Availability -> date the information if it changes.
- Statistic -> cite the source.

## The GEO test

For any important page, ask:

1. What unique information is here?
2. Which claims are supported by evidence?
3. Is the business or author clearly identified?
4. Are facts current?
5. Could another website copy this page in five minutes and lose nothing important?

If the answer to number five is yes, the page probably needs more first-hand value.

***Do this next:** Add one source-worthy asset to a core page: a real example, price range, process diagram, original photograph, measurement, checklist or case study detail.*

## Citation selection is only half the story

A useful 2026 research paper proposed separating two ideas that are often bundled together: **citation selection** and **citation absorption**.

Selection asks whether an AI search system chooses a page as a source. Absorption asks whether information from that page meaningfully contributes to the final answer.

For a business owner, the distinction is more practical than academic.

Imagine your guide is listed as one of several sources below an answer but none of its distinctive information appears in the response. That is still visibility, but it is different from a page whose data, explanation or comparison materially shapes what the system says.

You cannot reliably measure absorption across every platform from an ordinary analytics account, and the research is still early. But the idea gives us a better standard for writing.

Do not ask only, "How do I get cited?" Ask, "What is on this page that would actually help somebody answer the question?"

That pushes the work towards explicit definitions where the term is genuinely confusing, numerical facts with context, comparisons that state the basis of comparison, procedures that reflect real practice, original observations, evidence that can be checked and limitations that prevent overclaiming.

It also explains why a citation counter should never become the new version of a keyword-ranking obsession. A page can be cited for the wrong topic, cited in an answer nobody important sees or cited without producing any commercial effect.

The stronger long-term goal is **useful influence**. If your page helps a customer understand the problem correctly, narrows the options and makes your business easier to evaluate, it is doing valuable work whether the final journey contains one click, three clicks or no immediately measurable click at all.

## CHAPTER 9

## EVIDENCE, EXPERTISE AND ORIGINALITY

The internet has never been short of confident sentences.

"We pride ourselves on excellence."

"Customer satisfaction is our number one priority."

"We are a leading provider of innovative solutions."

These statements may be true. The problem is that they are almost impossible to distinguish from the same claims on ten thousand other websites.

### Replace self-description with proof

Whenever a page makes an important claim, ask what could support it.

**"We are experienced."**

Show years in the trade, representative projects, difficult jobs, lessons learned or the range of problems handled.

**"We cover Cheshire."**

State the actual service area and any limits.

**"Our prices are competitive."**

Give starting prices, example packages or explain what changes the quote.

**"We respond quickly."**

Explain normal response windows honestly and distinguish emergencies from planned work.

**"Customers trust us."**

Show genuine reviews, case studies and repeat-client examples.

The difference is not stylistic. Proof lowers risk for the buyer.

## First-hand knowledge is a competitive asset

Google's people-first guidance explicitly encourages content that demonstrates first-hand expertise and original information.

For a small business, that can be easier than it sounds because you already know things generic publishers do not.

A drainage contractor knows which problems repeatedly occur in certain property types. A web designer knows which mistakes cause small-business owners to lose control of domains. A mechanic knows which symptoms frequently point to a particular failure. A florist knows which wedding flowers cope badly with summer heat.

Those observations are content.

## Write the things competitors cannot safely invent

The best defence against commodity content is specificity.

Useful first-hand details include:

- why you chose one method over another;
- what went wrong on a job and how it was fixed;
- what customers commonly misunderstand;
- how you decide whether a job is suitable;
- what is included in a quote;
- what causes delays;
- what preparation helps;
- what signs mean a customer should not wait;
- what you would do differently next time.

A generic AI system can produce a plausible paragraph about almost any industry. It cannot truthfully invent your real experience.

## Original does not mean never said before

You do not need to discover a new law of physics.

A useful page can still cover common knowledge. Originality often comes from:

- your explanation;
- your examples;
- your local context;
- your photographs;
- your process;
- your data;
- your judgement.

A page about "how long a website rebuild takes" can be original if it explains your actual stages, typical delays and what a client can do to speed the project up.

## Be careful with credentials and regulated claims

Trust-building becomes risky when businesses exaggerate.

Do not imply accreditation you do not hold. Do not use logos without permission. Do not invent years of experience by adding together unrelated staff histories. Do not claim insurance, licensing or membership without checking the current status.

If an accreditation matters, name it precisely.

That accuracy helps people and machines distinguish evidence from marketing.

## Case studies are one of the strongest formats

A useful case study can include:

1. The situation.
2. The problem.
3. Constraints.
4. The decision.
5. The work completed.
6. The result.
7. What was learned.

Do not turn every case study into a victory parade. If there was a constraint or trade-off, explain it. Credibility often increases when the page sounds like a real job rather than a brochure.

## Original photography has unusual value

Real images can prove presence, process and workmanship.

For a local business, a real photograph of the team, equipment, premises, project or completed work can do more for trust than a polished stock image.

Add useful alt text where appropriate, compress images sensibly and make sure the written page still contains the key facts.

## Evidence creates a compounding effect

Strong evidence can support:

- better conversion;
- richer service pages;
- useful case studies;
- local relevance;
- links and mentions;
- review responses;
- social proof;
- AI retrieval and citation.

This is why "show, do not just tell" is more than a copywriting rule. It is a search visibility strategy.

***Do this next:*** Highlight every major claim on one important page. For each claim, add either evidence, a qualifier or a more accurate wording.

## Build an evidence ledger before you build more content

Most businesses have more proof than they realise, but it is scattered.

One team member has photographs on a phone. Another has a folder of completed quotations. Reviews sit on Google. Qualifications are in a drawer. A useful explanation is buried in an email somebody sends to every new customer. The website then says, "trusted experts delivering quality service" because that is all the person writing it can see.

An evidence ledger is simply a list of the things the business can legitimately prove.

For each important service, record real projects, permitted photographs, accurate measurements or outcomes, relevant qualifications or memberships, genuine reviews, problems solved, useful local experience and occasions when you advised a customer not to buy.

Then connect evidence to claims.

If you say you work across Cheshire, show real areas served. If the team is experienced, show what that experience enables. If the process is thorough, explain it. If customers trust you, use genuine reviews rather than an invented statistic.

The exercise also protects the business by exposing unsupported claims. Correct them even if search engines did not exist: evidence is not SEO decoration; it is why customers should believe you.

## CHAPTER 10

# INFORMATION ARCHITECTURE FOR HUMANS AND MACHINES

Information architecture is the way a website organises its knowledge.

For a small business, this sounds grander than it is. It means deciding which pages deserve to exist, what belongs on each page and how visitors move between them.

Good architecture reduces three kinds of confusion:

- customer confusion;
- crawler confusion;
- content-team confusion.

## Give each important page a clear job

A common website problem is one giant services page containing six unrelated offers and no depth.

The opposite problem is creating a separate page for every tiny keyword variation.

A sensible middle ground is to create pages around genuine customer decisions.

For example, a pest-control business might have:

- Home
- Residential pest control
- Commercial pest control
- Wasp treatment
- Rodent control
- Bed bug treatment
- Service areas
- Case studies

- Advice hub
- About
- Contact

The exact structure depends on the business, but each page should have a reason to exist.

## Think in topic clusters without making them robotic

A service page can be supported by useful guides.

**Core page:** Wasp nest treatment

### Supporting guides:

- How to tell whether you have wasps or bees
- What happens during a wasp treatment
- Can you remove a nest from a loft?
- When is a wasp nest likely to become inactive?
- What should you do before the technician arrives?

These guides answer genuine questions and naturally link back to the service.

## Internal links should make sense in the sentence

Avoid stuffing pages with blocks of keyword-rich links purely for SEO.

Link where the next page is genuinely useful.

A sentence such as "If you are unsure whether the insects are wasps or bees, see our identification guide" is better for a reader than a footer containing thirty town-and-service combinations.

## Build around decisions, not search volume alone

Some of the most useful pages may have low search volume.

A page explaining what happens after a quote may remove a major sales objection. A guide to domain ownership may prevent a client

from making an expensive mistake. A page explaining a specialist service may attract only a small audience but convert strongly.

Search volume is a demand signal, not the only measure of usefulness.

## Use descriptive headings

Headings should help a person scan the page.

Weak:

- Our approach
- More information
- Why us

Stronger:

- What is included in the survey?
- What affects the final quote?
- How long does installation normally take?
- What happens if we find additional damage?

Descriptive headings improve comprehension without needing to sound unnatural.

## Avoid burying facts in accordions

Collapsible sections can make long pages easier to use, but critical information should not become difficult to access.

If opening hours, legal restrictions, service exclusions or essential pricing information matters to the buying decision, make it easy to find.

## Navigation is a promise about priorities

The main menu tells visitors what the business considers important.

If a key service needs four clicks to find, the site is signalling that it is not central.

A small-business navigation should normally prioritise:

- core services;
- proof or projects;
- locations if genuinely relevant;
- about/trust;
- contact or booking.

## Create a content map before creating more content

Before adding another article, map what already exists.

For each URL, record:

- purpose;
- audience;
- primary topic;
- related service;
- status;
- whether it overlaps another page;
- whether it is still accurate;
- next action.

This simple inventory often reveals duplicate content and missing buyer questions faster than a keyword tool.

***Do this next:*** Draw your website as boxes on one sheet of paper. If you cannot explain the relationship between the boxes, neither can a new customer.

## Design the site around decisions, not your filing cabinet

Business owners naturally organise information according to how the company works internally. Customers rarely think that way.

A firm might have departments called Operations, Commercial, Maintenance and Projects. A customer thinks, "Can you repair my

leaking roof?", "Do you cover Northwich?" or "How much will this cost?"

Good information architecture translates the internal business into the customer's decision-making language.

Start by listing the handful of decisions a visitor must make:

1. Am I in the right place?
2. Do they provide the service I need?
3. Do they work in my area or situation?
4. Can I trust them?
5. What is likely to happen and cost?
6. How do I take the next step?

Your main navigation, service pages and internal links should make those decisions easier.

This is particularly useful in AI-assisted search because retrieval systems may land on a deep page rather than the home page. A service page should therefore make enough sense on its own. It should identify the service, the business, the relevant location or audience and the next useful action without assuming the visitor has read three other pages first.

At the same time, avoid turning every page into a complete copy of the website. Internal links exist so pages can share the job. A guide about damp in a roof can explain the issue and point to the repair service. The service page can explain the commercial offer and point to a case study. The case study can show proof and link back to the relevant service.

That network is better than stuffing everything into one enormous page.

A good site should guide people through the options. Visitors should always know where they are, why the information matters and what to do next.

## CHAPTER 11

## STRUCTURED DATA WITHOUT THE HYPE

Structured data is useful. It is also one of the most over-sold areas of modern search.

At its simplest, structured data is machine-readable markup that describes information already present on a page.

For example, `LocalBusiness` markup can identify a business name, address, telephone number, opening hours and other details in a standard format.

### What structured data can do

Correct structured data can help search engines interpret content and can make pages eligible for certain rich-result features.

For local businesses, useful types can include:

- `LocalBusiness` or a more specific subtype;
- `Organization`;
- `Product` where relevant;
- `Event` where relevant;
- `Article` for editorial content;
- `BreadcrumbList`;
- other Google-supported types that genuinely match the page.

The important phrase is **genuinely match the page**.

## What structured data cannot do

It cannot:

- guarantee a ranking;
- guarantee a rich result;
- make a weak page authoritative;
- create a fake review legitimately;
- tell an AI system to recommend your business;
- compensate for contradictory visible content.

Google explicitly says there is no special schema.org markup required for AI Overviews or AI Mode.

Its 2026 optimisation guidance also says businesses can ignore "special" AEO/GEO markup tricks that are not part of normal Search requirements.

## Visible content must match the markup

This is one of the safest rules in structured data.

If your markup says the shop closes at 6pm but the page says 5pm, you have created ambiguity.

If your schema says you have an aggregate rating that is not actually visible or valid under Google's policies, you create a compliance problem.

If your markup describes services the business does not actually offer, it is misleading.

Structured data should describe the page, not invent a better version of the business.

## Use the most specific sensible type

Google recommends using the most specific subtype that accurately fits the organisation.

For a physical local business, that may mean a specific `LocalBusiness` subtype rather than a generic organisation label.

Do not force a type because it sounds impressive. Accuracy is more useful than ambition.

## Local business details worth keeping aligned

Check consistency across:

- website;
- structured data;
- Google Business Profile;
- Bing Places;
- major directories;
- social profiles;
- booking systems;
- contact pages.

The goal is not to make every platform identical in design. It is to prevent factual contradictions.

## Beware fake review markup

Google has specific review-snippet rules. A local business cannot simply mark up self-controlled reviews on its own site and expect star snippets for itself.

This is a good example of why schema should not be implemented from a random blog post without checking current official documentation.

## Validate what you publish

Useful checks include:

- Google's Rich Results Test;
- Search Console enhancement reports where applicable;
- schema.org validator for general markup quality;
- manual comparison between markup and visible content.

Validation proves that the markup is syntactically understandable. It does not prove Google will show a feature.

***Do this next:** If your site uses structured data, compare the business name, telephone, address/service area, opening hours and main entity details against the visible page. Fix contradictions before adding more markup.*

## Treat structured data like maintained business data

The biggest structured-data mistake is not usually syntax. It is neglect.

A developer adds markup at launch, everybody forgets about it, and eighteen months later the website has new opening hours, a different telephone number or a renamed service while the structured data still describes the old business.

That is the opposite of what structured data is meant to achieve.

For a small local company, keep the job modest. Identify the types that genuinely apply, populate them from reliable business information and review them when the visible page changes.

A quarterly check can be as simple as:

- does the organisation name match the current brand?
- is the preferred logo URL still valid?
- are telephone and address details correct?
- are opening hours current where they are marked up?
- does the structured business type still describe what the company actually does?
- do breadcrumbs match the visible site structure?
- does every marked-up claim have a visible equivalent a customer can see?

Do not invent structured information because a testing tool says a property is recommended. A blank or inapplicable field is better than fabricated data.

It is also worth separating two ideas. Structured data can help a search engine understand information and can make a page eligible for supported rich features. That is valuable. But Google explicitly says there is no special structured data required to appear in its generative AI search features. Adding another schema type is therefore not a shortcut around weak content.

Good markup describes a good page. It does not rescue a bad one.

## Entity consistency beats markup volume

Structured data becomes much easier once you stop thinking of it as a bag of SEO tags and start thinking of it as one representation of the business.

Imagine a company that changed its trading name two years ago. The homepage uses the new name. The footer still contains the old name. The Facebook page has a third variation. One directory lists a previous telephone number and the JSON-LD still points at an old logo.

Adding more schema to that mess does not create clarity. It gives machines another version of the disagreement.

The stronger approach is to decide what the current business facts actually are and make those facts consistent wherever they are deliberately published.

For an ordinary local business, the core entity record might include:

- current trading name;
- main website URL;
- preferred logo;
- telephone number;
- public email address;
- physical address or accurate service-area model;
- opening hours where relevant;
- main business category;
- genuine social and profile links;
- parent or legal organisation details where they need to be stated.

Keep that record in a simple internal document or spreadsheet so the website, Business Profile, Bing Places, email signatures and future redesigns all start from the same source.

This matters especially after a move, merger, rebrand or new location, when old details can linger in PDFs, directories, staff bios and cached pages.

Google's current Organization guidance encourages useful real-world and online identity details such as name, address, telephone, URL and logo where relevant. Its LocalBusiness documentation similarly recommends using the most specific accurate business type.

The practical lesson is not "add every possible property". It is "make the properties you do publish dependable".

If a machine sees three different phone numbers, the problem is not a shortage of markup. If a customer sees them, the problem is bigger.

Treat entity consistency as maintenance. Re-check it after a rebrand, move, new phone system, acquisition, opening-hours change or major redesign - the moments when stale facts tend to survive.

## PART FOUR - THE LOCAL BUSINESS PLAYBOOK



### CHAPTER 12

## LOCAL SEARCH IN AN AI-SEARCH WORLD

Local search has always been a little different from general web search because the answer depends on the real world.

A useful result needs to know not only what a business does, but where it is, whether it is open, how far away it is, whether it serves the requested area and how prominent or trusted it appears.

Google describes local ranking mainly in terms of **relevance**, **distance and prominence**. Those three ideas remain useful even as AI is layered into the journey.

### Relevance: does the business actually fit the request?

A complete, accurate Google Business Profile helps Google understand what the business is and match it to relevant searches.

That means choosing the right categories, describing the real services and keeping the information current.

Do not add services you do not provide just because they have search volume.

A roofer who does not install solar panels should not pretend to be a solar installer to chase a broader query.

### Distance: local still means local

Distance cannot be optimised away.

If the customer wants a nearby café, a café thirty miles away is not suddenly the best result because it has better SEO.

For service-area businesses, geography is more complicated because the business travels to the customer. The website and profiles should make the genuine coverage clear without pretending to have offices in towns where no office exists.

### **Prominence: reputation travels across the web**

Google says prominence can be influenced by things such as links and reviews.

For a real local business, prominence is often built through ordinary activity:

- local press coverage;
- chamber or trade membership;
- supplier listings;
- sponsorships;
- community involvement;
- relevant directories;
- reviews;
- links from partners;
- customers mentioning the business online.

This is healthier than trying to manufacture hundreds of low-quality citations.

## Business Profiles are part of your information layer

A Business Profile is not merely a map pin.

It can contain:

- category;
- telephone number;
- opening hours;
- website;
- address or service area;
- photos;
- reviews;
- products or services;
- booking links;
- updates.

If those details contradict the website, uncertainty grows.

## Local AI answers raise the cost of inconsistency

Imagine a user asks:

"Who can repair a leaking flat roof around Middlewich this afternoon?"

A system may need to combine:

- service relevance;
- location;
- opening information;
- website content;
- profile data;
- reviews;
- possibly availability language.

If the website says the business covers Middlewich but the Business Profile lists a different region and the contact page still shows an old telephone number, the data is weaker than it needs to be.

Microsoft has explicitly advised local businesses to keep Bing Places information accurate for AI visibility. Google's AI guidance likewise tells site owners to keep Business Profile information up to date.

## Reviews answer questions that the business cannot answer about itself

A business can say it is reliable. Customers provide independent evidence.

Reviews can reveal themes such as:

- punctuality;
- communication;
- cleanliness;
- value;
- ability to handle unusual jobs;
- follow-up service;
- friendliness;
- speed.

Do not script reviews for customers. A repeated pattern of natural comments is more persuasive than a wall of identical praise.

## Photos help establish reality

For local businesses, photographs can show:

- premises;
- vehicles;
- team;
- completed work;
- products;
- local context.

Use real images when possible. They make a business feel present rather than theoretical.

## Your local consistency check

Check that the following agree wherever they appear:

- business name;
- main telephone number;
- address or service-area model;
- opening hours;
- core services;
- website URL;
- booking link;
- major credentials;
- business description.

Do not obsess over punctuation differences. Focus on factual consistency.

***Do this next:** Open your website, Google Business Profile and Bing Places side by side. Compare the details as though you were a customer trying to decide whether all three refer to the same active business.*

## A local-search scenario: what the customer is really asking

Consider somebody in Northwich asking an AI-assisted search tool:

*"I need a roofer to look at a leak around a chimney. Who works around Northwich, and what should I ask before booking?"*

That one question contains several separate needs. The person wants a local provider, but they also want reassurance about suitability, process and how to avoid making a poor choice.

A useful search or AI response may need to establish:

- which businesses genuinely serve Northwich;
- whether they actually handle roof leaks and chimney-related work;
- whether there is evidence of relevant jobs rather than a generic "all roofing undertaken" claim;
- how the customer can make contact;
- whether an inspection or call-out is normally required;
- and what information the customer should have ready.

No single trick supplies all of that. The answer may depend on a mixture of the business website, local profiles, reviews, independent mentions and current factual information.

This is where small businesses sometimes underestimate the value of ordinary operational detail. A page that says "**Roofing services in Northwich**" and little else gives a retrieval system very little to work with. A page that explains the actual work carried out in the area, common property types encountered, how inspections are arranged, the surrounding places covered and what affects urgency is far more useful.

The point is not to manufacture local trivia. Nobody needs a paragraph about the history of Northwich inserted into a roof-repair page merely to prove that the town exists. Useful local detail comes from doing the work.

For a service-area business, that might include:

- whether travel or minimum-job rules change at the edge of the area;
- realistic attendance arrangements rather than guaranteed response times;
- photographs from completed work where permission allows;
- common access or parking issues that genuinely affect the job;
- the nearest areas that are normally served together;
- and clear wording about where the business is **based** versus where it **works**.

There is also a customer-experience reason to get this right. A visitor should never need to phone simply to discover whether the business covers their postcode if that information can reasonably be published.

Good local visibility is not about pretending to be everywhere. It is about making the areas you genuinely serve easy to verify and easy to understand.

### **Worked example: a roofer serving Northwich without inventing a Northwich office**

This is an illustrative example, not a client case study. Imagine a roofing company that is based outside Northwich but regularly works there. Its old website has one sentence saying it covers Cheshire, a list of twenty towns in the footer and a Google Business Profile with the correct base address. The information is not exactly wrong, but it makes a customer do too much detective work.

A customer in Northwich with water coming through a bedroom ceiling does not really care whether the business has used the phrase "roofer Northwich" often enough. They want to know whether the company actually comes to Northwich, whether roof leaks are the sort of job it handles, what happens next and whether there is a reason to trust it.

The stronger version starts by fixing the source of truth. The website states where the business is based and names Northwich as a genuine service area. The Business Profile keeps the real address and current hours. Old directory entries with obsolete details are corrected where practical.

The Northwich service information then earns its place by being useful. It explains the types of repairs normally undertaken, whether an inspection is usually needed before a firm price can be given, what photographs are helpful when making an enquiry and what can

make a repair more complicated. If the business has permission to show completed work in the area, one or two real examples add far more value than a paragraph about Northwich landmarks.

A useful page might also say that emergency attendance is subject to availability rather than promising a response time the business cannot always meet. That small sentence matters. It is more honest for the customer and gives search systems a more accurate fact to work with.

Nothing in this example guarantees a ranking or an AI recommendation. The improvement is more basic than that: the business has become easier to verify. Its website, local profile and evidence now tell the same story. If a search system needs to establish who genuinely serves Northwich and deals with roof leaks, there is clearer information available. If a customer lands on the page, there is less uncertainty before the call.

That is what good local optimisation often looks like in practice. It is less about pretending to be more local than you are and more about removing ambiguity from the areas you genuinely serve.

## The local information triangle

For a local business, three information sources should tell the same basic story:

- 1. Your website** - services, areas, contact details, hours, proof and policies.
- 2. Your major business profiles** - particularly Google Business Profile and Bing Places where relevant.
- 3. The wider web** - reputable directories, local organisations, review platforms and mentions that describe the business independently.

You do not control every third-party source, but obvious contradictions are worth fixing.

Suppose the website says a business closes at 5pm, Google says 6pm and an old directory says the firm has moved. A customer sees inconvenience. A retrieval system sees uncertainty. Neither is helpful.

This is why local SEO is partly an information-management job.

Google continues to describe local ranking mainly through relevance, distance and prominence. Distance remains a factor you cannot simply optimise away. If a customer is standing ten miles from your business, a competitor two streets away may simply be a more natural local result. What you can control is whether Google has enough accurate information to understand relevance and whether the business has genuine prominence through reviews, links, mentions and real-world recognition.

AI answers do not remove those realities. In some ways they make accuracy more important because a generated response may combine hours, location, service availability and reputation into one recommendation-style answer.

For service-area businesses, be especially careful with language. "Based in Winsford and serving Northwich, Middlewich and surrounding Cheshire areas" is clearer than pretending to have an office in every town. Truthful geography is not a disadvantage. It gives the customer a more reliable expectation and avoids building a website around locations the business cannot genuinely support.

## CHAPTER 13

# LOCATION PAGES WITHOUT DOORWAY SPAM

Location pages are useful when location changes the answer.

They become spammy when a business creates dozens of almost identical pages and swaps only the town name.

The difference is not whether a town appears in the title. The difference is whether the page provides genuine value to someone in that place.

## Start with the truth about your operating model

There are three common cases.

### 1. You have a real location there.

A shop, office, clinic, branch or other staffed premises exists.

### 2. You are a service-area business.

You travel to customers in the town but do not have a public office there.

### 3. You do not meaningfully serve the area.

The town is being added only because a keyword tool showed demand.

Do not blur these categories.

## What makes a location page useful?

Useful local content can include:

- which services are available there;
- whether response times differ;
- travel or minimum-job policies;
- examples of work completed nearby;
- photographs with permission;
- local property or industry considerations that genuinely affect the service;
- parking, access or appointment information for a real branch;
- contact details and next steps;
- FAQs that are genuinely location-specific.

A page for "Pest control in Northwich" should not be the same as "Pest control in Winsford" with the place name changed six times.

## Local knowledge should be relevant, not decorative

Some pages include paragraphs about a town's history purely to look unique.

A customer looking for a drain repair does not need to know when the town hall was built.

Local detail is useful when it changes the job or decision.

Examples:

- common property types;
- access constraints;
- local planning or conservation issues where relevant;
- seasonal demand;
- service radius;
- known appointment patterns;
- branch facilities.

Do not invent local expertise you do not have.

## Thin location pages create a maintenance problem

Every extra page becomes another thing that can become inaccurate.

If the business changes its prices, opening hours, service limits or contact details, thirty near-duplicate pages may all need updating.

That is one reason fewer, stronger pages often outperform an uncontrolled location-page factory as a business asset.

## Google explicitly warns against scaled low-value content

Google's spam policies define scaled content abuse as producing many pages mainly to manipulate rankings without adding value, regardless of whether automation or AI was used.

This does not mean every templated page is spam. It means the purpose and value matter.

If each branch page represents a real branch and contains accurate branch information, templates can be sensible.

If each town page is a thin doorway leading to the same service, the risk is much higher.

## Use a hub when a page does not deserve to exist

Sometimes the best local structure is:

### **Areas we cover**

A useful overview listing real coverage and linking only to locations that have enough distinct information to justify their own page.

This is cleaner than forcing every village into a separate URL.

## A location-page quality test

Before publishing, ask:

1. Would this page still be useful if the town name were removed from the title?
2. Does it contain information specific to this location?
3. Can the business genuinely fulfil the service there?
4. Is the page meaningfully different from neighbouring location pages?
5. Does it help the customer decide or act?

If the first four answers are weak, consolidate.

***Do this next:** Export all location URLs. Group them into: genuinely useful, duplicate/thin, inaccurate, and unnecessary. Improve or merge the weak group rather than creating more.*

## A quick decision tree before creating another town page

Before creating service-in-town number twenty-three, answer these questions.

**Does the business genuinely serve the place?** If no, stop.

**Would a customer in that place need information that is meaningfully different from the main service page?** If no, a strong service-area hub may be better.

**Do you have something real to say?** That could be projects, access considerations, travel policies, local property types, booking arrangements, photographs or other relevant experience. If the only difference is the town name, the page is probably thin.

**Can you maintain it?** Twenty location pages multiply every pricing change, phone-number update, policy revision and design change. If nobody will look after them, quantity becomes a liability.

**Would you be comfortable showing the page to a customer from that town?** If the page feels like it was made for a crawler rather than a person, that is a warning sign.

A useful alternative is often a regional page that clearly lists the areas served and then links to a smaller number of genuinely substantial location pages. This can be cleaner for customers and easier to keep accurate.

The aim is not to minimise the number of pages. It is to make each page earn its place.

### What a genuinely useful location page sounds like

A good location page rarely needs to spend half its length describing the town.

The customer already knows where they live. What they need to know is how your service works **there**.

Compare two approaches.

A thin page says the company is proud to offer high-quality services in Town A, lists the same services shown everywhere else and finishes with a generic contact form. Swap Town A for Town B and nothing meaningful changes.

A useful page might explain that the business serves Town A on particular days, whether there is a call-out charge, what kinds of jobs are commonly booked there, how far in advance appointments normally need to be arranged, which neighbouring areas are covered on the same route and what the customer should prepare before the visit. If the business has genuine completed work in the area, it can show an appropriate example or photograph with permission.

Notice what makes the second page local: not a paragraph copied from Wikipedia, but operational detail.

For a service-area business, useful local information can include:

- whether the company is based in the town or travels into it;
- which services are actually available there;
- any genuine travel or minimum-job conditions;
- normal booking or response arrangements, described carefully rather than promised recklessly;
- relevant examples of completed work;
- photographs that really belong to the business;
- neighbouring areas covered on the same service route;
- links back to the core service pages so the visitor can get proper detail.

Not every point belongs on every page. A solicitor does not need to discuss travel routes in the same way as a landscaper. A salon with a physical premises has a different local story from a mobile pest controller. The page should reflect the operating model.

The safest writing test is brutally simple: **could you replace the town name and publish the page somewhere else without changing anything important?**

If yes, the page probably does not contain much local value.

There is also nothing wrong with deciding that a place does not deserve its own page. A strong "Areas we cover" hub can be clearer than fifteen near-identical town pages. It gives the customer one trustworthy place to check coverage and saves the business from maintaining fifteen versions of the same facts.

Location pages should exist because they help a local customer make a decision. Ranking potential is a benefit, not a sufficient reason for the page to exist.

## CHAPTER 14

# REVIEWS, REPUTATION, LINKS AND MENTIONS

Trust on the web is distributed.

Your own website is one source. Customers, suppliers, directories, professional bodies, local media and other websites can all provide additional signals.

This matters because a business describing itself is not the same as other people independently recognising it.

## Reviews are evidence with context

A star rating is a summary. The written review often contains the useful detail.

A good review might reveal that the business:

- arrived when promised;
- explained a problem clearly;
- respected the property;
- handled an emergency calmly;
- kept to the quote;
- fixed a mistake;
- communicated delays;
- supported the customer afterwards.

These are the themes that influence real decisions.

## Ask for reviews ethically

Make review requests part of a normal follow-up process.

Do not:

- offer rewards for positive reviews without proper disclosure;
- write reviews for customers;
- pressure people to use particular wording;
- hide legitimate criticism;
- create fake accounts.

A healthy reputation includes the occasional imperfect review and a sensible response.

## Respond like a business, not a courtroom

A review response is public customer service.

Thank people for useful feedback. Correct factual errors calmly when necessary. Do not publish private customer details in an argument.

A measured response to criticism can increase trust because it shows how the business behaves when something goes wrong.

## Links are still useful when they reflect real relationships

A link from a relevant site can help discovery and prominence. The best links usually come from something that exists beyond SEO.

Examples:

- manufacturer installer directory;
- professional association;
- local business group;
- supplier case study;
- community sponsorship;
- event page;
- partner page;
- local press story;
- useful guide cited by another organisation.

The link is valuable because the relationship is real.

## **Mentions without links can still matter commercially**

A person may hear about the business on a Facebook group, a local forum, a trade community or an AI answer, then search the brand directly.

That discovery path may never appear as a neat referral in analytics.

This is why brand-search trends and enquiry-source questions are useful alongside conventional link reports.

## **Directories: fewer good ones beat hundreds of bad ones**

Prioritise directories that customers or search systems might reasonably use.

Relevant examples can include:

- Google Business Profile;
- Bing Places;
- respected trade directories;
- local chamber or business associations;
- niche industry directories;
- major mapping and navigation ecosystems where relevant.

Do not pay for a bulk package simply because it promises "500 citations".

## **Reputation should be consistent with the website**

If reviews repeatedly praise emergency support but the website never explains that service, add a useful page.

If customers frequently mention a particular strength, consider whether it deserves evidence on the site.

Reputation can improve content strategy because it reveals what people actually value.

## Build a proof library

Create a simple folder containing:

- approved testimonials;
- project photographs;
- accreditations;
- press mentions;
- awards;
- case-study notes;
- common review themes;
- permission records where needed.

Then reuse that evidence accurately across service pages instead of inventing new marketing claims each time.

***Do this next:** Read your last twenty reviews. List the five themes customers mention most often. Compare those themes with the promises on your homepage. Are you promoting the things customers actually value?*

## Reputation works best as a loop, not a campaign

Many businesses treat reviews as something to think about twice a year when somebody notices the star rating. A better approach is to make reputation part of the normal customer journey.

Do the work. Ask at an appropriate moment. Make it easy for the customer to leave an honest review. Respond like a human. Learn from recurring comments. Feed useful evidence back into the website.

That last step is often missed.

If several customers praise communication, perhaps the site should explain how updates are handled during a job. If people repeatedly mention cleanliness, show the working process. If reviews keep answering the same question about parking, children, nervous pets,

turnaround time or aftercare, that is a clue that the website should answer it too.

Reviews are not just ranking material. They are customer research.

The same applies to links and mentions. A legitimate link from a local trade body, supplier, community organisation, event, partner or publication usually exists because something happened in the real world. Those relationships are more defensible than buying hundreds of anonymous directory links.

A good question for any reputation tactic is: **would this still make sense if search engines stopped counting it tomorrow?**

If the answer is yes because it helps customers find you, verifies the business, supports a relationship or builds real recognition, it is probably worth more attention than a tactic designed only to manufacture a signal.

## How to ask for useful reviews without writing them for the customer

Businesses often want more detailed reviews, then accidentally make the request so awkward that the customer either writes nothing or copies the wording they were given.

You do not need to script the review. You can make the task easier by reminding the customer what kinds of details are helpful.

A natural request might say that an honest review is appreciated and, if they are not sure what to write, they could mention what work was done, how the process went and anything that stood out to them.

That leaves the opinion with the customer.

The distinction matters. "Please mention that we were fast, professional and good value" is steering the evidence. "If useful, mention what you asked us to do and how you found the experience" is a prompt to remember the job.

Detailed reviews are commercially useful because they answer the questions future customers actually have. A five-star rating says somebody was pleased. A sentence explaining that the engineer called before arriving, explained two repair options and left the kitchen tidy says far more about what dealing with the company feels like.

Look for recurring language rather than isolated compliments. If ten customers independently mention communication, that is stronger evidence than the company repeatedly calling itself communicative. If people mention a particular part of the process as confusing, that is equally valuable. Reputation work should improve the business, not just decorate the website.

A sensible review routine can be very light-touch:

- ask after the customer has had enough time to judge the work;
- provide a direct link so they are not hunting for the right profile;
- ask for an honest review rather than a positive one;
- never invent wording on their behalf;
- respond to criticism calmly and specifically;
- record useful themes for future service-page improvements;
- keep permission and attribution straight when reusing testimonials on your own site.

The final point is easily overlooked. A review platform and your own marketing material are not the same place. If you want to reproduce a customer's words in a brochure or case study, handle that properly rather than assuming a public review gives unlimited permission for every reuse.

Good reputation work should feel like normal customer service. If it starts to feel like a scheme for manufacturing praise, it has gone off course.

## CHAPTER 15

# SERVICE PAGES THAT ACTUALLY CONVERT

A service page has two audiences, but one job.

Search systems need to understand it. Customers need to trust it. The job is to move a suitable customer towards a sensible next step.

## Open with clarity

The first screen should answer three questions quickly:

1. What is the service?
2. Who or where is it for?
3. What should the visitor do next?

Avoid starting with a paragraph about how passionate the company is.

A clearer opening might say:

**Commercial pest control for shops, hospitality and small industrial premises across Cheshire. Routine contracts and one-off treatment available. Call or request a survey.**

The customer knows immediately whether the page is relevant.

## Explain the problem you solve

Do not assume the visitor knows exactly what they need.

Describe common situations:

- symptoms;
- triggers;
- warning signs;
- when a job is urgent;
- when another service is more appropriate.

This helps capture customers earlier in the decision journey.

## Explain what is included

Ambiguity creates weak leads.

A good service page can explain:

- survey or diagnosis;
- preparation;
- materials or method;
- timescale;
- clean-up;
- reporting;
- follow-up;
- warranty or limits;
- exclusions.

You do not need to reveal proprietary detail. You do need to set expectations.

## Price transparency is a spectrum

Not every service can publish a fixed price.

Useful alternatives include:

- starting price;
- typical range;
- example job;
- package comparison;
- minimum charge;
- list of factors that affect the quote.

"Contact us for a price" gives the customer no help.

"Most domestic call-outs start from £X. The final price depends on access, severity and whether follow-up treatment is required" is more informative without pretending every job is identical.

## Show proof near the claim

Do not put all evidence on a separate testimonials page and expect visitors to go looking for it.

If the page says you handle complex commercial work, show a relevant commercial case study.

If the page says you cover a location, show a real job there where appropriate.

If the page says the process is tidy, use photographs that demonstrate it.

## Handle objections before the form

Common objections include:

- "Will this cost more than expected?"
- "How long will I be without the service?"
- "Do I have to be there?"
- "Will you damage anything?"
- "What happens if the problem comes back?"
- "Can you work around opening hours?"

Answering those questions can improve conversion and AEO at the same time.

## Make the call to action specific

"Contact us" is functional but vague.

Better calls to action describe the next step:

- Request a site survey
- Send photos for an initial assessment
- Check appointment availability
- Ask for a migration review
- Book a consultation
- Get a written quote

Then explain what happens after the click.

## A service-page blueprint

A strong page often includes:

1. Clear service statement.
2. Who it is for and where it is available.
3. Common problems or use cases.
4. What the service includes.
5. Process and timescale.
6. Price or price factors.
7. Evidence and case study.
8. FAQs and objections.
9. Clear next step.

You do not need every section on every page. Use what the customer needs.

***Do this next:** Compare your most important service page with the blueprint above. Add the single missing section that would remove the biggest buyer uncertainty.*

## A before-and-after service page, in words

Imagine a pest-control page that currently reads:

*Fast, professional pest control across Cheshire. We handle all pests. Contact us today for a free quote.*

Nothing there is necessarily false. The problem is that almost any competitor could say it.

A stronger page would open by naming the actual service and area, then immediately answer the main buying questions. It might explain which pests are covered, whether inspection is needed first, what customers should do before the visit, how treatment around children

or pets is handled, what affects price and whether follow-up visits are sometimes required.

It would then show evidence: genuine photographs where appropriate, the technician or company behind the work, relevant qualifications, real reviews and a short case example.

The call to action might change from "Contact us" to something more useful:

**Tell us what you have seen and where you are. If you can safely take a photo, send that too. We will tell you the likely next step and whether a visit is needed.**

That small change reduces uncertainty. It also produces a better enquiry because the customer knows what information to provide.

The same pattern works in almost every sector: state clearly what the service is; explain who it is suitable for; answer the main price, timing and process questions; show why the business is credible; deal with the obvious objections; and make the next step specific.

A service page should feel less like a poster and more like the first five minutes of a useful conversation.

## **Worked example: turning a vague pest-control page into a decision page**

This is another illustrative example. Picture a small pest-control firm whose existing wasp page says little more than "fast, professional wasp treatment across Cheshire" followed by a phone number. The page is short, but short is not the problem. The problem is that the customer still has almost every important question unanswered.

The owner starts by writing down the questions that arrive on the phone in summer. Is it definitely a wasp nest? Can the nest be in a loft or wall cavity? Should people block the entrance? Is it safe to stay in the house? What happens if the nest cannot be reached directly? Does

the price change with height or access? How quickly can somebody attend?

Those questions become the structure of the page. The opening says exactly what the service is, the area genuinely covered and the next step. A short identification section explains common signs while making it clear that photographs do not always provide a certain diagnosis. A preparation section tells the customer what not to do before the visit. The price section explains the factors that can change the quote rather than hiding behind "contact us for pricing".

Proof sits beside the claims. If the page says the firm works at height, a suitable job photograph is more persuasive than a stock image of a wasp. If the business holds a relevant qualification or membership, it is named accurately. If customers regularly mention tidy work or clear explanations in genuine reviews, selected quotations can support those points without pretending every review says the same thing.

The call to action also changes. Instead of a bare "Contact us" button, the page asks the customer to send the location, what they have seen and a photograph if it can be taken safely. That makes the next step clearer and can improve the quality of the enquiry before anyone picks up the phone.

For AEO, the page now contains direct answers to real questions. For GEO, it contains clearer facts, boundaries and evidence. For ordinary SEO, it is more relevant and useful. For the customer, which is the bit that matters most, it feels like dealing with a competent person rather than reading an advert.

The important lesson is that none of this requires adding hundreds of words for their own sake. The page becomes stronger because it removes uncertainty in the order the customer actually experiences it.

## When one service deserves more than one page

There is a temptation at both extremes. Some websites squeeze every service into one enormous page. Others create dozens of thin pages because a keyword tool has found dozens of variations.

Neither approach is automatically right.

A separate page earns its place when the customer has a meaningfully different problem, buying decision or process to understand.

A pest-control business, for example, may reasonably have separate pages for wasps, rodents and bed bugs because the signs, urgency, treatment process, preparation and customer questions are different. A garage may separate MOT testing, diagnostics and air-conditioning servicing for the same reason. An accountant may distinguish bookkeeping, payroll and year-end accounts when each service has its own audience and decision process.

By contrast, creating separate pages for **"website designer Winsford"**, **"web designer Winsford"**, **"website design Winsford"** and **"web design company Winsford"** would usually be dividing language rather than dividing a real customer need.

A useful test is to ask four questions:

1. Would a customer reasonably expect a different explanation for this service?
2. Are the price factors, process or suitability questions meaningfully different?
3. Can we provide evidence or examples specific to it?
4. Would the page still deserve to exist if search engines did not?

If the answer is yes to most of those, a dedicated page may help. If the only difference is the phrase you hope to rank for, it probably belongs as wording within a stronger existing page.

This matters even more in an AI-search environment because retrieval systems benefit from clear topical boundaries, but they do not need a new URL for every way of phrasing the same idea. A coherent page can answer several closely related questions perfectly well.

There is a practical maintenance benefit too. Every extra page creates another place where prices, hours, credentials, service areas and policies can become stale. A smaller set of strong pages is often easier for a small business to keep accurate than a sprawling library built during an enthusiastic SEO campaign and forgotten six months later.

Think of your website as a useful filing cabinet, not a bag of keywords. Give genuinely different subjects their own drawer. Keep variations of the same subject together.

## The conversion test

Before publishing, give the page to somebody who knows nothing about the business and ask them five questions:

1. What exactly is being offered?
2. Is it available to someone like you?
3. What will probably happen next?
4. What makes the business believable?
5. What would you do if you wanted to proceed?

If they cannot answer those after a quick read, there is more work to do.

## A useful AI-assisted writing workflow

AI can be genuinely helpful when you are improving a business website. It can also produce an impressive amount of bland material very quickly. The difference usually comes down to what you ask it to work **from**.

A weak workflow begins with: "Write me an SEO page for emergency plumbers in Cheshire." The system has no real knowledge of the business, so it fills the gaps with common phrases, assumptions and generic advice. The result may sound polished while saying very little that another plumber could not publish.

A stronger workflow begins with the business's own raw material:

- notes from real customer calls;
- quotations and the questions people ask before accepting them;
- service checklists;
- photographs and project notes;
- actual price factors;
- common reasons a job becomes more complicated;
- service-area rules;
- cancellation, booking or warranty policies;
- and the owner's explanations of what customers tend to misunderstand.

Then use AI as an editor, organiser or interviewer rather than an imaginary expert.

You might ask it to group fifty real customer questions into themes, identify where two answers contradict one another, turn rough notes into a clearer first draft, suggest headings, or challenge a page by listing the questions a cautious buyer could still have after reading it.

The business owner or subject expert still supplies the facts and makes the final judgement.

A practical workflow is:

**Collect -> organise -> draft -> verify -> humanise -> publish -> revisit.**

**Collect:** Gather the real source material before asking for prose.

**Organise:** Group it around the customer's decision, not around a list of keywords.

**Draft:** Use AI if it saves time, but tell it which facts are fixed and which must not be invented.

**Verify:** Check every price, location, credential, technical claim and promise against a real source.

**Humanise:** Put back the phrases, judgement and useful caveats that sound like the people who actually do the work.

**Publish:** Make the page easy to navigate and give it a clear next step.

**Revisit:** Update it when the business changes or when new customer questions keep appearing.

The test is not whether a detector thinks a human or a machine wrote the page. Those tools are not a sensible publishing strategy. The test is whether the page is accurate, useful, distinctive, maintainable and recognisably about the real business.

AI is most valuable when it helps you uncover and communicate the expertise you already have. It is least valuable when it is asked to manufacture expertise on your behalf.

## PART FIVE - MEASUREMENT, CONTROL AND ACTION



### CHAPTER 16

# MEASURING VISIBILITY WHEN CLICKS ARE NO LONGER THE WHOLE STORY

For years, organic search reporting was built around a familiar chain:

**rankings -> impressions -> clicks -> sessions -> conversions**

That chain still matters, but AI-assisted discovery adds new steps and missing data.

A person can see a brand in a generated answer and never click. They can search the brand later. They can open Maps. They can ask a follow-up question. They can use an assistant that sends a referral with its own tracking label.

Measurement therefore needs a wider lens.

### Keep the classic metrics

Do not throw away:

- impressions;
- rankings;
- clicks;
- landing-page visits;
- conversions;
- calls;
- bookings;
- revenue where traceable.

These still tell you whether search visibility is producing business.

## Add generative visibility where reliable data exists

By September 2026, both Google and Microsoft had introduced more explicit reporting around AI visibility.

Google announced dedicated generative AI performance reports in Search Console in June 2026 and later noted that the insights had rolled out worldwide by the end of August.

Bing Webmaster Tools introduced AI Performance reporting that shows how pages are cited across Microsoft Copilot, Bing AI summaries and selected partner experiences.

These tools make it possible to move beyond screenshots and anecdotes.

## Track ChatGPT referral traffic separately

OpenAI says links from ChatGPT search include `utm_source=chatgpt.com`, which makes referral traffic easier to identify in analytics.

That does not measure every possible brand exposure inside ChatGPT. It measures referrals that actually arrive on the site.

## Direct AI traffic may be small while influence is larger

Direct referral traffic from AI tools can be measured where a platform passes a recognisable referrer or tracking parameter, but those visits are only one part of the picture. A business should not abandon Google, Maps, direct traffic or referrals because a dashboard shows a new AI channel. The commercial significance of AI discovery may arrive through influence and assisted journeys before it dominates direct traffic.

## Build a sensible modern scorecard

Track five groups.

### 1. Search visibility

Impressions, rankings, query themes and pages gaining or losing visibility.

### 2. Website engagement

Useful visits, service-page journeys, form starts and content that assists enquiries.

### 3. Leads and revenue

Calls, forms, bookings, quote requests and sales where practical.

### 4. Local actions

Business Profile interactions, map visibility, directions, calls and local enquiries.

### 5. AI visibility and referrals

Search Console generative AI reports, Bing AI citations, cited URLs, ChatGPT referrals and branded searches that may follow exposure.

## Ask customers how they found you

Analytics will never capture every journey.

A simple optional question on a form or during a sales call can be revealing:

"How did you first hear about us?"

Possible answers might include Google, Maps, recommendation, social media, ChatGPT/AI, previous customer or other.

Do not force the customer to reconstruct a perfect attribution model. Use the answers as directional evidence.

## Watch brand search

If a brand begins appearing in AI answers, users may search the name directly later.

Monitor branded queries in Search Console and compare them with marketing activity, seasonality and PR.

This is not proof of causation, but it can expose trends that referral data misses.

## Do not chase vanity citations

Ten citations for irrelevant informational questions are not automatically valuable.

A citation that helps a qualified customer discover a service is more useful than a hundred mentions with no commercial relationship to the business.

Measurement should answer:

- Are the right people finding us?
- Are they understanding us?
- Are we being considered?
- Are useful enquiries increasing?

***Do this next:*** Add one monthly line to your reporting for AI visibility and one for branded search. Keep them beside, not instead of, leads and revenue.

## Attribution without fooling yourself

Modern discovery makes neat attribution harder. That does not mean measurement is pointless; it means you need to be honest about what each metric can prove.

A ChatGPT referral with `utm_source=chatgpt.com` proves that a visit arrived through a tracked ChatGPT link. It does not tell you how

many people saw the business in an answer and later typed the name into Google.

A rise in branded searches can suggest increasing awareness. It cannot prove which channel caused it.

A Bing citation count shows that pages were referenced in supported AI experiences. It does not tell you that each citation was equally prominent or commercially useful.

Google's dedicated generative AI performance reports, rolled out worldwide in 2026, add another useful view of visibility. Again, visibility is not the same thing as a sale.

The answer is not to invent precision. Use a **layered scorecard**:

- business outcomes: leads, bookings, sales and revenue;
- discovery: impressions, local actions, citations and generative visibility;
- behaviour: visits, engagement and conversion paths;
- brand demand: branded searches, direct visits and repeat enquiries;
- qualitative evidence: what customers say when asked how they found you.

Look for patterns across those layers rather than trying to force every customer into one perfect last-click story.

Microsoft expanded its AI visibility tooling again in June 2026 with additional views around topics, intent, citation share and comparisons. That direction is worth watching because it shows where measurement is heading: not simply "did we get a click?" but "where does our information participate in the decision?"

The discipline is to keep the exciting new metrics beside the boring commercial ones. If AI visibility doubles while qualified enquiries fall, the dashboard should not be allowed to declare victory.

## Worked example: when fewer clicks do not automatically mean worse marketing

Imagine an accountancy practice reviewing a month in which organic clicks are down. The first reaction is obvious: something must have gone wrong. But the rest of the picture is less tidy. Branded searches are slightly stronger, calls from the website are steady and two new clients say they first encountered the firm while comparing accountants in an AI-assisted search experience.

That does not prove the AI exposure caused the enquiries. It would be equally wrong to pretend it did. What it does show is why a single metric can no longer carry the whole story.

The practice would review the classic search data first: which pages lost clicks, whether rankings or impressions changed and whether the decline was concentrated in broad informational queries or high-intent service queries. Then it would check local actions, branded demand, generative visibility reports where available and the source information collected on enquiry forms or during sales calls.

If informational clicks fell but qualified enquiries and revenue remained healthy, the right response might be to keep watching rather than launch a panicked redesign. If service-page visibility and enquiries both fell, that is a different problem. The same headline number - fewer clicks - can describe two very different commercial situations.

This is why the scorecard in this chapter deliberately keeps business outcomes at the centre. AI citations, generative impressions and referral sessions are useful evidence about discovery. They are not the final objective. The final objective is still to help the right people understand the business and take a sensible next step.

## Do not compare unlike AI metrics

The arrival of dedicated AI visibility reports creates a new temptation: put every number into one spreadsheet column and pretend the platforms are measuring the same thing.

They are not.

Google's Search Console generative AI reporting gives site owners a dedicated view of visibility within Google's generative Search features. Microsoft exposes a different set of signals in Bing Webmaster Tools, including citations, cited pages, grounding-query samples and, in later 2026 updates, views around topics, intent, citation share and comparisons.

Those are useful datasets, but a Google generative impression is not the same unit as a Bing citation. Neither is equivalent to a traditional average position.

This sounds obvious until a monthly report says something like:

"AI visibility increased 42%."

Forty-two per cent of what?

A better report keeps each platform's native metric intact and labels it plainly.

For example:

- Google generative Search impressions: up or down;
- Bing AI citations: up or down;
- number of Bing pages cited: up or down;
- ChatGPT referral sessions: up or down;
- branded search demand: up or down;
- qualified enquiries: up or down.

You can then add interpretation underneath rather than inventing a universal "AI score".

This also protects you from false confidence. A sharp rise in generative impressions may simply mean the site is appearing more often for broad informational questions. A smaller rise in citations to a high-intent service page may be commercially more interesting. Neither conclusion can be reached from the headline number alone.

Use the new reports to answer practical questions:

- Which pages are appearing?
- Which topics or query themes are associated with them?
- Are the visible pages accurate and useful?
- Does the exposure relate to services the business actually wants to sell?
- Is brand demand changing at the same time?
- Are enquiries changing?

The reporting is getting better quickly. That is good news. It is not permission to abandon measurement discipline.

## CHAPTER 17

# CRAWLERS, PREVIEW CONTROLS AND CONTENT PERMISSIONS

AI discovery has introduced a new question for website owners: **which automated systems should be allowed to access which content, and for what purpose?**

The answer is not "allow everything" or "block everything".

The right choice depends on the business model.

## Search discovery and model training can be different activities

OpenAI's current publisher guidance distinguishes OAI-SearchBot, which is used for search discovery, from GPTBot, which publishers can block if they do not want content used for potential model training.

That distinction is useful because it allows a site owner to think about search visibility separately from training preferences.

Other providers have their own user agents and policies. Do not assume one rule controls every AI system.

## Google Search has its own preview controls

Google's AI Overviews and AI Mode are part of Search. Google says normal Search controls such as `noindex`, `nosnippet`, `max-snippet` and `data-nosnippet` can affect how content is displayed in its AI features.

More restrictive preview settings can reduce how much content is available for AI-generated previews.

## Bing supports data-nosnippet

Bing introduced support for the data-nosnippet HTML attribute so publishers can exclude selected sections from search snippets and AI-generated answers while keeping the rest of a page discoverable.

This can be useful for:

- subscriber-only material;
- legal boilerplate;
- unstable A/B test text;
- promotional sections;
- sensitive or irrelevant content.

Do not use it casually on the main information you hope the system will use.

## Robots.txt is not a copyright contract

Robots.txt is a crawler instruction mechanism. It is not a substitute for understanding licensing, contracts, paywalls or legal rights.

If the business publishes valuable proprietary material, decide what should be public before deciding how to optimise it.

## Public information should be intentionally public

A surprising number of websites expose:

- staging areas;
- old PDFs;
- staff documents;
- unlinked quote files;
- test pages;
- duplicate exports;
- outdated policy documents.

Before worrying about AI, perform a basic content exposure audit.

Search engines and crawlers can only be managed sensibly when the site owner knows what is public.

## CDN and bot protection can block legitimate crawlers

A robots.txt file may allow a crawler while a firewall or bot-management service blocks it.

OpenAI specifically advises publishers to check web protection layers and verify crawler IP ranges where needed.

The same principle applies broadly: diagnose the whole request path.

## Decide your policy, then document it

For each category of public content, decide:

- Can normal search engines crawl it?
- Can search-focused AI crawlers access it?
- Is it eligible for indexing?
- Can snippets be shown?
- Is there a training preference you want to express?
- Does the page contain anything that should not have been public in the first place?

Keep a simple record so the policy survives staff or supplier changes.

***Do this next:** Review robots.txt, meta robots directives and CDN bot rules together. Make sure the access policy you think you have is the policy the server actually enforces.*

## Write a crawler policy before adding random blocks

It is increasingly common to find a robots.txt file that has been edited by several plugins, hosting tools and well-meaning people over the years. New AI crawlers arrive and another line gets added without anybody deciding what the business actually wants.

A better approach is to write a short policy first.

For each type of automated access, decide whether the business wants classic search discovery, inclusion in AI-assisted search and

summaries, model-training access where a separate control exists, preview/snippet use, agent access to public pages and no access to private, account or administrative areas.

Then map the policy to the controls each platform actually supports.

OpenAI, for example, distinguishes OAI-SearchBot for search discovery from GPTBot for potential model training. Google Search uses Googlebot and Search preview controls for its search experiences, while Google-Extended relates to some other generative AI uses. Bing supports familiar indexing/snippet controls and added data-nosnippet support for selectively excluding material from previews and AI-generated answers while leaving the page discoverable.

The exact names matter. A blanket rule copied from a blog post can block something the business actually wanted.

Also remember that `robots.txt` is only one layer. Firewalls, CDNs and bot-management products can return a 403 response even when the robots file says `Allow: /`. When diagnosing crawler access, test the actual response rather than assuming the text file tells the whole story.

For most ordinary small-business sites, the policy can stay simple: public service and information pages should normally be genuinely public; private customer areas should not be crawlable; and crawler permissions should be deliberate rather than accidental.

### **Do not confuse OpenAI's search crawler, training crawler and user-triggered fetches**

OpenAI's current crawler documentation is a useful example of why the phrase "block AI bots" is too vague to be a policy.

There are separate identities for separate jobs.

**OAI-SearchBot** is used for search discovery. If a publisher wants public pages to be eligible to appear in ChatGPT search answers, this is the crawler OpenAI tells site owners to allow. OpenAI also publishes IP ranges so hosts and CDNs can verify the traffic rather than trusting a user-agent string on its own.

**GPTBot** is associated with crawling content that may be used to improve and train OpenAI's generative foundation models. A publisher can express a different preference for GPTBot from OAI-SearchBot. In other words, a site can choose to participate in search discovery while opting out of potential model-training use.

**ChatGPT-User** is different again. It is used for some user-initiated actions when ChatGPT or a custom GPT visits a page. OpenAI's documentation says it is not an automatic web crawler and is not the control used to determine Search inclusion. Because the request is triggered by a user, normal robots.txt expectations may not apply in the same way.

That distinction is worth understanding before copying a block of robots rules from social media.

A small-business owner might reasonably decide:

- public service pages should be available to normal search and OAI-SearchBot;
- GPTBot should be allowed or disallowed according to the owner's training preference;
- private client areas should be protected by authentication and not rely on crawler rules at all;
- staging sites should not be publicly exposed in the first place;
- the CDN should not accidentally block approved search crawlers while the robots file appears to allow them.

None of this requires a grand AI governance committee for a five-page website. It requires somebody to know what the rules mean.

There is a practical security point too. User-agent strings can be spoofed. If crawler identity matters to your infrastructure, use the provider's published verification method or IP information rather than assuming that anything calling itself GPTBot really is GPTBot.

And remember that robots.txt is public. Do not use it as a list of secret directories you hope nobody will notice. Sensitive material should be protected by real access controls.

The policy question comes first: **what do we want to be public, searchable, reusable and private?** Only then should you translate that decision into crawler rules.

## CHAPTER 18

## AGENT-READY WEBSITES

Search is moving beyond answers towards actions.

An AI agent may not simply tell a user which restaurant is open. It may help compare options, inspect availability or interact with a booking flow.

For small businesses, this creates a new design question:

**Can an automated assistant understand and use the site without breaking the experience for humans?**

### Agent-readiness starts with ordinary usability

A website does not need a futuristic protocol to become easier for agents.

Start with:

- clear buttons;
- descriptive form labels;
- logical headings;
- accessible menus;
- useful error messages;
- standard links;
- predictable page structure;
- important information in text;
- sensible confirmation pages.

This is good web design regardless of AI.

### Accessibility semantics can help agents

OpenAI says browser agents can use ARIA roles, labels and states to interpret interactive elements. Google's 2026 guidance notes that browser agents may inspect the DOM and accessibility tree.

That means accessibility work can improve machine interaction as a side effect.

A button labelled "Submit" is less informative than "Request a quote".

An icon-only control with no accessible name is harder for both screen readers and agents.

## Forms should ask only for what is needed

Agents and humans both struggle with unnecessarily complicated forms.

Remove fields that do not affect the next step.

For a simple enquiry, you might need:

- name;
- contact detail;
- service;
- postcode;
- short description;
- optional photo upload.

You probably do not need the visitor's date of birth, job title and marketing preferences before they can ask a question.

## Make state changes obvious

After a form is sent, say what happened.

After a booking is confirmed, display the booking details.

If payment failed, explain the error.

If a slot has expired, do not leave the user on a dead screen.

Agents rely on the same feedback loops people do.

## Protect actions that should require human judgement

Agent-readiness does not mean removing safeguards.

High-impact actions may still need:

- confirmation;
- identity checks;
- payment verification;
- terms acceptance;
- human review;
- capacity checks.

The goal is predictable interaction, not unrestricted automation.

## Keep pricing and availability machine-readable where practical

If prices or times are dynamic, present them in a structured interface that still makes sense visually.

For a booking system, use clear dates, times and status labels.

For products, keep stock and price data current.

For services, avoid showing one price in a graphic and another in a hidden script.

## Test with assistive technology, not only agents

A website that works with keyboard navigation and a screen reader is often structurally healthier.

Accessibility testing provides value now, even if agentic web use develops more slowly than expected.

## Agent-ready is a direction, not a badge

There is no universal "agent-ready certification" for a normal small-business website.

Treat it as a design principle:

- clear meaning;
- reliable controls;
- accessible semantics;
- stable state;
- accurate data;
- safe actions.

***Do this next:** Review your main form. Replace vague labels, remove unnecessary fields and make the success state explicit. You improve human conversion today and machine usability tomorrow.*

## The next customer may arrive with a browser agent

Agent-ready design becomes easier to understand when you stop imagining a science-fiction robot and picture a customer delegating a boring task.

"Find three local garages that service this model, check whether they have Saturday opening and tell me how to book."

"Find a dog groomer near Winsford that handles nervous dogs and has an appointment next week."

"Compare these two accountants and tell me what each package includes."

A browser agent attempting those tasks needs many of the same things a person needs: clear service information, explicit labels, sensible navigation, accessible controls and a predictable booking or enquiry journey.

The difference is that visual guesswork is less dependable. A human may infer that a stylised card is clickable; an agent is more likely to depend on text labels, control states, the DOM and accessibility information. That is why the unglamorous basics - real buttons, real labels and clear states - matter so much.

That makes a few boring practices unusually future-proof:

- use proper buttons for actions;
- give form fields real labels;
- make validation errors explicit;
- show whether an action succeeded;
- do not hide essential information behind hover effects;
- make prices, availability and service constraints readable as text where practical;
- keep authentication and confirmation around actions that carry financial or legal consequences.

Do not weaken security to make an agent happy. A booking, payment, account change or contract acceptance should still have the safeguards appropriate to the action.

The goal is not to make your website controllable by every automated system. It is to make legitimate public journeys structurally clear enough that people, assistive technology and authorised agents can understand them.

## Design actions so they can be inspected before they are committed

The most important agent-friendly design principle may turn out to be one that already matters to people: make consequences clear before an action becomes irreversible.

A human can often recover from a vague interface by guessing, clicking Back and trying again. An automated agent may act more quickly and at greater scale. That makes ambiguous actions expensive.

Consider a booking form. There is a big difference between:

**Check availability**

and

**Book and charge £85.**

The interface should make that difference obvious in the visible text, the accessible name of the control and the confirmation flow.

The same applies to quote requests, account changes, subscriptions, cancellations and purchases.

A robust action flow usually has four stages:

1. **Understand** - the service, price, date, scope and conditions are readable before the action.
2. **Prepare** - the form asks for the information genuinely required.
3. **Confirm** - the user or authorised agent can review the important details before committing.
4. **Acknowledge** - the site clearly states what happened and what happens next.

This is not an argument for adding confirmation screens to every harmless click. It is an argument for matching friction to consequence.

A polished visual control with no meaningful accessible name is not as robust as it looks. The same interface can be attractive, accessible and easier for authorised agents to understand; those goals do not need to compete.

Agent-friendly design also benefits from explicit state. If an appointment is no longer available, say "That slot is no longer available" rather than silently refreshing the calendar. If a payment has succeeded, show a clear confirmation number or status. If a form is waiting for human review, say that rather than implying the job is booked.

For a small business, you can test this without any specialist agent tooling. Read the page as though you could not infer anything from colour, position or brand familiarity.

- Is the action named clearly?
- Can you tell whether money will be taken?
- Can you tell whether a booking is provisional or confirmed?
- Can you identify required fields?
- Are date and time formats unambiguous?
- Is the success state unmistakable?
- Is there a safe way to correct a mistake before commitment?

If the answers are yes, you have probably improved the experience for humans as well as future agents.

## CHAPTER 19

# COMMON MYTHS, MISTAKES AND EXPENSIVE DISTRACTIONS

Fast-moving markets create certainty merchants.

When business owners feel behind, somebody will sell them a shortcut.

Here are the claims worth challenging.

## Myth 1: SEO is dead

It is not.

Google's own current documentation says SEO remains relevant to generative AI Search because those experiences use core Search ranking and quality systems.

The presentation layer is changing. The foundation has not disappeared.

## Myth 2: You need a secret GEO schema

You do not.

Google says there is no special schema.org markup required for AI Overviews or AI Mode.

Use normal structured data accurately where it helps describe real visible content.

## Myth 3: You must create an llms.txt file or AI text file to appear in Google AI results

Google's 2026 optimisation guidance explicitly says businesses can ignore unnecessary AI text files as an AEO/GEO tactic for Google Search.

That does not mean no AI product will ever support specialised files. It means you should not treat them as a universal ranking requirement.

### **Myth 4: Publish 1,000 AI pages before competitors do**

This is one of the fastest routes to a weak site.

Google's scaled content abuse policy applies regardless of whether content is produced by people, scripts or generative AI. The problem is mass production primarily for manipulation with little added value.

Automation can help with research, outlines, editing, data processing and repetitive publishing tasks. It does not remove the need for value.

### **Myth 5: Add an FAQ for every possible question**

No.

Answer meaningful questions. Consolidate duplicates. Do not make the page painful to read because a keyword tool produced a spreadsheet.

### **Myth 6: AI citations are the new rankings**

Citations are useful signals, not a direct replacement for rankings.

Different AI products cite differently. A citation may not drive a click. A citation may be irrelevant commercially. A page can also influence an answer in ways that are not captured by a simple count.

Measure business outcomes.

### **Myth 7: If an AI tool mentions a competitor, copy its page**

Copying the visible winner does not reveal why the system selected it.

The competitor may have stronger external reputation, fresher information, better relevance, more first-hand evidence or simply have been selected in that particular response.

Use competitor analysis to learn, not clone.

### **Myth 8: Longer content always wins**

Length is not value.

A 400-word page with exact prices, opening hours and process detail can be more useful than a 3,000-word generic essay.

Write enough to satisfy the job.

### **Myth 9: Put every town in a separate page**

Only when the page deserves to exist.

Thin location pages are a maintenance burden and can become scaled low-value content.

### **Myth 10: A tool can guarantee AI recommendations**

No ethical provider can guarantee a specific ranking, citation or recommendation across changing AI systems.

Tools can help monitor and diagnose. They cannot control a black-box platform.

### **Mistake: optimising away the human**

The deepest mistake is producing pages that machines can parse but nobody wants to read.

Awkward headings, repetitive definitions and mechanically placed keyword phrases can make a technically optimised page feel untrustworthy.

The customer is still the person who pays.

## Mistake: rebuilding everything because of an acronym

If the website already has strong service pages, real proof, good technical health and accurate local information, do not throw it away to "become GEO-ready".

Improve what exists.

## Mistake: measuring only the exciting new channel

A small business can become obsessed with AI citations while missing a broken enquiry form or declining map visibility.

Novelty is not priority.

***Do this next:*** For every new SEO/AEO/GEO recommendation you receive, ask two questions: "What user problem does this solve?" and "Which official or credible evidence supports it?"

## The five-minute smell test for SEO and GEO claims

When somebody recommends a new tactic, run it through five questions before touching the website.

**1. What exact problem does this solve?** "AI visibility" is too vague.

Does the tactic improve crawlability, clarity, evidence, freshness, measurement or user experience? If nobody can say, be cautious.

**2. Is the claim supported by the platform, independent testing or only the seller?** Google now explicitly warns site owners to evaluate third-party AEO/GEO advice critically. No external tool has Google's private ranking data.

**3. Would this still help a real visitor?** Not every technical task has to be visible to a customer, but a tactic that makes the site worse for people deserves a very strong justification.

**4. What is the downside if the theory is wrong?** Adding a harmless monitoring tag is different from deleting useful pages, changing every URL or publishing thousands of thin articles.

**5. Can we measure the result?** If success can only be demonstrated using the consultant's own proprietary score, you may be buying a scoreboard rather than an outcome.

This test will not catch every bad idea, but it slows down the most expensive kind of SEO mistake: acting on certainty where the evidence is weak.

The internet will keep producing new acronyms. Your job is not to be first to obey each one. Your job is to understand what changed, decide whether it matters to your customers and improve the business accordingly.

## Myth 11: the new AI dashboard is a ranking report

The new reporting tools are useful, but they do not turn generative search into a neat league table.

Google's generative AI reports in Search Console are visibility reports. Bing's AI Performance and later AI visibility views expose citation and topic-related information. OpenAI can send identifiable referral traffic from ChatGPT search.

None of those numbers means "you rank number three in AI".

Generated answers can change with wording, context, follow-up questions, location, freshness and model behaviour. A page may be used as a source without occupying a fixed position that can be tracked like an old-fashioned keyword ranking.

Be suspicious of any tool that takes a complicated, changing set of observations and turns it into a precise universal score without explaining the sampling, platform coverage and limitations.

Use the dashboards for what they are good at: spotting pages, topics and trends worth investigating. Then return to the commercial question: did the extra visibility help the right people understand, consider or contact the business?

## Myth 12: AI can write the whole website while I get on with running the business

It can write **words** for the whole website. That is not the same as building a trustworthy source of business information.

The difficult parts of a good site are often the bits that require judgement: what you really include in the service, what you do not do, which jobs are unsuitable, what changes the price, how far you travel, what a customer must prepare, which promise is safe to make and which claim needs evidence.

If those decisions have not been supplied, an AI system either has to stay vague or guess. Vagueness produces forgettable content. Guessing can produce inaccurate content.

A better use of automation is to reduce the labour around expertise rather than replace the expertise itself. Let software help organise notes, normalise formatting, identify repeated questions, suggest alternative headings or produce an initial structure. Then have somebody who understands the business check the substance.

For a small company, this can still save a surprising amount of time. It simply avoids the false economy of publishing fifty pages quickly and spending the next year correcting them.

The website does not need to prove that every sentence was typed manually. It needs to prove that the business behind the sentences knows what it is talking about.

## CHAPTER 20

# YOUR 90-DAY SEARCH VISIBILITY PLAN

You do not need to transform the whole website this week.

A better approach is to work in layers: fix access, clarify important pages, add evidence, strengthen local information and then measure what changed.

## Days 1-15: establish the baseline

### Technical checks

- Verify Google Search Console and Bing Webmaster Tools.
- Check important pages are indexed.
- Review robots.txt and noindex directives.
- Confirm HTTPS, canonical domain and sitemap.
- Test main forms, phone links and booking paths.
- Check mobile usability.

### Business information checks

- Confirm business name and contact details.
- Review Google Business Profile.
- Review Bing Places.
- Check opening hours and service areas.
- Remove or update obsolete information.

## Measurement baseline

Record:

- organic impressions;
- clicks;
- core rankings or query themes;
- branded searches;
- enquiries;
- Business Profile activity;
- Bing AI citation data if available;
- Google generative AI visibility data if available;
- AI referral traffic.

Do not judge success yet. Establish a starting point.

## Days 16-30: fix the money pages

Choose the three pages most closely tied to revenue.

For each page:

- rewrite the opening for clarity;
- make the service and area explicit;
- explain what is included;
- add price guidance or price factors;
- explain the process;
- answer major objections;
- add genuine evidence;
- improve the call to action;
- add relevant internal links.

Do not add words for the sake of length.

## Days 31-45: build the question map

Collect customer questions from real conversations.

Group them into:

- price;
- timing;
- suitability;
- process;
- risk;
- comparison;
- preparation;
- aftercare;
- local availability.

Decide which questions belong on service pages and which deserve separate guides.

Publish two or three high-value pieces rather than twenty thin ones.

## Days 46-60: build the proof layer

Create or improve:

- one case study;
- real project photographs;
- team/about information;
- qualification details;
- transparent pricing examples;
- review collection process;
- external profile consistency.

This stage often produces the biggest trust improvement because it turns claims into evidence.

## Days 61-75: strengthen local visibility

Audit:

- Business Profile categories;
- service details;
- opening hours;
- photos;
- booking links;
- review themes;
- Bing Places;
- location pages;
- local directories that genuinely matter;
- links from real local relationships.

Remove thin location pages that exist only to repeat keywords.

## Days 76-90: measure, learn and repeat

Compare the new period with the baseline, allowing for seasonality.

Look for:

- pages gaining impressions;
- new query themes;
- changes in branded search;
- enquiries by landing page;
- map and profile activity;
- AI citations or generative visibility;
- referral traffic from AI assistants;
- conversion changes.

Then choose the next three pages to improve.

## The weekly maintenance rhythm

After the 90 days, a small business can keep momentum with a lightweight routine.

**Weekly**

- respond to reviews;
- check enquiries and forms;
- record repeated customer questions;
- fix obvious inaccuracies.

**Monthly**

- review search performance;
- update one important page;
- publish or improve one useful piece of evidence/content;
- check local profiles.

**Quarterly**

- technical crawl/index review;
- location and service audit;
- proof/case-study update;
- competitor and market review;
- AI visibility review.

**What success should look like**

The goal is not to "complete GEO".

The goal is a website that becomes more useful and more defensible over time.

You should see:

- clearer service pages;
- fewer contradictory facts;
- stronger proof;
- better customer questions answered;
- healthier local profiles;
- better measurement;
- more confidence in what to improve next.

Search will keep changing. A useful source gives you something stable to build on.

***Final action:*** Put the first fifteen days into the calendar. The plan only becomes strategy when somebody owns the next step.

## Give the plan an owner

A 90-day plan fails surprisingly often because everybody agrees with it and nobody owns it.

For a very small business, the owner may be the person doing the work. In a larger team, one person should still be responsible for keeping the list moving, even if a developer, copywriter or agency completes some of the tasks.

You do not need a complicated project-management system. A simple table is enough:

| Task                                 | Owner | Due     | Evidence of completion           |
|--------------------------------------|-------|---------|----------------------------------|
| Check indexing for core pages        | Alex  | 12 Sept | Search Console screenshots/notes |
| Rewrite boiler service page          | Sam   | 18 Sept | Published URL                    |
| Collect five real customer questions | Jo    | 20 Sept | Question list                    |
| Update Business Profile hours        | Alex  | 21 Sept | Profile checked                  |
| Publish first case study             | Sam   | 30 Sept | Published URL                    |

The last column matters. "Done" should mean something observable.

This is particularly important when several suppliers are involved. A website company might fix technical issues, the business owner might supply photographs and examples, and an administrator might update local profiles. Without ownership, each person assumes the missing information belongs to somebody else.

### Change fewer things at once

When a website is underperforming, there is a temptation to rewrite every page, change the design, move hosting, alter URLs, install a new SEO plugin and publish twenty articles in the same fortnight.

Apart from the operational risk, this makes learning difficult. If performance changes, you have no idea which change helped or hurt.

Prioritise in this order:

1. **Breakages and access problems** - forms, indexing, mobile failures, wrong contact details.
2. **High-value pages** - the pages closest to enquiries and revenue.
3. **Trust gaps** - missing evidence, unclear ownership, weak reviews or contradictory business information.
4. **Content gaps** - important customer questions nobody answers well.
5. **Nice-to-have refinements** - polishing pages that already do their job.

This is not a laboratory experiment; you cannot isolate every variable. But fewer simultaneous changes make it easier to understand what happened.

## The end-of-90-days review

At the end of the plan, do not begin by asking, "Did rankings go up?"

Start with the business.

- Are enquiries more relevant?
- Are customers asking fewer basic questions that the website should have answered?
- Are sales conversations starting with better-informed people?
- Are the main service pages clearer?
- Is there more real proof on the site?
- Are local profiles accurate?
- Can you now see search and AI visibility more clearly than three months ago?

Then look at the search data.

A page may gain impressions before it gains clicks. A revised service page may convert better without moving dramatically in average position. A new case study may receive few visits but become useful evidence in sales conversations. A clearer Business Profile may produce calls that never touch the website.

The review should therefore produce three lists.

**Keep** - things that are working and should become routine.

**Improve** - useful work that is incomplete or needs another iteration.

**Stop** - activities that consumed time without supporting customers, visibility or revenue.

Be willing to put something in the Stop column. A strategy becomes stronger when it learns what not to do.

## What the second 90 days should look like

The second cycle should not simply repeat the first.

By then, the foundations should be healthier. Spend more time on the gaps the data and customer conversations exposed.

One business may discover that local visibility is strong but service pages are weak. Another may have excellent content and terrible conversion. A third may be losing enquiries because nobody maintains opening hours and booking links. A fourth may be appearing in AI citations for informational topics that never lead to commercially relevant visits.

The next cycle should follow the evidence.

Choose three themes, not thirty tasks. For example:

- strengthen three high-value service pages;
- publish two evidence-rich case studies;
- improve local visibility in the two towns that genuinely matter.

That is enough to create momentum without turning marketing into a second full-time job.

## The principle that survives the acronyms

SEO will change. AEO may become a standard term or fade into ordinary SEO language. GEO may develop into a mature

measurement discipline or be replaced by another acronym. Agents will become more capable. Search interfaces will change again.

The business still needs to do four things well:

**Be findable.** Make the public information technically accessible and easy to discover.

**Be understandable.** Explain what you do in language people actually use.

**Be believable.** Support claims with evidence, consistency and real-world reputation.

**Be easy to choose.** Give the customer a clear, low-friction next step.

Those four jobs are boring enough to survive fashion. That is exactly why they are useful.

# APPENDIX A - SERVICE PAGE BLUEPRINT

Use this as a planning sheet, not a rigid template.

## 1. Clear opening

**Service:** What exactly is being offered?

**Audience:** Who is it for?

**Area:** Where is it available?

**Action:** What should a suitable customer do next?

Example structure:

[Service] for [customer type] across [real area]. [Key availability/process point]. [Specific next step].

## 2. Problems and use cases

List the situations that lead customers to this service.

- symptom/problem 1;
- symptom/problem 2;
- symptom/problem 3;
- when another service may be more appropriate.

## 3. What is included

Explain the actual work.

- initial assessment;
- preparation;
- main service;
- reporting;
- clean-up;
- follow-up;
- exclusions.

## 4. Process

1. Enquiry.
2. Assessment.
3. Quote or booking.
4. Work.
5. Completion.
6. Follow-up.

Change the stages to match reality.

## 5. Price guidance

Use one or more:

- fixed price;
- starting price;
- typical range;
- package table;
- price factors;
- minimum charge;
- example job.

## 6. Evidence

Add the proof that fits the page.

- case study;
- photographs;
- review;
- qualification;
- data;
- guarantee terms;
- real example.

## 7. Buying questions

Answer the questions that delay a decision.

- How long does it take?
- What should I do before you arrive?
- What can change the quote?
- Do I need to be present?
- What happens if the problem comes back?
- Is there a guarantee?
- Do you cover my area?

## 8. Next step

Use a specific call to action.

**Request a survey**

**Send photos for an initial assessment**

**Check availability**

**Book a consultation**

Then explain what happens after submission.

## APPENDIX B - CUSTOMER QUESTION BANK

Use these prompts to discover useful content ideas. Do not publish every question automatically.

### Price and value

- How much does the service cost?
- What affects the final price?
- Is there a call-out charge?
- What is included?
- What costs extra?
- Is a cheaper option suitable?
- Do you offer packages?
- When is payment due?

### Timing

- How quickly can you attend?
- How long does the job take?
- How far ahead should I book?
- Do you work evenings or weekends?
- What happens if the job takes longer?

### Suitability

- Is this service right for my situation?
- What are the alternatives?
- When should I not choose this option?
- What information do you need before advising me?

## Process

- What happens first?
- Do you need a site visit?
- Do I need to prepare anything?
- Will there be disruption?
- Do I need to be present?
- What happens afterwards?

## Risk

- What can go wrong?
- What happens if you find another problem?
- Is the work guaranteed?
- Are there limitations?
- What should I do if the issue returns?

## Local

- Do you cover my postcode?
- Is there a travel charge?
- Are response times different in my area?
- Do you have a branch here or do you travel in?

## Comparison

- Option A vs Option B: which is better for whom?
- Repair vs replacement?
- One-off service vs ongoing plan?
- DIY vs professional service?

## Trust

- Who will do the work?
- What qualifications do you hold?
- Are you insured?
- Can I see examples?
- What do previous customers say?
- What happens if I am unhappy?

# APPENDIX C - TECHNICAL AND LOCAL AUDIT CHECKLIST

## Crawl and index

- ☐ Important pages return a successful response.
- ☐ HTTPS works correctly.
- ☐ Preferred domain redirects consistently.
- ☐ Robots.txt does not block important pages unintentionally.
- ☐ Core pages are not accidentally noindex.
- ☐ Canonical tags make sense.
- ☐ XML sitemap contains important indexable URLs.
- ☐ Google Search Console is verified.
- ☐ Bing Webmaster Tools is verified.

## Page content

- ☐ Main service is clear in the opening screen.
- ☐ Important facts exist as text, not only images.
- ☐ Headings describe real sections.
- ☐ Price or price factors are explained where useful.
- ☐ Process and timescale are clear.
- ☐ Claims have evidence.
- ☐ Calls to action are specific.
- ☐ Internal links connect related content.

## Experience

- ☐ Mobile layout works.
- ☐ Forms work.
- ☐ Phone links work.
- ☐ Buttons have descriptive labels.
- ☐ Important images have sensible alt text.
- ☐ Navigation is keyboard-accessible where practical.
- ☐ Success/error states are clear.
- ☐ Page speed is reasonable on mobile.

## Local

- ☐ Business name is consistent.
- ☐ Telephone number is current.
- ☐ Address/service-area model is accurate.
- ☐ Opening hours are current.
- ☐ Google Business Profile categories are accurate.
- ☐ Bing Places is current.
- ☐ Website services match profile services.
- ☐ Location pages are genuinely useful.
- ☐ Reviews are collected ethically.
- ☐ Real local photos are used where appropriate.

## AI and crawler controls

- ☐ Search crawler policy has been reviewed.
- ☐ OAI-SearchBot access matches your ChatGPT search preference.
- ☐ GPTBot/training preference has been considered separately.
- ☐ Google preview controls match your intent.
- ☐ Sensitive or private content is not accidentally public.
- ☐ CDN/WAF rules do not accidentally block legitimate crawlers you intend to allow.

## Measurement

- ☐ Organic visibility is tracked.
- ☐ Leads and calls are tracked where practical.
- ☐ Branded search is monitored.
- ☐ Local profile activity is reviewed.
- ☐ Google generative AI performance data is reviewed if available.
- ☐ Bing AI Performance is reviewed if available.
- ☐ AI referral traffic is separated in analytics.

## APPENDIX D - GLOSSARY

### AEO - Answer Engine Optimisation

An industry term for making information easier to use in direct answers, snippets, assistants and other answer-oriented experiences.

### AI citation

A source reference shown alongside or within an AI-generated response.

### AI Overview

Google's generated summary experience that can appear in Search for queries where Google determines it adds value.

### AI Mode

A Google Search experience designed for more exploratory, complex and follow-up interactions.

### Canonical URL

The preferred URL for a piece of content when multiple URLs are similar or duplicated.

### Crawlability

Whether a crawler can access a page.

**Entity**

A distinct thing a search system can identify, such as a business, person, place, product or organisation.

**GEO - Generative Engine Optimisation**

An academic and industry term for improving how content participates in generative search and AI answers, including retrieval, citation and source use.

**Grounding**

Connecting an AI-generated response to current external evidence, often retrieved from search indexes or other trusted sources.

**Index**

A search engine's processed collection of web content used for retrieval and ranking.

**IndexNow**

A protocol that lets websites notify participating search engines when URLs are added, updated or removed.

**Internal link**

A link from one page on a website to another page on the same site.

**LocalBusiness structured data**

Schema markup that can describe a local business and properties such as name, address and opening hours.

**Noindex**

A directive asking compliant search engines not to include a page in their index.

**OAI-SearchBot**

OpenAI's crawler associated with search discovery. OpenAI advises publishers not to block it if they want public content eligible for inclusion in ChatGPT search summaries and citations.

**People-first content**

Google's term for content created primarily to help people rather than manipulate search rankings.

**Query fan-out**

Google's term for a technique used by some AI search experiences to issue multiple related searches across subtopics and sources while building an answer.

**RAG - Retrieval-Augmented Generation**

A technique that improves generated answers by retrieving relevant external information and using it as context.

**Robots.txt**

A file that provides crawler access instructions for parts of a website.

**SEO - Search Engine Optimisation**

The practice of improving a website's discoverability, understanding, relevance and performance in search.

**Structured data**

Machine-readable markup that describes the meaning of visible page information in a standard format.

**UTM parameter**

A tracking parameter added to a URL so analytics can identify referral sources or campaigns.

## APPENDIX E - RESEARCH NOTES AND FURTHER READING

This book deliberately prioritises first-party platform documentation and neutral research where possible. Industry studies are used selectively and are labelled as such.

Source audit completed 8 September 2026. The principal Google Search, Google Business Profile, Microsoft Bing, OpenAI and W3C references below were re-checked against their live documentation on that date. Fast-changing product controls and reporting features should still be re-checked before any future edition.

The wording in the book is original. These sources support factual claims, product behaviour and the broader conclusions discussed in the chapters.

### Chapter source map

The chapter notes have been gathered here so the main text reads like a book rather than a research report. They show the main evidence base for each chapter; the full source details follow below.

**Chapter 1.** This chapter draws on Google's official documentation for AI features in Search, Microsoft's explanation of grounding and the evolving search index, OpenAI's publisher guidance for search discovery, and Pew Research Center studies of AI summaries and click behaviour. Full references appear in Appendix E.

**Chapter 2.** This chapter draws on Google Search Central's 2026 generative AI optimisation guidance, Microsoft's 2026 grounding and AI Performance documentation, and the original GEO research paper. The definitions here are practical working definitions, not claims that all platforms use these labels identically.

**Chapter 3.** This chapter synthesises Google documentation on AI features and query fan-out, Microsoft's material on grounding and index design, Bing Webmaster Tools AI Performance, Google's generative AI Search Console reporting, OpenAI publisher guidance and recent academic work on citation selection and absorption.

**Chapter 4.** This chapter is grounded mainly in Google Search Central's people-first content guidance, 2026 generative AI optimisation guide and spam policies, with local prominence principles from Google Business Profile help.

**Chapter 5.** This chapter draws on Google's technical and generative AI search guidance, Bing Webmaster Tools principles and OpenAI's current publisher FAQ distinguishing search discovery from training controls.

**Chapter 6.** Google's current AI search guidance continues to emphasise good page experience and accessible, crawlable content. OpenAI's developer guidance for browser agents highlights ARIA and accessible interface semantics as useful for machine interaction.

**Chapter 7.** This chapter reflects Google's emphasis on helpful, satisfying, people-first content and its warning against creating pages for every query variation. It also aligns with Microsoft's guidance that AI retrieval benefits from clear, verifiable and applicable information.

**Chapter 8.** This chapter draws on the original GEO paper, Microsoft's 2026 grounding and AI Performance material, Bing guidance on freshness and IndexNow, Google's emphasis on unique non-commodity content, and 2026 academic work separating citation selection from citation absorption.

**Chapter 9.** This chapter builds on Google's people-first content guidance, general structured-data quality principles and broader search policies against misleading content. The practical frameworks are original to this book.

**Chapter 10.** This chapter reflects established SEO principles around internal linking, duplication and findability, as reinforced in Google's current AI search guidance. The architecture model and exercises are original to this book.

**Chapter 11.** This chapter relies on Google's general structured data guidelines, LocalBusiness and Organization documentation, review-snippet rules and 2026 AI search guidance stating that no special AI schema is required.

**Chapter 12.** This chapter draws on Google Business Profile help for relevance, distance and prominence; Google's local structured data and AI search guidance; and Microsoft's Bing Places and AI Performance guidance for local businesses.

**Chapter 13.** This chapter draws primarily on Google's current spam policies, including scaled content abuse, together with general people-first content guidance and local-business accuracy principles.

**Chapter 14.** This chapter aligns with Google's local prominence guidance and review policies, including current restrictions around self-controlled review markup for local businesses. The reputation framework is original to this book.

**Chapter 15.** This chapter synthesises people-first content principles, AEO-style answer clarity and conversion best practice. The service-page blueprint is original to this book.

**Chapter 16.** This chapter uses current Google Search Console generative AI reporting, Bing Webmaster Tools AI Performance and visibility insights, OpenAI referral tracking guidance and Pew click research. Independent studies are presented as directional evidence rather than universal benchmarks.

**Chapter 17.** This chapter uses current official guidance from Google Search, Bing and OpenAI. Crawler names and controls can change, so re-check the provider documentation before making production changes.

**Chapter 18.** This chapter draws on Google's 2026 guidance on agentic experiences and OpenAI's developer FAQ describing how browser agents use ARIA and accessible page semantics.

**Chapter 19.** This chapter draws heavily on Google's 2026 AI optimisation guide, current spam policies and structured data guidance, with supporting context from Bing and academic GEO research.

**Chapter 20.** The 90-day plan is an original PixelGem framework built from the principles and current platform guidance cited throughout this book.

## Google Search Central

### AI Features and Your Website

Google Search Central documentation. Covers AI Overviews, AI Mode, eligibility, query fan-out and the continued relevance of SEO fundamentals.

<https://developers.google.com/search/docs/appearance/ai-features>

### Optimizing Your Website for Generative AI Features on Google Search

Google Search Central, 2026. Covers non-commodity content, crawlability, page experience, warnings against scaled query pages and unnecessary AEO/GEO hacks, generative AI measurement and agentic experiences.

<https://developers.google.com/search/docs/fundamentals/ai-optimization-guide>

### Top Ways to Ensure Your Content Performs Well in Google's AI Experiences on Search

Google Search Central Blog, 21 May 2025.

<https://developers.google.com/search/blog/2025/05/succeeding-in-ai-search>

## **Creating Helpful, Reliable, People-First Content**

Google Search Central.

<https://developers.google.com/search/docs/fundamentals/creating-helpful-content>

## **Spam Policies for Google Web Search**

Includes scaled content abuse and other manipulative practices.

<https://developers.google.com/search/docs/essentials/spam-policies>

## **General Structured Data Guidelines**

Google Search Central.

<https://developers.google.com/search/docs/appearance/structured-data/sd-policies>

## **LocalBusiness Structured Data**

Google Search Central.

<https://developers.google.com/search/docs/appearance/structured-data/local-business>

## **Organization Structured Data**

Google Search Central.

<https://developers.google.com/search/docs/appearance/structured-data/organization>

## **Review Snippet Structured Data**

Google Search Central. Includes rules affecting self-controlled reviews for local businesses and organisations.

<https://developers.google.com/search/docs/appearance/structured-data/review-snippet>

## **Introducing Search Generative AI Performance Reports in Search Console**

Google Search Central Blog, 3 June 2026, with later rollout update.

<https://developers.google.com/search/blog/2026/06/gen-ai-performance-reports>

### **Google Search Documentation Updates - May 2026**

Includes the new generative AI optimisation guide and the May 2026 deprecation of FAQ rich results.

<https://developers.google.com/search/updates>

### **Google Business Profile**

#### **Tips to Improve Your Local Ranking on Google**

Google Business Profile Help. Explains relevance, distance and prominence.

<https://support.google.com/business/answer/7091>

### **Microsoft Bing**

#### **Introducing AI Performance in Bing Webmaster Tools Public Preview**

Bing Webmaster Blog, 10 February 2026. Covers AI citation reporting, local business information, freshness and IndexNow.

<https://blogs.bing.com/webmaster/February-2026/Introducing-AI-Performance-in-Bing-Webmaster-Tools-Public-Preview>

#### **New AI Visibility Insights in Bing Webmaster Tools: Intent, Topics, Citation Share, Compare**

Bing Search Blog, 16 June 2026. Describes expanded AI visibility reporting around intent, topics, citation share and comparison views.

<https://blogs.bing.com/search/?p=7113>

#### **Elevating the Role of Grounding on the AI Web**

Bing Search Blog, 12 February 2026. Discusses grounding and GEO.

<https://blogs.bing.com/search/February-2026/Elevating-the-Role-of-Grounding-on-the-AI-Web>

## **Evolving Role of the Index: From Ranking Pages to Supporting Answers**

Bing Search Blog, 6 May 2026. Explains differences between ranking for human choice and grounding for generated answers.

<https://blogs.bing.com/search/May-2026/Evolving-role-of-the-index-From-ranking-pages-to-supporting-answers>

## **Bing Introduces Support for the data-nosnippet HTML Attribute**

Bing Webmaster Blog, 15 October 2025.

<https://blogs.bing.com/webmaster/October-2025/Bing-Introduces-Support-for-the-data-nosnippet-HTML-Attribute>

## **OpenAI**

### **Publishers and Developers - FAQ**

OpenAI Help Center. Covers OAI-SearchBot, GPTBot distinctions, ChatGPT search referrals and agent accessibility guidance.

<https://help.openai.com/en/articles/12627856>

### **Overview of OpenAI Crawlers**

OpenAI Developers. Distinguishes OAI-SearchBot, GPTBot and ChatGPT-User, with current crawler purposes and published verification information.

<https://developers.openai.com/api/docs/bots>

### **Searching the Web with ChatGPT**

OpenAI Help Center. Covers website eligibility for ChatGPT search and OAI-SearchBot access.

<https://help.openai.com/en/articles/9237897>

## Accessibility standards

### Web Content Accessibility Guidelines (WCAG) 2.2

W3C Recommendation. A technology-neutral accessibility standard covering perceivable, operable, understandable and robust web content, with additional criteria for focus, target size, authentication and other practical interaction issues.

<https://www.w3.org/TR/WCAG22/>

## Independent research

### GEO: Generative Engine Optimization

Pranjal Aggarwal, Vishvak Murahari, Tanmay Rajpurohit, Ashwin Kalyan, Karthik Narasimhan and Ameet Deshpande. arXiv:2311.09735, first submitted 2023. Introduced a formal GEO framework and GEO-bench.

<https://arxiv.org/abs/2311.09735>

### From Citation Selection to Citation Absorption: A Measurement Framework for Generative Engine Optimization Across AI Search Platforms

Kai Zhang, Xinyue He and Jingang Yao. arXiv:2604.25707, 2026. Early research distinguishing source selection from how strongly a source shapes a generated answer.

<https://arxiv.org/abs/2604.25707>

### Google Users Are Less Likely to Click on Links When an AI Summary Appears in the Results

Pew Research Center, 22 July 2025. Analysis of browsing behaviour from 900 U.S. adults.

<https://www.pewresearch.org/short-reads/2025/07/22/google-users-are-less-likely-to-click-on-links-when-an-ai-summary-appears-in-the-results/>

## **A Majority of Americans Say They Read AI Summaries at the Top of Search Results**

Pew Research Center, 17 June 2026. Survey data from U.S. adults.

<https://www.pewresearch.org/chart/a-majority-of-americans-say-they-read-ai-summaries-at-the-top-of-search-results/>

## **Americans and AI 2026: Chatbots, Smart Devices and Views on Impact**

Pew Research Center, 17 June 2026. A survey of 5,119 U.S. adults; useful context on chatbot use and AI search-summary adoption.

<https://www.pewresearch.org/internet/2026/06/17/americans-and-ai-2026-chatbots-smart-devices-and-views-on-impact/>

## **A note about fast-changing guidance**

Before publishing a future edition, re-check the platform documentation above. Search products change faster than printed books.

## APPENDIX F - PRACTICAL EXAMPLES BY BUSINESS TYPE

The principles in this book become easier to use when you translate them into the language of a real business. These examples are not keyword lists to copy. They are prompts for finding questions, evidence and pages that deserve to exist.

### 1. Roofer

#### Useful customer questions

- Is this likely to be a repair or a replacement?
- How quickly does a leaking roof need attention?
- Can you work on slate, tile and flat roofs?
- What affects the price of a roof repair?
- Will scaffolding be required?
- What happens if you uncover rotten timber?

#### Source-worthy evidence

- before-and-after photographs;
- explanation of repair decisions;
- case study showing hidden damage;
- real price factors;
- materials used and why.

#### Useful page idea

A guide showing three real roof-repair scenarios, what the visible symptom was, what was found after inspection and what changed the final price.

## 2. Electrician

### Useful customer questions

- When does a consumer unit need replacing?
- What does an EICR actually check?
- How long will the power be off?
- Do I need an electrician for this job?
- What information do you need before quoting?

### Source-worthy evidence

- qualification details;
- example EICR findings explained in plain English;
- photographs of compliant completed work;
- clear process and safety boundaries.

### Useful page idea

"EICR results explained: what C1, C2 and FI findings mean for a homeowner" - written carefully and checked against current electrical guidance.

### 3. Pest controller

#### Useful customer questions

- Is this a wasp, bee or hoverfly?
- Can you treat a nest in a loft?
- What should I do before you arrive?
- Is the treatment safe around pets?
- Why has the problem returned?

#### Source-worthy evidence

- identification photographs;
- real treatment process;
- service-area and attendance information;
- proofing examples;
- seasonal advice based on real jobs.

#### Useful page idea

A visual guide to the most common pest enquiries received locally, with clear boundaries around when identification requires an inspection.

## 4. Accountant

### Useful customer questions

- What does a small limited company accountant normally do?
- What is included in the monthly fee?
- When should I move from sole trader to limited company?
- What records do you need from me?
- Can you take over from another accountant?

### Source-worthy evidence

- package comparison;
- onboarding process;
- deadline calendar;
- anonymised examples of common bookkeeping problems;
- explanation of what is and is not included.

### Useful page idea

A "first 30 days with a new accountant" guide that explains handover, authorisations, software access and what the client needs to provide.

## 5. Hairdresser or beauty salon

### Useful customer questions

- How long does the appointment take?
- Do I need a consultation first?
- What should I do before colouring?
- What happens if I have a skin sensitivity?
- Can I book two services together?

### Source-worthy evidence

- real treatment photographs with consent;
- stylist profiles;
- patch-test policy;
- service duration and price ranges;
- aftercare instructions.

### Useful page idea

A treatment comparison that explains who each option is suitable for, typical maintenance and realistic results rather than only showing polished photographs.

## 6. Independent garage

### Useful customer questions

- What is included in a service?
- Do you use manufacturer-equivalent parts?
- How do you diagnose an intermittent fault?
- Can I wait while the work is done?
- What happens before extra work is authorised?

### Source-worthy evidence

- diagnostic process;
- example fault investigation;
- labour or inspection charges;
- workshop photographs;
- policy on additional work.

### Useful page idea

A real diagnostic case study showing symptoms, tests, what was ruled out and why the final repair was recommended.

## 7. Restaurant or café

### Useful customer questions

- Do you cater for allergies?
- Is booking required?
- Is there parking nearby?
- Do you have vegetarian or vegan options?
- Can you accommodate a large group?

### Source-worthy evidence

- current menu;
- opening hours;
- genuine photographs;
- booking policy;
- accessibility information;
- clear allergy process.

### Useful page idea

A practical "planning your visit" page with parking, access, booking, group size and dietary information kept genuinely current.

## 8. Builder or landscaper

### Useful customer questions

- Do I need planning permission?
- How long will the work take?
- What is included in the quote?
- How are changes handled once work starts?
- Who removes waste?

### Source-worthy evidence

- project timelines;
- before/during/after photographs;
- quotation example;
- material choices;
- explanation of variations and unforeseen work.

### Useful page idea

A project diary showing the stages of one genuine job from survey to completion, including a problem encountered and how it was resolved.

## 9. Solicitor

### Useful customer questions

- Do I need a solicitor for this issue?
- How are fees calculated?
- What happens at the first appointment?
- What documents should I bring?
- How long might the process take?

### Source-worthy evidence

- clear practice areas;
- regulated credentials;
- fee models where appropriate;
- process guides reviewed by qualified staff;
- boundaries around general information and legal advice.

### Useful page idea

A plain-English process map for a common matter, clearly distinguishing general educational information from advice on an individual case.

## 10. Web designer or IT provider

### Useful customer questions

- Who owns my domain?
- Can you take over an existing website?
- What happens to SEO during a rebuild?
- Is hosting included?
- What if my current developer disappears?
- How do backups and updates work?

### Source-worthy evidence

- published packages or starting prices;
- migration checklist;
- before-and-after examples;
- hosting and backup policy;
- real performance improvements;
- explanation of ownership and access.

### Useful page idea

A website handover checklist that tells a business exactly which logins, ownership records and backups it should hold before changing supplier.

### What the examples have in common

Notice what is missing: none of these ideas begins with "write 2,000 words about the keyword".

They begin with real decisions. The most useful content often comes from a point where a customer is uncertain, a quote could change, a risk needs explaining or the business can show evidence that a generic publisher cannot.

# APPENDIX G - MONTHLY SEARCH VISIBILITY REVIEW

Use this once a month. Keep it short enough that it actually gets done.

## 1. Business outcomes

Enquiries this month: \_\_\_\_\_

Qualified enquiries: \_\_\_\_\_

Bookings/sales attributed to web activity where known: \_\_\_\_\_

Main services generating enquiries: \_\_\_\_\_

## 2. Classic search

Organic clicks: \_\_\_\_\_

Organic impressions: \_\_\_\_\_

Top gaining pages: \_\_\_\_\_

Top declining pages: \_\_\_\_\_

New useful queries noticed: \_\_\_\_\_

## 3. Local visibility

Google Business Profile calls/actions: \_\_\_\_\_

Review count change: \_\_\_\_\_

Average rating: \_\_\_\_\_

New review themes: \_\_\_\_\_

Opening hours/service details checked: Yes / No

#### 4. Generative and AI-assisted visibility

Google generative AI visibility reviewed: Yes / No / Not available

Bing AI Performance reviewed: Yes / No / Not available

AI citations or cited pages worth noting: \_\_\_\_\_

ChatGPT/AI referral sessions: \_\_\_\_\_

Branded search trend: Up / Flat / Down / Unsure

#### 5. Content and evidence

Customer questions heard repeatedly:

\_\_\_\_\_

Page updated this month: \_\_\_\_\_

New evidence added: \_\_\_\_\_

Outdated fact found and corrected: \_\_\_\_\_

#### 6. Next month's three priorities

1. \_\_\_\_\_
2. \_\_\_\_\_
3. \_\_\_\_\_

***Keep the review commercial.** A new citation is interesting. A better-qualified enquiry is useful. Do not let a new metric replace the reason the website exists.*

## ABOUT PIXELGEM

PixelGem helps small businesses build, improve and look after websites that have a job to do.

That can mean a new website, an existing-site rescue, hosting and care, local search improvements, technical fixes or simply helping a business explain itself more clearly online.

The approach behind this book is the same approach we bring to websites: practical foundations first, no scare tactics, no invented guarantees and no chasing fashionable tricks at the expense of the customer.

**Based in Cheshire, helping small businesses locally and across the UK.**  
[pixelgem.co.uk](http://pixelgem.co.uk)

# THE LAST PAGE

Search will keep changing.

The acronyms will keep changing.

What customers need from a good business is much more stable.

**Be easy to find.**

**Be easy to understand.**

**Give useful answers.**

**Show your evidence.**

**Keep your information accurate.**

**Make it obvious what somebody should do next.**

Do that well and you are not merely optimising for an algorithm.

You are building a better business website.

**PixelGem**

pixelgem.co.uk